

starting out with >>>

JAVA™

From Control Structures
through Data Structures

4TH EDITION



STARTING OUT WITH JAVA™

From Control Structures through Data Structures

FOURTH EDITION

Tony Gaddis

Haywood Community College

Godfrey Muganda

North Central College



330 Hudson Street, NY NY 10013

Senior Vice President Courseware Portfolio Management: Marcia J. Horton

Director, Portfolio Management: Engineering, Computer Science & Global Editions: Julian Partridge

Portfolio Manager: Matt Goldstein

Portfolio Management Assistant: Meghan Jacoby

Product Marketing Manager: Yvonne Vannatta

Field Marketing Manager: Demetrius Hall

Marketing Assistant: Jon Bryant

Managing Content Producer: Scott Disanno

Content Producer: Amanda Brands

Inventory Manager: Bruce Boundy

Manufacturing Buyer, Higher Ed, Lake Side Communications, Inc. (LSC): Maura Zaldivar-Garcia

Cover Designer: Joyce Wells

Cover Image: phloen/Alamy Stock Photo

Manager, Rights Management: Ben Ferrini

Full-Service Project Management: Sasibalan Chidambaram, SPi Global

Composition: SPi Global

Cover Printer: Phoenix Color/Hagerstown

Printer/Bindery: LSC Communications, Inc.

Copyright © 2019, 2016, 2012, 2007 by Pearson Education, Inc. or its affiliates. All Rights Reserved. Printed in the United States of America. This publication is protected by copyright, and permission should be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise. For information regarding permissions, request forms and the appropriate contacts within the Pearson Education Global Rights & Permissions department, please visit www.pearsoned.com/permissions/.

Many of the designations by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed in initial caps or all caps.

Unless otherwise indicated herein, any third-party trademarks that may appear in this work are the property of their respective owners and any references to third-party trademarks, logos or other trade dress are for demonstrative or descriptive purposes only. Such references are not intended to imply any sponsorship, endorsement, authorization, or promotion of Pearson's products by the owners of such marks, or any relationship between the owner and Pearson Education, Inc. or its affiliates, authors, licensees or distributors.

Microsoft and/or its respective suppliers make no representations about the suitability of the information contained in the documents and related graphics published as part of the services for any purpose. All such documents and related graphics are provided "as is" without warranty of any kind. Microsoft and/or its respective suppliers hereby disclaim all warranties and conditions with regard to this information, including all warranties and conditions of merchantability. Whether express, implied or statutory, fitness for a particular purpose, title and non-infringement. In no event shall microsoft and/or its respective suppliers be liable for any special, indirect or consequential damages or any damages whatsoever resulting from loss of use, data or profits, whether in an action of contract. Negligence or other tortious action, arising out of or in connection with the use or performance of information available from the services.

The documents and related graphics contained herein could include technical inaccuracies or typographical errors changes are periodically added to the information herein. Microsoft and/or its respective suppliers may make improvements and/or changes in the product(s) and/or the program(s) described herein at any time partial screen shots may be viewed in full within the software version specified.

Microsoft® Windows®, and Microsoft Office® are registered trademarks of the Microsoft Corporation in the U.S.A. and other countries. This book is not sponsored or endorsed by or affiliated with the Microsoft Corporation.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Library of Congress Cataloging-in-Publication Data

Names: Gaddis, Tony, author. | Muganda, Godfrey, author.

Title: Starting out with Java. From control structures through data structures / Tony Gaddis, Haywood Community College, Godfrey Muganda, North Central College.

Description: Fourth edition. | Pearson, [2019] | Includes bibliographical references and index.

Identifiers: LCCN 2018002219| ISBN 9780134787961 | ISBN 013478796X

Subjects: LCSH: Java (Computer program language) | Data structures (Computer science)

Classification: LCC QA76.73.J38 G33 2019 | DDC 005.13/3--dc23 LC record available at <https://lccn.loc.gov/2018002219>

1 18



ISBN-13: 978-0-13-478796-1

ISBN-10: 0-13-478796-X

Contents in Brief

1. [Chapter 1 Introduction to Computers and Java 1](#)
2. [Chapter 2 Java Fundamentals 27](#)
3. [Chapter 3 Decision Structures 111](#)
4. [Chapter 4 Loops and Files 189](#)
5. [Chapter 5 Methods 269](#)
6. [Chapter 6 A First Look at Classes 317](#)
7. [Chapter 7 Arrays and the `ArrayList` Class 403](#)
8. [Chapter 8 A Second Look at Classes and Objects 493](#)
9. [Chapter 9 Text Processing and More about Wrapper Classes 557](#)
10. [Chapter 10 Inheritance 611](#)
11. [Chapter 11 Exceptions and Advanced File I/O 701](#)
12. [Chapter 12 JavaFX: GUI Programming and Basic Controls 759](#)
13. [Chapter 13 JavaFX: Advanced Controls 823](#)
14. [Chapter 14 JavaFX: Graphics, Effects, and Media 909](#)
15. [Chapter 15 Recursion 999](#)
16. [Chapter 16 Sorting, Searching, and Algorithm Analysis 1027](#)
17. [Chapter 17 Generics 1079](#)
18. [Chapter 18 Collections and the Stream API 1125](#)
19. [Chapter 19 Linked Lists 1195](#)
20. [Chapter 20 Stacks and Queues 1245](#)
21. [Chapter 21 Binary Trees, AVL Trees, and Priority Queues 1287](#)
1. [Index 1353](#)

2. Appendices A–M Companion Website
3. Chapters 22–25 Companion Website
4. Case Studies 1–7 Companion Website

Contents

1. [Preface xxv](#)
1. [Chapter 1 Introduction to Computers and Java 1](#)
 1. [1.1 Introduction 1](#)
 2. [1.2 Why Program? 1](#)
 3. [1.3 Computer Systems: Hardware and Software 2](#)
 1. [Hardware 2](#)
 2. [Software 5](#)
 4. [1.4 Programming Languages 6](#)
 1. [What Is a Program? 6](#)
 2. [A History of Java 8](#)
 5. [1.5 What Is a Program Made Of? 8](#)
 1. [Language Elements 8](#)
 2. [Lines and Statements 11](#)
 3. [Variables 11](#)
 4. [The Compiler and the Java Virtual Machine 12](#)
 5. [Java Software Editions 13](#)
 6. [Compiling and Running a Java Program 14](#)
 6. [1.6 The Programming Process 16](#)
 1. [Software Engineering 18](#)
 7. [1.7 Object-Oriented Programming 19](#)
 1. [Review Questions and Exercises 21](#)
 2. [Programming Challenge 25](#)

2. [Chapter 2 Java Fundamentals 27](#)
 1. [2.1 The Parts of a Java Program 27](#)
 2. [2.2 The `print` and `println` Methods, and the Java API 33](#)
 3. [2.3 Variables and Literals 39](#)
 1. [Displaying Multiple Items with the `+` Operator 40](#)
 2. [Be Careful with Quotation Marks 41](#)
 3. [More about Literals 42](#)
 4. [Identifiers 42](#)
 5. [Class Names 44](#)
 4. [2.4 Primitive Data Types 44](#)
 1. [The Integer Data Types 46](#)
 2. [Floating-Point Data Types 47](#)
 3. [The `boolean` Data Type 50](#)
 4. [The `char` Data Type 50](#)
 5. [Variable Assignment and Initialization 52](#)
 6. [Variables Hold Only One Value at a Time 53](#)
 5. [2.5 Arithmetic Operators 54](#)
 1. [Integer Division 57](#)
 2. [Operator Precedence 57](#)
 3. [Grouping with Parentheses 59](#)
 4. [The `Math` Class 62](#)
 6. [2.6 Combined Assignment Operators 63](#)
 7. [2.7 Conversion between Primitive Data Types 65](#)
 1. [Mixed Integer Operations 67](#)

2. [Other Mixed Mathematical Expressions 68](#)
8. [2.8 Creating Named Constants with `final` 69](#)
9. [2.9 The `String` Class 70](#)
 1. [Objects Are Created from Classes 71](#)
 2. [The `String` Class 71](#)
 3. [Primitive Type Variables and Class Type Variables 71](#)
 4. [Creating a `String` Object 72](#)
10. [2.10 Scope 76](#)
11. [2.11 Comments 78](#)
12. [2.12 Programming Style 83](#)
13. [2.13 Reading Keyboard Input 85](#)
 1. [Reading a Character 89](#)
 2. [Mixing Calls to `nextLine` with Calls to Other `Scanner` Methods 89](#)
14. [2.14 Dialog Boxes 93](#)
 1. [Displaying Message Dialogs 93](#)
 2. [Displaying Input Dialogs 94](#)
 3. [An Example Program 94](#)
 4. [Converting String Input to Numbers 96](#)
15. [2.15 Common Errors to Avoid 99](#)
 1. [Review Questions and Exercises 100](#)
 2. [Programming Challenges 106](#)
3. [Chapter 3 Decision Structures 111](#)
 1. [3.1 The `if` Statement 111](#)
 1. [Using Relational Operators to Form Conditions 113](#)

2. [Putting It All Together 114](#)
 3. [Programming Style and the `if` Statement 117](#)
 4. [Be Careful with Semicolons 117](#)
 5. [Having Multiple Conditionally Executed Statements 118](#)
 6. [Flags 118](#)
 7. [Comparing Characters 119](#)
2. [3.2 The `if-else` Statement 121](#)
 3. [3.3 Nested `if` Statements 124](#)
 4. [3.4 The `if-else-if` Statement 128](#)
 5. [3.5 Logical Operators 134](#)
 1. [The Precedence of Logical Operators 139](#)
 2. [Checking Numeric Ranges with Logical Operators 140](#)
 6. [3.6 Comparing String Objects 142](#)
 1. [Ignoring Case in String Comparisons 146](#)
 7. [3.7 More about Variable Declaration and Scope 147](#)
 8. [3.8 The Conditional Operator \(Optional\) 149](#)
 9. [3.9 The `switch` Statement 150](#)
 10. [3.10 Displaying Formatted Output with `System.out.printf` and `String.format` 160](#)
 1. [Format Specifier Syntax 163](#)
 2. [Precision 164](#)
 3. [Specifying a Minimum Field Width 164](#)
 4. [Flags 167](#)
 5. [Formatting String Arguments 170](#)
 6. [The `String.format` Method 172](#)

11. [3.11 Common Errors to Avoid 174](#)
 1. [Review Questions and Exercises 175](#)
 2. [Programming Challenges 181](#)
4. [Chapter 4 Loops and Files 189](#)
 1. [4.1 The Increment and Decrement Operators 189](#)
 1. [The Difference between Postfix and Prefix Modes 192](#)
 2. [4.2 The `while` Loop 193](#)
 1. [The `while` Loop Is a Pretest Loop 196](#)
 2. [Infinite Loops 196](#)
 3. [Don't Forget the Braces with a Block of Statements 197](#)
 4. [Programming Style and the `while` Loop 198](#)
 3. [4.3 Using the `while` Loop for Input Validation 200](#)
 4. [4.4 The `do-while` Loop 204](#)
 5. [4.5 The `for` Loop 207](#)
 1. [The `for` Loop Is a Pretest Loop 210](#)
 2. [Avoid Modifying the Control Variable in the Body of the `for` Loop 211](#)
 3. [Other Forms of the Update Expression 211](#)
 4. [Declaring a Variable in the `for` Loop's Initialization Expression 211](#)
 5. [Creating a User Controlled `for` Loop 212](#)
 6. [Using Multiple Statements in the Initialization and Update Expressions 213](#)
 6. [4.6 Running Totals and Sentinel Values 216](#)
 1. [Using a Sentinel Value 219](#)
 7. [4.7 Nested Loops 221](#)
 8. [4.8 The `break` and `continue` Statements \(Optional\) 229](#)

9. [4.9 Deciding Which Loop to Use 229](#)
10. [4.10 Introduction to File Input and Output 230](#)
 1. [Using the `PrintWriter` Class to Write Data to a File 230](#)
 2. [Appending Data to a File 236](#)
 3. [Specifying the File Location 237](#)
 4. [Reading Data from a File 237](#)
 5. [Reading Lines from a File with the `nextLine` Method 238](#)
 6. [Adding a `throws` Clause to the Method Header 241](#)
 7. [Checking for a File's Existence 245](#)
11. [4.11 Generating Random Numbers with the `Random` Class 249](#)
12. [4.12 Common Errors to Avoid 255](#)
 1. [Review Questions and Exercises 256](#)
 2. [Programming Challenges 262](#)
5. [Chapter 5 Methods 269](#)
 1. [5.1 Introduction to Methods 269](#)
 1. [void Methods and Value-Returning Methods 270](#)
 2. [Defining a `void` Method 271](#)
 3. [Calling a Method 272](#)
 4. [Layered Method Calls 276](#)
 5. [Using Documentation Comments with Methods 277](#)
 2. [5.2 Passing Arguments to a Method 279](#)
 1. [Argument and Parameter Data Type Compatibility 281](#)
 2. [Parameter Variable Scope 282](#)
 3. [Passing Multiple Arguments 282](#)

4. [Arguments Are Passed by Value 284](#)
 5. [Passing Object References to a Method 285](#)
 6. [Using the @param Tag in Documentation Comments 288](#)
 3. [5.3 More about Local Variables 290](#)
 1. [Local Variable Lifetime 292](#)
 2. [Initializing Local Variables with Parameter Values 292](#)
 4. [5.4 Returning a Value from a Method 293](#)
 1. [Defining a Value-Returning Method 293](#)
 2. [Calling a Value-Returning Method 294](#)
 3. [Using the @return Tag in Documentation Comments 296](#)
 4. [Returning a boolean Value 300](#)
 5. [Returning a Reference to an Object 300](#)
 5. [5.5 Problem Solving with Methods 302](#)
 1. [Calling Methods That Throw Exceptions 305](#)
 6. [5.6 Common Errors to Avoid 305](#)
 1. [Review Questions and Exercises 306](#)
 2. [Programming Challenges 311](#)
6. [Chapter 6 A First Look at Classes 317](#)
 1. [6.1 Objects and Classes 317](#)
 1. [Classes: Where Objects Come From 318](#)
 2. [Classes in the Java API 319](#)
 3. [Primitive Variables vs. Objects 321](#)
 2. [6.2 Writing a Simple Class, Step by Step 324](#)
 1. [Accessor and Mutator Methods 338](#)

2. [The Importance of Data Hiding 338](#)
3. [Avoiding Stale Data 339](#)
4. [Showing Access Specification in UML Diagrams 339](#)
5. [Data Type and Parameter Notation in UML Diagrams 339](#)
6. [Layout of Class Members 340](#)
3. [6.3 Instance Fields and Methods 341](#)
4. [6.4 Constructors 346](#)
 1. [Showing Constructors in a UML Diagram 348](#)
 2. [Uninitialized Local Reference Variables 348](#)
 3. [The Default Constructor 348](#)
 4. [Writing Your Own No-Arg Constructor 349](#)
 5. [The `String` Class Constructor 350](#)
5. [6.5 Passing Objects as Arguments 358](#)
6. [6.6 Overloading Methods and Constructors 370](#)
 1. [The `BankAccount` Class 372](#)
 2. [Overloaded Methods Make Classes More Useful 378](#)
7. [6.7 Scope of Instance Fields 378](#)
 1. [Shadowing 379](#)
8. [6.8 Packages and `import` Statements 380](#)
 1. [Explicit and Wildcard `import` Statements 380](#)
 2. [The `java.lang` Package 381](#)
 3. [Other API Packages 381](#)
9. [6.9 Focus on Object-Oriented Design: Finding the Classes and Their Responsibilities 382](#)
 1. [Finding the Classes 382](#)

2. [Identifying a Class's Responsibilities 385](#)
 3. [This Is Only the Beginning 388](#)
10. [6.10 Common Errors to Avoid 388](#)
 1. [Review Questions and Exercises 389](#)
 2. [Programming Challenges 394](#)
7. [Chapter 7 Arrays and the `ArrayList` Class 403](#)
 1. [7.1 Introduction to Arrays 403](#)
 1. [Accessing Array Elements 405](#)
 2. [Inputting and Outputting Array Contents 406](#)
 3. [Java Performs Bounds Checking 409](#)
 4. [Watch Out for Off-by-One Errors 410](#)
 5. [Array Initialization 411](#)
 6. [Alternate Array Declaration Notation 412](#)
 2. [7.2 Processing Array Elements 413](#)
 1. [Array Length 415](#)
 2. [The Enhanced `for` Loop 416](#)
 3. [Letting the User Specify an Array's Size 417](#)
 4. [Reassigning Array Reference Variables 419](#)
 5. [Copying Arrays 420](#)
 3. [7.3 Passing Arrays as Arguments to Methods 422](#)
 4. [7.4 Some Useful Array Algorithms and Operations](#)
 1. [Comparing Arrays 426](#)
 2. [Summing the Values in a Numeric Array 427](#)
 3. [Getting the Average of the Values in a Numeric Array 428](#)

4. [Finding the Highest and Lowest Values in a Numeric Array 428](#)
5. [The `SalesData` Class 429](#)
6. [Partially Filled Arrays 437](#)
7. [Working with Arrays and Files 438](#)
5. [7.5 Returning Arrays from Methods 439](#)
6. [7.6 `String` Arrays 441](#)
 1. [Calling `String` Methods from an Array Element 443](#)
7. [7.7 Arrays of Objects 444](#)
8. [7.8 The Sequential Search Algorithm 447](#)
9. [7.9 Two-Dimensional Arrays 450](#)
 1. [Initializing a Two-Dimensional Array 454](#)
 2. [The `length` Field in a Two-Dimensional Array 455](#)
 3. [Displaying All the Elements of a Two-Dimensional Array 457](#)
 4. [Summing All the Elements of a Two-Dimensional Array 457](#)
 5. [Summing the Rows of a Two-Dimensional Array 458](#)
 6. [Summing the Columns of a Two-Dimensional Array 458](#)
 7. [Passing Two-Dimensional Arrays to Methods 459](#)
 8. [Ragged Arrays 461](#)
10. [7.10 Arrays with Three or More Dimensions 462](#)
11. [7.11 The Selection Sort and the Binary Search Algorithms 463](#)
 1. [The Selection Sort Algorithm 463](#)
 2. [The Binary Search Algorithm 466](#)
12. [7.12 Command-Line Arguments and Variable-Length Argument Lists 468](#)
 1. [Command-Line Arguments 469](#)

2. [Variable-Length Argument Lists 470](#)
13. [7.13 The `ArrayList` Class 472](#)
 1. [Creating and Using an `ArrayList` Object 473](#)
 2. [Using the Enhanced `for` Loop with an `ArrayList` 474](#)
 3. [The `ArrayList` Class's `toString` method 475](#)
 4. [Removing an Item from an `ArrayList` 476](#)
 5. [Inserting an Item 477](#)
 6. [Replacing an Item 478](#)
 7. [Capacity 479](#)
 8. [Storing Your Own Objects in an `ArrayList` 479](#)
 9. [Using the Diamond Operator for Type Inference 480](#)
14. [7.14 Common Errors to Avoid 481](#)
 1. [Review Questions and Exercises 481](#)
 2. [Programming Challenges 486](#)
8. [Chapter 8 A Second Look at Classes and Objects 493](#)
 1. [8.1 Static Class Members 493](#)
 1. [A Quick Review of Instance Fields and Instance Methods 493](#)
 2. [Static Members 494](#)
 3. [Static Fields 494](#)
 4. [Static Methods 497](#)
 2. [8.2 Passing Objects as Arguments to Methods 500](#)
 3. [8.3 Returning Objects from Methods 503](#)
 4. [8.4 The `toString` Method 505](#)
 5. [8.5 Writing an `equals` Method 509](#)

6. [8.6 Methods That Copy Objects 512](#)
 1. [Copy Constructors 514](#)
7. [8.7 Aggregation 515](#)
 1. [Aggregation in UML Diagrams 523](#)
 2. [Security Issues with Aggregate Classes 523](#)
 3. [Avoid Using `null` References 525](#)
8. [8.8 The `this` Reference Variable 528](#)
 1. [Using `this` to Overcome Shadowing 529](#)
 2. [Using `this` to Call an Overloaded Constructor from Another Constructor 530](#)
9. [8.9 Enumerated Types 531](#)
 1. [Enumerated Types Are Specialized Classes 532](#)
 2. [Switching On an Enumerated Type 538](#)
10. [8.10 Garbage Collection 540](#)
 1. [The `finalize` Method 542](#)
11. [8.11 Focus on Object-Oriented Design: Class Collaboration 542](#)
 1. [Determining Class Collaborations with CRC Cards 545](#)
12. [8.12 Common Errors to Avoid 546](#)
 1. [Review Questions and Exercises 547](#)
 2. [Programming Challenges 551](#)
9. [Chapter 9 Text Processing and More about Wrapper Classes 557](#)
 1. [9.1 Introduction to Wrapper Classes 557](#)
 2. [9.2 Character Testing and Conversion with the `Character` Class 558](#)
 1. [Character Case Conversion 563](#)
 3. [9.3 More `String` Methods 566](#)

1. [Searching for Substrings 566](#)
2. [Extracting Substrings 572](#)
3. [Methods That Return a Modified `String` 576](#)
4. [The Static `valueOf` Methods 577](#)

4. [9.4 The `StringBuilder` Class 579](#)

1. [The `StringBuilder` Constructors 580](#)
2. [Other `StringBuilder` Methods 581](#)
3. [The `toString` Method 584](#)

5. [9.5 Tokenizing Strings 589](#)

6. [9.6 Wrapper Classes for the Numeric Data Types 594](#)

1. [The Static `toString` Methods 594](#)
2. [The `toBinaryString`, `toHexString`, and `toOctalString` Methods 594](#)
3. [The `MIN\_VALUE` and `MAX\_VALUE` Constants 595](#)
4. [Autoboxing and Unboxing 595](#)

7. [9.7 Focus on Problem Solving: The `TestScoreReader` Class 597](#)

8. [9.8 Common Errors to Avoid 601](#)

1. [Review Questions and Exercises 601](#)
2. [Programming Challenges 605](#)

10. [Chapter 10 Inheritance 611](#)

1. [10.1 What Is Inheritance? 611](#)

1. [Generalization and Specialization 611](#)
2. [Inheritance and the “Is a” Relationship 612](#)
3. [Inheritance in UML Diagrams 620](#)
4. [The Superclass’s Constructor 621](#)

5. [Inheritance Does Not Work in Reverse 623](#)
2. [10.2 Calling the Superclass Constructor 624](#)
 1. [When the Superclass Has No Default or No-Arg Constructors 630](#)
 2. [Summary of Constructor Issues in Inheritance 631](#)
3. [10.3 Overriding Superclass Methods 632](#)
 1. [Overloading versus Overriding 637](#)
 2. [Preventing a Method from Being Overridden 640](#)
4. [10.4 Protected Members 641](#)
 1. [Package Access 646](#)
5. [10.5 Chains of Inheritance 647](#)
 1. [Class Hierarchies 653](#)
6. [10.6 The `Object` Class 653](#)
7. [10.7 Polymorphism 655](#)
 1. [Polymorphism and Dynamic Binding 656](#)
 2. [The “Is-a” Relationship Does Not Work in Reverse 658](#)
 3. [The `instanceof` Operator 659](#)
8. [10.8 Abstract Classes and Abstract Methods 660](#)
 1. [Abstract Classes in UML 666](#)
9. [10.9 Interfaces 667](#)
 1. [An Interface is a Contract 669](#)
 2. [Fields in Interfaces 673](#)
 3. [Implementing Multiple Interfaces 673](#)
 4. [Interfaces in UML 673](#)
 5. [Default Methods 674](#)

6. [Polymorphism and Interfaces 676](#)

10. [10.10 Anonymous Inner Classes 681](#)

11. [10.11 Functional Interfaces and Lambda Expressions 684](#)

12. [10.12 Common Errors to Avoid 689](#)

1. [Review Questions and Exercises 690](#)

2. [Programming Challenges 696](#)

11. [Chapter 11 Exceptions and Advanced File I/O 701](#)

1. [11.1 Handling Exceptions 701](#)

1. [Exception Classes 702](#)

2. [Handling an Exception 703](#)

3. [Retrieving the Default Error Message 707](#)

4. [Polymorphic References to Exceptions 710](#)

5. [Using Multiple `catch` Clauses to Handle Multiple Exceptions 710](#)

6. [The `finally` Clause 718](#)

7. [The Stack Trace 720](#)

8. [Handling Multiple Exceptions with One `catch` Clause 721](#)

9. [When an Exception Is Not Caught 723](#)

10. [Checked and Unchecked Exceptions 724](#)

2. [11.2 Throwing Exceptions 725](#)

1. [Creating Your Own Exception Classes 728](#)

2. [Using the `@exception` Tag in Documentation Comments 731](#)

3. [11.3 Advanced Topics: Binary Files, Random Access Files, and Object Serialization 732](#)

1. [Binary Files 732](#)

2. [Random Access Files 739](#)

3. [Object Serialization 744](#)
4. [Serializing Aggregate Objects 748](#)

4. [11.4 Common Errors to Avoid 749](#)

1. [Review Questions and Exercises 749](#)
2. [Programming Challenges 755](#)

12. [Chapter 12 JavaFX: GUI Programming and Basic Controls 759](#)

1. [12.1 Graphical User Interfaces 759](#)

1. [Event-Driven GUI Programs 761](#)

2. [12.2 Introduction to JavaFX 762](#)

1. [Controls 762](#)
2. [Stages and Scenes 763](#)
3. [The Application Class 763](#)

3. [12.3 Creating Scenes 765](#)

1. [Creating Controls 766](#)
2. [Creating Layout Containers 766](#)
3. [Creating a Scene Object 767](#)
4. [Adding the Scene Object to the Stage 768](#)
5. [Setting the Size of the Scene 770](#)
6. [Aligning Controls in an HBox Layout Container 770](#)

4. [12.4 Displaying Images 772](#)

1. [Loading Images from an Internet Location 775](#)
2. [Setting the Size of an Image 775](#)
3. [Preserving the Image's Aspect Ratio 775](#)
4. [Changing an ImageView's Image 776](#)

5. [12.5 More about the `HBox`, `VBox`, and `GridPane` Layout Containers 776](#)
 1. [The `HBox` Layout Container 777](#)
 2. [The `VBox` Layout Container 782](#)
 3. [The `GridPane` Layout Container 784](#)
 4. [Using Multiple Layout Containers in the Same Screen 791](#)
6. [12.6 `Button` Controls and Events 792](#)
 1. [Handling Events 794](#)
 2. [Writing Event Handlers 794](#)
 3. [Registering an Event Handler 795](#)
7. [12.7 Reading Input with `TextField` Controls 799](#)
8. [12.8 Using Anonymous Inner Classes and Lambda Expressions to Handle Events 803](#)
 1. [Using Anonymous Inner Classes to Create Event Handlers 803](#)
 2. [Using Lambda Expressions to Create Event Handlers 806](#)
9. [12.9 The `BorderPane` Layout Container 808](#)
10. [12.10 The `ObservableList` Interface 812](#)
11. [12.11 Common Errors to Avoid 814](#)
 1. [Review Questions and Exercises 814](#)
 2. [Programming Challenges 818](#)
13. [Chapter 13 JavaFX: Advanced Controls 823](#)
 1. [13.1 Styling JavaFX Applications with CSS 823](#)
 1. [Type Selector Names 824](#)
 2. [Style Properties 825](#)
 3. [Applying a Stylesheet to a JavaFX Application 826](#)
 4. [Applying Styles to the Root Node 830](#)

5. [Specifying Multiple Selectors in the Same Style Definition 832](#)
6. [Working with Colors 832](#)
7. [Creating a Custom Style Class Name 834](#)
8. [ID Selectors 836](#)
9. [Inline Style Rules 837](#)
2. [13.2 RadioButton Controls 838](#)
 1. [Determining in Code Whether a RadioButton Is Selected 839](#)
 2. [Selecting a RadioButton in Code 839](#)
 3. [Responding to RadioButton Clicks 844](#)
3. [13.3 CheckBox Controls 848](#)
 1. [Determining in Code Whether a CheckBox Is Selected 848](#)
 2. [Selecting a CheckBox in Code 849](#)
 3. [Responding to CheckBox Clicks 853](#)
4. [13.4 ListView Controls 853](#)
 1. [Retrieving the Selected Item 855](#)
 2. [Retrieving the Index of the Selected Item 856](#)
 3. [Responding to Item Selection with an Event Handler 859](#)
 4. [Adding Items versus Setting Items 860](#)
 5. [Initializing a ListView with an Array or an ArrayList 861](#)
 6. [Selection Modes 862](#)
 7. [Retrieving Multiple Selected Items 863](#)
 8. [Working With the Elements of an ObservableList 866](#)
 9. [Converting an ObservableList to an Array 867](#)
 10. [Using Code to Select an Item in a ListView 868](#)

11. [List View Orientation 868](#)

12. [Creating List Views of Objects Other Than String 869](#)

5. [13.5 ComboBox Controls 874](#)

1. [Retrieving the Selected Item 875](#)

2. [Responding to ComboBox Item Selection with an Event Handler 877](#)

3. [Editable ComboBoxes 879](#)

6. [13.6 Slider Controls 880](#)

7. [13.7 TextArea Controls 885](#)

8. [13.8 Menus 887](#)

1. [Assigning Mnemonics to Menu Items 895](#)

9. [13.9 The FileChooser Class 897](#)

1. [Displaying a FileChooser Dialog Box 898](#)

10. [13.10 Using Console Output to Debug a GUI Application 898](#)

11. [13.11 Common Errors to Avoid 902](#)

1. [Review Questions and Exercises 902](#)

2. [Programming Challenges 906](#)

14. [Chapter 14 JavaFX: Graphics, Effects, and Media 909](#)

1. [14.1 Drawing Shapes 909](#)

1. [The Screen Coordinate System 909](#)

2. [The Shape Class and Its Subclasses 910](#)

3. [The Line Class 911](#)

4. [Changing the Stroke Color 914](#)

5. [The Circle Class 914](#)

6. [The Rectangle Class 918](#)

7. [The `Ellipse` Class 921](#)
8. [The `Arc` Class 924](#)
9. [The `Polygon` Class 928](#)
10. [The `Polyline` Class 931](#)
11. [The `Text` Class 933](#)
12. [Rotating Nodes 936](#)
13. [Scaling Nodes 938](#)
2. [14.2 Animation 940](#)
 1. [The `TranslateTransition` Class 941](#)
 2. [The `RotateTransition` Class 944](#)
 3. [The `ScaleTransition` Class 949](#)
 4. [The `StrokeTransition` Class 952](#)
 5. [The `FillTransition` Class 953](#)
 6. [The `FadeTransition` Class 954](#)
 7. [Controlling the Animation 956](#)
 8. [Specifying an Interpolator 956](#)
3. [14.3 Effects 958](#)
 1. [The `DropShadow` Class 958](#)
 2. [The `InnerShadow` Class 961](#)
 3. [The `ColorAdjust` Class 962](#)
 4. [The `BoxBlur`, `GaussianBlur`, and `MotionBlur` Classes 962](#)
 5. [The `SepiaTone` Class 965](#)
 6. [The `Glow` Class 966](#)
 7. [The `Reflection` Class 967](#)

8. [Combining Effects 967](#)

4. [14.4 Playing Sound Files 969](#)

1. [Registering an `EndOfMedia` Event Handler 971](#)

5. [14.5 Playing Videos 974](#)

6. [14.6 Handling Key Events 979](#)

1. [Using an Anonymous Inner Class to Register a Key Event Handler to the Scene 980](#)

2. [Using a Lambda Expression to Register a Key Event Handler to the Scene 981](#)

7. [14.7 Handling Mouse Events 986](#)

8. [14.8 Common Errors to Avoid 992](#)

1. [Review Questions and Exercises 992](#)

2. [Programming Challenges 995](#)

15. [Chapter 15 Recursion 999](#)

1. [15.1 Introduction to Recursion 999](#)

2. [15.2 Solving Problems with Recursion 1002](#)

1. [Direct and Indirect Recursion 1006](#)

3. [15.3 Examples of Recursive Methods 1007](#)

1. [Summing a Range of Array Elements with Recursion 1007](#)

2. [Drawing Concentric Circles 1008](#)

3. [The Fibonacci Series 1011](#)

4. [Finding the Greatest Common Divisor 1012](#)

4. [15.4 A Recursive Binary Search Method 1014](#)

5. [15.5 The Towers of Hanoi 1017](#)

6. [15.6 Common Errors to Avoid 1021](#)

1. [Review Questions and Exercises 1022](#)

2. [Programming Challenges 1025](#)
16. [Chapter 16 Sorting, Searching, and Algorithm Analysis 1027](#)
 1. [16.1 Introduction to Sorting Algorithms 1027](#)
 1. [The Bubble Sort 1027](#)
 2. [Using the Bubble Sort to Sort Objects 1033](#)
 3. [The Selection Sort 1036](#)
 4. [Using the Selection Sort to Sort Objects 1040](#)
 5. [The Insertion Sort 1042](#)
 6. [Using the Insertion Sort to Sort Objects 1047](#)
 7. [The Quicksort Algorithm 1048](#)
 8. [Using Quicksort to Sort Objects 1052](#)
 2. [16.2 Introduction to Search Algorithms 1052](#)
 1. [The Sequential Search Algorithm 1052](#)
 2. [Average Case Analysis of the Sequential Search 1055](#)
 3. [The Binary Search 1055](#)
 4. [Efficiency of the Binary Search 1059](#)
 5. [A Recursive Binary Search 1059](#)
 3. [16.3 Analysis of Algorithms 1063](#)
 1. [Computational Problems and Basic Steps 1064](#)
 2. [Complexity of Algorithms 1065](#)
 3. [Worst Case Complexity of Algorithms 1066](#)
 4. [Average Case Complexity 1067](#)
 5. [Asymptotic Complexity and the Big O Notation 1068](#)
 4. [16.4 Common Errors to Avoid 1072](#)

1. [Review Questions and Exercises 1073](#)
2. [Programming Challenges 1076](#)
17. [Chapter 17 Generics 1079](#)
 1. [17.1 Introduction to Generics 1079](#)
 1. [Using the Diamond Operator for Type Inference 1082](#)
 2. [17.2 Writing a Generic Class 1082](#)
 1. [Only Reference Types Can Be Passed to Type Parameters 1087](#)
 2. [Instantiating a Generic Class without Specifying a Type Argument 1088](#)
 3. [Commonly Used Type Parameter Names 1090](#)
 3. [17.3 Passing Objects of a Generic Class to a Method 1090](#)
 1. [Constraining a Type Parameter 1092](#)
 2. [Defining a Type Parameter in a Method Header 1094](#)
 3. [The `extends` Key Word Constrains a Type to an Upper Bound 1095](#)
 4. [The `super` Key Word Constrains a Type to a Lower Bound 1095](#)
 4. [17.4 Writing Generic Methods 1096](#)
 5. [17.5 Constraining a Type Parameter in a Generic Class 1097](#)
 6. [17.6 Inheritance and Generic Classes 1100](#)
 7. [17.7 Defining Multiple Type Parameters 1104](#)
 8. [17.8 Generics and Interfaces 1107](#)
 1. [Constraining a Type Parameter to a Type That Implements an Interface 1108](#)
 9. [17.9 Erasure 1113](#)
 10. [17.10 Restrictions on the Use of Generic Types 1116](#)
 11. [17.11 Common Errors to Avoid 1118](#)
 1. [Review Questions and Exercises 1119](#)

2. [Programming Challenges 1122](#)
18. [Chapter 18 Collections and the Stream API 1125](#)
 1. [18.1 Introduction to the Java Collections Framework 1125](#)
 1. [Lists, Sets, and Maps 1125](#)
 2. [Java Functional Interfaces 1126](#)
 3. [JCF Interfaces 1126](#)
 4. [Iterating a Collection 1127](#)
 5. [The `Iterable` Interface and the `forEach` method 1127](#)
 6. [Iterators 1128](#)
 7. [The Enhanced `For` Loop 1131](#)
 8. [The `Collection` Interface 1131](#)
 2. [18.2 Lists 1133](#)
 1. [The `List` Interface 1134](#)
 2. [The `ArrayList` and `LinkedList` Classes 1135](#)
 3. [Creating and Initializing Lists 1136](#)
 4. [List Iterators 1137](#)
 5. [Methods of the `LinkedList` Class 1140](#)
 6. [Using an Interface Variable to Reference a `Collection` Object 1141](#)
 3. [18.3 Sets 1143](#)
 1. [The `Set` Interface 1143](#)
 2. [The `HashSet` Class 1143](#)
 3. [Implementing Your Own `hashCode` Method 1148](#)
 4. [The `LinkedHashSet` Class 1152](#)
 5. [The `SortedSet` Interface and the `TreeSet` Class 1153](#)

6. [Using a Comparator to Compare Objects 1156](#)

4. [18.4 Maps 1158](#)

1. [The `HashMap` Class 1160](#)

2. [The `LinkedHashMap` Class 1166](#)

3. [The `SortedMap` Interface and the `TreeMap` Class 1168](#)

5. [18.5 The `Collections` Class 1169](#)

6. [18.6 The Stream API 1171](#)

1. [The Stream Package 1172](#)

2. [Creating a Stream 1172](#)

3. [Terminal Operations 1173](#)

4. [Intermediate Operations on Streams 1177](#)

5. [Intermediate Stream Operations are Lazy 1179](#)

6. [The Stream `map` Operation 1181](#)

7. [The Stream `reduce` Methods 1182](#)

8. [Stream Collectors 1184](#)

7. [18.7 Common Errors to Avoid 1189](#)

1. [Review Questions and Exercises 1190](#)

2. [Programming Challenges 1193](#)

19. [Chapter 19 Linked Lists 1195](#)

1. [19.1 Introduction to Linked Lists 1195](#)

1. [Creating Lists and Adding Nodes 1196](#)

2. [Removing a Node 1198](#)

3. [Traversing Linked Lists 1199](#)

2. [19.2 Operations on Linked Lists 1202](#)

1. [The isEmpty Method 1202](#)
2. [The size Method 1203](#)
3. [The add Methods 1203](#)
4. [The remove Methods 1204](#)
5. [The toString Method 1204](#)
6. [A Graphical Interface for the Linked List Program 1210](#)

3. [19.3 Doubly Linked and Circularly Linked Lists 1215](#)

1. [Adding a Node to a Doubly Linked List 1216](#)
2. [Removing a Node from a Doubly Linked List 1217](#)
3. [Circularly Linked Lists 1223](#)

4. [19.4 Recursion on Linked Lists 1224](#)

1. [A Recursive Implementation of size 1224](#)
2. [A Recursive Implementation of the add\(String e, Node list\) Method 1225](#)
3. [A Recursive Implementation of add\(int index, String e, Node list\) 1226](#)
4. [A Recursive Implementation of the remove Methods 1226](#)

5. [19.5 Common Errors to Avoid 1238](#)

1. [Review Questions and Exercises 1238](#)
2. [Programming Challenges 1241](#)

20. [Chapter 20 Stacks and Queues 1245](#)

1. [20.1 Stacks and Their Applications 1245](#)

1. [Examples and Applications of Stacks 1245](#)
2. [Stack Operations 1246](#)
3. [Stacks of Primitive Types 1247](#)

2. [20.2 Array Implementation of Stacks 1248](#)

1. [The Stack Push Operation 1249](#)
2. [The Stack_{empty} Method 1249](#)
3. [The Stack_{peek} and _{pop} Methods 1249](#)
4. [Stacks of Objects 1253](#)
3. [20.3 Linked Implementation of Stacks 1254](#)
 1. [Implementation of Stack Methods 1255](#)
4. [20.4 Queues and Their Applications 1259](#)
5. [20.5 Array Implementation of Queues 1259](#)
 1. [GUI Front End for the Queue Demo Program 1266](#)
6. [20.6 Linked List Implementation of Queues 1270](#)
 1. [Queue Initialization and Enqueuing of Items 1270](#)
 2. [Dequeuing Items 1271](#)
7. [20.7 Generic Implementation of Stacks and Queues 1275](#)
8. [20.8 Queues and Breadth-First Search 1278](#)
9. [20.9 Common Errors to Avoid 1281](#)
 1. [Review Questions and Exercises 1282](#)
 2. [Programming Challenges 1284](#)
21. [Chapter 21 Binary Trees, AVL Trees, and Priority Queues 1287](#)
 1. [21.1 Binary Trees and Their Applications 1287](#)
 1. [Binary Tree Concepts 1288](#)
 2. [Applications of Binary Trees 1289](#)
 3. [Representing Nodes of Binary Trees 1289](#)
 4. [Traversing a Binary Tree 1290](#)
 5. [Graphical Display of Binary Trees 1293](#)

6. [Class Implementation of Binary Trees 1297](#)
2. [21.2 Binary Search Trees 1297](#)
 1. [Adding a Value to a Binary Search Tree 1299](#)
 2. [Removing a Value from a Binary Search Tree 1299](#)
 3. [Implementation of Binary Search Trees 1303](#)
 4. [Graphical Display of Binary Search Trees 1308](#)
 5. [A Program for Testing the Binary Search Tree Class 1310](#)
3. [21.3 AVL Trees 1315](#)
 1. [Adding New Elements to AVL Trees 1316](#)
 2. [Implementation of AVL Trees 1319](#)
4. [21.4 Priority Queues 1328](#)
 1. [Applications of Priority Queues 1330](#)
 2. [Using Comparators with Priority Queues 1331](#)
 3. [Analysis of Heapsort 1333](#)
 4. [Implementation of Priority Queues 1333](#)
 5. [Binary Trees with the Heap Order Property 1333](#)
 6. [Complete Binary Trees and Heaps 1334](#)
 7. [The Depth of Complete Binary Trees 1335](#)
 8. [Storing a Complete Binary Tree in an Array 1335](#)
 9. [Adding an Item to a Heap 1336](#)
 10. [Removing the Minimum Element 1337](#)
5. [21.5 Common Errors to Avoid 1345](#)
 1. [Review Questions and Exercises 1345](#)
 2. [Programming Challenges 1349](#)

1. [Index 1353](#)
2. The following appendices, online chapters, and online case studies are available on the book's online resource page at www.pearson.com/gaddis.
3. **Online Appendices**
 1. [Appendix A: The ASCII/Unicode Characters](#)
 2. [Appendix B: Operator Precedence and Associativity](#)
 3. [Appendix C: Java Key Words](#)
 4. [Appendix D: Installing the JDK and Using the JDK Documentation](#)
 5. [Appendix E: Using the javadoc Utility](#)
 6. [Appendix F: More about the Math Class](#)
 7. [Appendix G: Packages](#)
 8. [Appendix H: Working with Records and Random Access Files](#)
 9. [Appendix I: Configuring JavaDB and Installing Apache Derby](#)
 10. [Appendix J: The QuickSort Algorithm](#)
 11. [Appendix K: Named Colors](#)
 12. [Appendix L: Answers to Checkpoints](#)
 13. [Appendix M: Answers to Odd-Numbered Review Questions](#)
14. **Online Chapters**
 15. [Chapter 22: Databases](#)
 16. [Chapter 23: A First Look at GUI Applications with Swing](#)
 17. [Chapter 24: Advanced Swing GUI Applications](#)
 18. [Chapter 25: Applets and More](#)
19. **Online Case Studies**
 20. [Case Study 1: Calculating Sales Commission](#)
 21. [Case Study 2: The Amortization Class](#)

- 22. [Case Study 3: The `PinTester` Class](#)
- 23. [Case Study 4: `Parallel Arrays`](#)
- 24. [Case Study 5: The `FeetInches` Class](#)
- 25. [Case Study 6: The `SerialNumber` Class](#)
- 26. [Case Study 7: A Simple Text Editor Application](#)

LOCATION OF VIDEONOTES IN THE TEXT

Compiling and Running a Java Program, p. [14](#)

[Chapter 1](#) Using an IDE, p. [14](#)

Your First Java Program, p. [25](#)

Displaying Console Output, p. [33](#)

Declaring Variables, p. [39](#)

[Chapter 2](#)

Simple Math Expressions, p. [55](#)

The Miles-per-Gallon Problem, p. [107](#)

The if Statement, p. [111](#)

The if-else Statement, p. [120](#)

[Chapter 3](#)

The if-else-if Statement, p. [129](#)

The Time Calculator Problem, p. [182](#)

The while Loop, p. [193](#)

[Chapter 4](#)

The Pennies for Pay Problem, p. [263](#)

Passing Arguments to a Method, p. [279](#)

[Chapter 5](#) Returning a Value from a Method, p. [293](#)

The Retail Price Calculator Problem, p. [311](#)

Writing Classes and Creating Objects, p. [325](#)

[Chapter 6](#) Initializing an Object with a Constructor, p. [346](#)

The Personal Information Class Problem, p. [395](#)

Accessing Array Elements in a Loop, p. [407](#)

[Chapter 7](#) Passing an Array to a Method, p. [422](#)

The Charge Account Validation Problem, p. [487](#)

Returning Objects from Methods, p. [503](#)

[Chapter 8](#) Aggregation, p. [515](#)

The BankAccount, Class Copy Constructor Problem, p. [552](#)

[Chapter 9](#) The Sentence Capitalizer Problem, p. [605](#)

Inheritance, p. [611](#)

[Chapter 10](#) Polymorphism, p. [655](#)

The Employee and Productionworker Classes Problem, p. [696](#)

Handling Exceptions, p. [701](#)

[Chapter 11](#)

The Exception Project Problem, p. [757](#)

Introduction to JavaFX, p. [762](#)

Creating Scenes, p. [765](#)

Displaying Images, p. [772](#)

The HBox Layout Container, p. [777](#)

The VBox Layout Container, p. [782](#)

[Chapter 12](#) The GridPane Layout Container, p. [784](#)

Button Controls and Events, p. [792](#)

The TextField Control, p. [799](#)

Using Anonymous Inner Classes as Event Handlers, p. [803](#)

Using Lambda Expressions as Event Handlers, p. [806](#)

The Latin Translator Problem, p. [818](#)

JavaFX and CSS, p. [823](#)

RadioButton Controls, p. [838](#)

CheckBox Controls, p. [848](#)

[Chapter 13](#) ListView Controls, p. [853](#)

ComboBox Controls, p. [874](#)

Slider Controls, p. [880](#)

The Dorm and Meal Plan Calculator Problem, p. [906](#)

Drawing Shapes with JavaFX, p. [909](#)

JavaFX Animation, p. [940](#)

JavaFX Effects, p. [958](#)

Playing Sound Files with JavaFX, p. [970](#)

[Chapter 14](#)

Playing Videos with JavaFX, p. [974](#)

Handling Key Events in JavaFX, p. [979](#)

Handling Mouse Events in JavaFX, p. [986](#)

The Coin Toss Problem, p. [996](#)

Reducing a Problem with Recursion, p. [1003](#)

[Chapter 15](#) The Recursive Power Problem, p. [1026](#)

The Bubble Sort Algorithm, p. [1027](#)

[Chapter 16](#) Searching for Objects with the Binary Search Algorithm, p. [1076](#)

Inheritance and Generic classes, p. [1100](#)

[Chapter 17](#) Highest and Lowest Elements, p. [1123](#)

[Chapter 18](#) The HashMap Class, p. [1160](#)

Adding a Node to a Doubly-Linked List, p. [1216](#)

[Chapter 19](#) Recursive Linked List Operations, p. [1242](#)

Array Implementation of Queues, p. [1259](#)

[Chapter 20](#) Array-Based Deque, p. [1285](#)

Removing a Value from a Binary Search Tree, p. [1299](#)

[Chapter 21](#) Prefix Calculator, p. [1350](#)

Preface

Welcome to *Starting Out with Java: From Control Structures through Data Structures, Fourth Edition*. This book is intended for a traditional two-semester CS1/CS2 sequence of courses. The first half of the book, intended for a CS1 class, teaches fundamental programming and problem-solving concepts using Java. The second half of the book, intended for a CS2 class, teaches advanced concepts and provides an introduction to algorithms and data structures. The book is written for students with no prior programming background, but experienced students will also benefit from its depth of detail.

Control Structures First, Then Objects, Then Data Structures

The text introduces students to the fundamentals of data types, input/output, control structures, methods, and objects created from standard library classes. Next, students learn to use arrays. Then, a series of more advanced topics are presented, including inheritance, polymorphism, GUI development, recursion, searching, sorting, algorithm analysis, and generics. Finally, students are introduced to the world of data structures, beginning with an overview of the Java Collections Framework (JCF), and continuing through a series of chapters in which students learn to develop their own data structures, including linked lists, stacks, queues, priority queues, binary trees, and AVL trees. Students are shown how to implement data structures either with or without generics, which gives instructors a great deal of flexibility in presenting the material.

As with all the books in the *Starting Out With* series, the hallmark of the text is its clear, friendly, and easy-to-understand writing. In addition, it is rich in concise, practical example programs.

Changes in This Edition

This book's pedagogy, organization, and clear writing style remain the same as in the previous edition. The most significant change in this edition is the switch from Swing to JavaFX in the chapters that focus on GUI development. Although Swing is not officially deprecated, Oracle has announced that JavaFX has replaced Swing as the standard GUI library for Java¹.

¹ <http://www.oracle.com/technetwork/java/javafx/overview/faq-1446554.html#6>

In this edition, we have added the following new chapters:

- **[Chapter 12](#) JavaFX: GUI Programming and Basic Controls:** This chapter presents the basics of developing graphical user interface (GUI) applications with JavaFX. Fundamental controls,

layout containers, and the basic concepts of event-driven programming are covered.

- **[Chapter 13 JavaFX: Advanced Controls](#):** This chapter discusses CSS styling and advanced user interface controls.
- **[Chapter 14 JavaFX: Graphics, Effects, and Media](#):** This chapter discusses 2D shapes, animation, visual effects, playing audio and video, and responding to mouse and keyboard events.

The Swing and Applet material that appeared in the previous edition is still available on the book's companion Web site, as the following online chapters:

- The previous [Chapter 12](#) *A First Look At GUI Applications* is now available online as [Chapter 23](#).
- The previous [Chapter 13](#) *Advanced GUI Applications* is now available online as [Chapter 24](#).
- The previous [Chapter 14](#) *Applets and More* is now available online as [Chapter 25](#).

In addition to the new JavaFX chapters, the online Database chapter, which is now [Chapter 22](#), has been updated to use JavaFX instead of Swing for its GUI applications. We have also added several new, motivational programming problems throughout the book.

Organization of the Text

The text teaches Java step by step. Each chapter covers a major set of topics and builds knowledge as students progress through the book. Although the chapters can be easily taught in their existing sequence, there is some flexibility. [Figure P-1](#) shows chapter dependencies. Each box represents a chapter or a group of chapters. An arrow points from a chapter to the chapter that must be previously covered.




Figure P-1 Chapter dependencies

[Figure P-1 Full Alternative Text](#)

Brief Overview of Each Chapter

Chapter 1: Introduction to Computers and Java.

This chapter provides an introduction to the field of computer science, and covers the fundamentals of hardware, software, and programming languages. The elements of a program, such as key words, variables, operators, and punctuation are discussed by examining a simple program. An overview of entering source code, compiling, and executing a program is presented. A brief history of Java is also given.

Chapter 2: Java Fundamentals.

This chapter gets students started in Java by introducing data types, identifiers, variable declarations, constants, comments, program output, and simple arithmetic operations. The conventions of programming style are also introduced. Students learn to read console input with the `Scanner` class and with dialog boxes using `JOptionPane`.

Chapter 3: Decision Structures.

In this chapter students explore relational operators and relational expressions, and are shown how to control the flow of a program with the `if`, `if-else`, and `if-else-if` statements. Nested `if` statements, logical operators, the conditional operator, and the `switch` statement are also covered. The chapter discusses how to compare `String` objects with the `equals`, `compareTo`, `equalsIgnoreCase`, and `compareToIgnoreCase` methods. Formatting numeric output with the `System.out.printf` method and the `String.format` method is discussed.

Chapter 4: Loops and Files.

This chapter covers Java's repetition control structures. The `while` loop, `do-while` loop, and `for` loop are taught, along with common uses for these devices. Counters, accumulators, running totals, sentinels, and other application-related topics are discussed. Simple file operations for reading and writing text files are included.

Chapter 5: Methods.

In this chapter students learn how to write void methods, value-returning methods, and methods that do and do not accept arguments. The concept of functional decomposition is discussed.

Chapter 6: A First Look at Classes.

This chapter introduces students to designing classes for the purpose of instantiating objects. Students learn about class fields and methods, and UML diagrams are introduced as a design tool. Then constructors and overloading are discussed. A `BankAccount` class is presented as a case study, and a section on object-oriented design is included. This section leads the students through the process of identifying classes and their responsibilities within a problem domain. There is also a section that briefly explains packages and the `import` statement.

Chapter 7: Arrays and the ArrayList Class.

In this chapter students learn to create and work with single and multi-dimensional arrays. Numerous array-processing techniques are demonstrated, such as summing the elements in an array, finding the highest and lowest values, and sequentially searching an array are also discussed. Other topics, including ragged arrays and variable-length arguments (varargs) are also discussed. The `ArrayList` class is introduced and Java's generic types are briefly discussed and demonstrated.

Chapter 8: A Second Look at Classes and Objects.

This chapter shows students how to write classes with added capabilities. Static methods and fields, interaction between objects, passing objects as arguments, and returning objects from methods are discussed. Aggregation and the “has a” relationship is covered, as well as enumerated types. A section on object-oriented design shows how to use CRC cards to determine the collaborations among classes.

Chapter 9: Text Processing and More about Wrapper Classes.

This chapter discusses the numeric and `Character` wrapper classes. Methods for converting numbers to strings, testing the case of characters, and converting the case of characters are covered. Autoboxing and unboxing are also discussed. More `String` class methods are covered, including using the `split` method to tokenize strings. The chapter also covers the `StringBuilder` class.

Chapter 10: Inheritance.

The study of classes continues in this chapter with the subjects of inheritance and polymorphism. The topics covered include superclasses, subclasses, how constructors work in inheritance, method overriding, polymorphism and dynamic binding, protected and package access, class hierarchies, abstract classes, abstract methods, anonymous inner classes, interfaces, and lambda expressions.

Chapter 11: Exceptions and Advanced File I/O.

In this chapter students learn to develop enhanced error trapping techniques using exceptions. Handling exceptions is covered, as well as developing and throwing custom exceptions. The chapter discusses advanced techniques for working with sequential access, random access, text, and binary files.

Chapter 12: JavaFX: GUI Programming and Basic Controls.

This chapter presents the basics of developing graphical user interface (GUI) applications with JavaFX. Fundamental controls, layout containers, and the basic concepts of event-driven programming are covered.

Chapter 13: JavaFX: Advanced Controls.

This chapter discusses CSS styling and advanced user interface controls, such as `RadioButtons`, `CheckBoxes`, `ListViews`, `ComboBoxes`, `Sliders`, and `TextAreas`. Menu systems and `FileChoosers` are also covered.

Chapter 14: JavaFX: Graphics, Effects, and Media.

This chapter discusses 2D shapes, animation, visual effects, playing audio and video, and responding to mouse and keyboard events.

Chapter 15: Recursion.

This chapter presents recursion as a problem-solving technique. Numerous examples of recursive methods are demonstrated.

Chapter 16: Sorting, Searching, and Algorithm Analysis.

In this chapter students learn the basics of sorting arrays and searching for data stored in them. The chapter covers the bubble sort, selection sort, insertion sort, Quicksort, sequential search, and binary search algorithms. A section on algorithm analysis is also provided, which covers basic steps, complexity, and Big O notation.

Chapter 17: Generics.

This chapter shows students how to write generic classes and methods. Topics include erasure, passing instances of a generic class as arguments to a method, constraining type parameters, generics and inheritance, and generics and interfaces. Restrictions on the use of generic types are also discussed.

Chapter 18: Collections and the Stream API.

This chapter introduces students to the Java Collections Framework (JCF). Lists, sets, maps, and the Collections class are discussed. The chapter concludes with a discussion of the classes and interfaces in the `java.util.stream` package.

Chapter 19: Linked Lists.

This chapter introduces concepts and techniques for writing a linked list class. Basic linked list operations, such as adding, inserting, removing, and traversing, are covered. Doubly-linked lists, circularly-linked lists, and using recursion with linked lists are discussed.

Chapter 20: Stacks and Queues.

In this chapter students learn about the operations of stacks and queues, and how to write array-based and linked list-based stack and queue classes.

Chapter 21: Binary Trees, AVL Trees, and Priority Queues.

This chapter covers various aspects of binary trees and binary tree operations. The specific binary tree implementations covered are binary search trees, AVL trees, and priority queues.

Chapter 22 (Online): Databases.

This chapter is available on the book's companion Web site (at www.pearson.com/gaddis). It introduces the students to database programming. The basic concepts of database management systems and SQL are first introduced. Then the students learn to use JDBC to write database applications in Java. Relational data is covered, and numerous example programs are presented throughout the chapter.

Chapters 23 – 25 (Online): Legacy Swing GUI Chapters.

The previous edition's chapters on Swing GUI and Applet programming are now available on the book's companion website (at www.pearson.com/gaddis) as Chapters 23 through 25.

Features of the Text

Concept Statements.

Each major section of the text starts with a concept statement that concisely summarizes the focus of the section.

Example Programs.

The text has an abundant number of complete and partial example programs, each designed to highlight the current topic. In most cases the programs are practical, real-world examples.

Program Output.

Each example program is followed by a sample of its output, which shows students how the program functions.

Checkpoints.

Checkpoints, highlighted by the checkmark icon, appear at intervals throughout each chapter. They are

designed to check students' knowledge soon after learning a new topic. Answers for all Checkpoint questions can be downloaded from the book's companion Web site at www.pearson.com/gaddis.

NOTE:

Notes appear at several places throughout the text. They are short explanations of interesting or often misunderstood points relevant to the topic at hand.

TIP:

Tips advise the student on the best techniques for approaching different programming problems and appear regularly throughout the text.

WARNING!

Warnings caution students about certain Java features, programming techniques, or practices that can lead to malfunctioning programs or lost data.

In the Spotlight.

Many of the chapters provide an *In the Spotlight* section that presents a programming problem, along with detailed, step-by-step analysis showing the student how to solve it.

VideoNotes.

A series of online videos, developed specifically for this book, are available for viewing at www.pearson.com/gaddis. Icons appear throughout the text, alerting the student to videos about specific topics.

Case Studies.

Case studies that simulate real-world business applications are introduced throughout the text and are available for download from the Gaddis resource page at www.pearson.com/gaddis.

Common Errors to Avoid.

Most chapters provide a list of common errors and explanations of how to avoid them.

Review Questions and Exercises.

Each chapter presents a thorough and diverse set of review questions and exercises. They include Multiple Choice and True/False, Find the Error, Algorithm Workbench, and Short Answer.

Programming Challenges.

Each chapter offers a pool of programming challenges designed to solidify students' knowledge of topics at hand. In most cases the assignments present real-world problems to be solved.

Supplements

Student Online Resources

Many student resources are available for this book from the publisher. The following items are available on the Gaddis Series resource page at www.pearson.com/gaddis:

- The source code for each example program in the book
- Access to the book's companion VideoNotes
- Appendixes A–M (listed in the Contents)
- A collection of seven valuable Case Studies (listed in the Contents)
- Links to download the Java™ Edition Development Kit
- Links to download numerous programming environments, including jGRASP™, Eclipse™, TextPad™, NetBeans™, JCreator, and DrJava

Instructor Resources

The following supplements are available to qualified instructors:

- Answers to all of the Review Questions
- Solutions for the Programming Challenges
- PowerPoint presentation slides for each chapter

Acknowledgments

There have been many helping hands in the development and publication of this book. We would like to thank the following faculty reviewers for their helpful suggestions and expertise:

- John Bono
- George Mason University
- Irene Bruno
- George Mason University
- Richard Burns
- Westchester University
- Jackie Horton
- University of Vermont
- Mohammad T. Islam
- Southern Connecticut State University
- David Krebs
- University of Pittsburgh
- Michael Ruth
- Roosevelt University
- Christian Servin
- El Paso Community College
- Li Yang
- Western Michigan University

Reviewers for Previous Editions

- Ahmad Abuhejleh
- University of Wisconsin, River Falls
- Colin Archibald
- Valencia Community College
- Ijaz Awani
- Savannah State University
- Bill Bane
- Tarleton State University
- N. Dwight Barnette
- Virginia Tech
- Asoke Bhattacharyya
- Saint Xavier University, Chicago
- Marvin Bishop
- Manhattan College
- eather Booth
- University of Tennessee, Knoxville
- David Boyd
- Valdosta State University
- Julius Brandstatter
- Golden Gate University
- Kim Cannon
- Greenville Tech
- Jesse Cecil

- College of the Siskiyous
- James Chegwiddden
- Tarrant County College
- Kay Chen
- Bucks County Community College
- Brad Chilton
- Tarleton State University
- Diane Christie
- University of Wisconsin, Stout
- Cara Cocking
- Marquette University
- Jose Cordova
- University of Louisiana, Monroe
- Walter C. Daugherty
- Texas A&M University
- Michael Doherty
- University of the Pacific
- Jeanne M. Douglas
- University of Vermont
- Sander Eller
- California Polytechnic University, Pomona
- Brooke Estabrook-Fishinghawk
- Mesa Community College
- Mike Fry

- Lebanon Valley College
- David Goldschmidt
- College of St. Rose
- Georgia R. Grant
- College of San Mateo
- Carl Stephen Guynes
- University of North Texas
- Nancy Harris
- James Madison University
- Chris Haynes
- Indiana University
- Ric Heishman
- Northern Virginia Community College
- Deedee Herrera
- Dodge City Community College
- Mary Hovik
- Lehigh Carbon Community College
- Brian Howard
- DePauw University
- Alan Jackson
- Oakland Community College (MI)
- Norm Jacobson
- University of California, Irvine
- Zhen Jiang

- West Chester University
- Stephen Judd
- University of Pennsylvania
- Neven Jurkovic
- Palo Alto College
- Dennis Lang
- Kansas State University
- Jiang Li
- Austin Peay State University
- Harry Lichtbach
- Evergreen Valley College
- Michael A. Long
- California State University, Chico
- Cheng Luo
- Coppin State University
- Tim Margush
- University of Akron
- Blayne E. Mayfield
- Oklahoma State University
- Scott McLeod
- Riverside Community College
- Dean Mellas
- Cerritos College
- Georges Merx

- San Diego Mesa College
- Martin Meyers
- California State University, Sacramento
- Pati Milligan
- Baylor University
- Laurie Murphy
- Pacific Lutheran University
- Steve Newberry
- Tarleton State University
- Lynne O’Hanlon
- Los Angeles Pierce College
- Merrill Parker
- Chattanooga State Technical Community College
- Bryson R. Payne
- North Georgia College and State University
- Rodney Pearson
- Mississippi State University
- Peter John Polito
- Springfield College
- Charles Robert Putnam
- California State University, Northridge
- Y. B. Reddy
- Grambling State University
- Elizabeth Riley

- Macon State College
- Felix Rodriguez
- Naugatuck Valley Community College
- Diane Rudolph
- John A Logan College
- Carolyn Schauble
- Colorado State University
- Bonnie Smith
- Fresno City College
- Daniel Spiegel
- Kutztown University
- Caroline St. Clair
- North Central College
- Karen Stanton
- Los Medanos College
- Peter van der Goes
- Rose State College
- Timothy Urness
- Drake University
- Tuan A Vo
- Mt. San Antonio College
- Xiaoying Wang
- University of Mississippi
- Yu Wu

- University of North Texas
- Zijiang Yang
- Western Michigan University

The authors would like to thank their families for their tremendous support during the development of this book. We also want to thank everyone at Pearson for making the *Starting Out With* series so successful. We are extremely fortunate to have Matt Goldstein as our editor and Meghan Jacoby as editorial assistant. Their support and encouragement makes it a pleasure to write chapters and meet deadlines. We are also fortunate to have Demetrius Hall as our marketing manager. He consistently does a great job getting this book out to the academic community. We had a great production team, led by Amanda Brands. She worked tirelessly to make this edition a reality. Thanks to you all!

About the Authors

Tony Gaddis is the principal author of the *Starting Out With* series of textbooks. Tony has nearly 20 years experience teaching computer science courses at Haywood Community College in North Carolina. He is a highly acclaimed instructor who was previously selected as the North Carolina Community College Teacher of the Year and has received the Teaching Excellence award from the National Institute for Staff and Organizational Development. The *Starting Out With* series includes introductory books using the C++ programming language, the Java™ programming language, Microsoft® Visual Basic®, Microsoft® C#®, Python, Programming Logic and Design, MIT App Inventor, and Alice, all published by Pearson.

Godfrey Muganda is Professor of Computer Science at North Central College. He teaches a wide variety of courses at the undergraduate and graduate levels including courses in Linux and Unix programming, Windows and .NET programming, web application development, web services, data structures, and algorithms. He is a past winner of the North Central College faculty award for outstanding scholarship. His primary research interests are in the area of fuzzy sets and systems.

Chapter 1 Introduction to Computers and Java

TOPICS

1. [1.1 Introduction](#)
2. [1.2 Why Program?](#)
3. [1.3 Computer Systems: Hardware and Software](#)
4. [1.4 Programming Languages](#)
5. [1.5 What Is a Program Made Of?](#)
6. [1.6 The Programming Process](#)
7. [1.7 Object-Oriented Programming](#)

1.1 Introduction

This book teaches programming using Java. Java is a powerful language that runs on practically every type of computer. It can be used to create large applications, small programs, mobile applications, and code that powers a Web site. Before plunging right into learning Java, however, this chapter will review the fundamentals of computer hardware and software, and then take a broad look at computer programming in general.

1.2 Why Program?

CONCEPT: Computers can do many different jobs because they are programmable.

Every profession has tools that make the job easier to do. Carpenters use hammers, saws, and measuring tapes. Mechanics use wrenches, screwdrivers, and ratchets. Electronics technicians use probes, scopes, and meters. Some tools are unique and can be categorized as belonging to a single profession. For example, surgeons have certain tools that are designed specifically for surgical operations. Those tools probably aren't used by anyone other than surgeons. There are some tools, however, that are used in several professions. Screwdrivers, for instance, are used by mechanics, carpenters, and many others.

The computer is a tool used by so many professions that it cannot be easily categorized. It can perform so many different jobs that it is perhaps the most versatile tool ever made. To the accountant, computers balance books, analyze profits and losses, and prepare tax reports. To the factory worker, computers control manufacturing machines and track production. To the mechanic, computers analyze the various systems in an automobile and pinpoint hard-to-find problems. The computer can do such a wide variety of tasks because it can be *programmed*. It is a machine specifically designed to follow instructions. Because of the computer's programmability, it doesn't belong to any single profession. Computers are designed to do whatever job their programs, or *software*, tell them to do.

Computer programmers do a very important job. They create software that transforms computers into the specialized tools of many trades. Without programmers, the users of computers would have no software, and without software, computers would not be able to do anything.

Computer programming is both an art and a science. It is an art because every aspect of a program should be carefully designed. Here are a few of the things that must be designed for any real-world computer program:

- The logical flow of the instructions
- The mathematical procedures
- The layout of the programming statements
- The appearance of the screens
- The way information is presented to the user
- The program's "user friendliness"
- Manuals, help systems, and/or other forms of written documentation

There is also a science to programming. Because programs rarely work right the first time they are

written, a lot of analyzing, experimenting, correcting, and redesigning is required. This demands patience and persistence of the programmer. Writing software demands discipline as well. Programmers must learn special languages such as Java because computers do not understand English or other human languages. Programming languages have strict rules that must be carefully followed.

Both the artistic and scientific nature of programming makes writing computer software like designing a car: Both cars and programs should be functional, efficient, powerful, easy to use, and pleasing to look at.

1.3 Computer Systems: Hardware and Software

CONCEPT: All computer systems consist of similar hardware devices and software components.

Hardware

Hardware refers to the physical components that a computer is made of. A computer, as we generally think of it, is not an individual device, but a system of devices. Like the instruments in a symphony orchestra, each device plays its own part. A typical computer system consists of the following major components:

- The central processing unit (CPU)
- Main memory
- Secondary storage devices
- Input devices
- Output devices

The organization of a computer system is shown in [Figure 1-1](#).




Figure 1-1 The organization of a computer system 1-1a (TL) iko/Shutterstock, 1-1b (TL) Nikita Rogul/Shutterstock, 1-1c (TLC) Feng Yu/Shutterstock , 1-1d (TLC) Chiyacat/Shutterstock, 1-1e Elkostas/Shutterstock, 1-1f (TLB) tkemot/Shutterstock, 1-1g (BC) Vitaly Korovin/Shutterstock, 1-1h (TRB) Lusoimages/Shutterstock, 1-1i (TRC) jocic/Shutterstock, 1-1j (TR) Best Pictures here/Shutterstock, 1-1k (MC) Peter

Guess/Shutterstock, 1-11 (TC)

Aquila/Shutterstock

Let's take a closer look at each of these devices.

The CPU

At the heart of a computer is its *central processing unit*, or *CPU*. The CPU's job is to fetch instructions, follow the instructions, and produce some resulting data. Internally, the central processing unit consists of two parts: the *control unit* and the *arithmetic and logic unit (ALU)*. The control unit coordinates all of the computer's operations. It is responsible for determining where to get the next instruction and regulating the other major components of the computer with control signals. The arithmetic and logic unit, as its name suggests, is designed to perform mathematical operations. The organization of the CPU is shown in [Figure 1-2](#).




Figure 1-2 The organization of the CPU

A program is a sequence of instructions stored in the computer's memory. When a computer is running a program, the CPU is engaged in a process known formally as the *fetch/decode/execute cycle*. The steps in the fetch/decode/execute cycle are as follows:

Fetch The CPU's control unit fetches, from main memory, the next instruction in the sequence of program instructions.

Decode The instruction is encoded in the form of a number. The control unit decodes the instruction and generates an electronic signal.

Execute The signal is routed to the appropriate component of the computer (such as the ALU, a disk drive, or some other device). The signal causes the component to perform an operation.

These steps are repeated as long as there are instructions to perform.

Main Memory

Commonly known as *random access memory*, or *RAM*, the computer’s main memory is a device that holds information. Specifically, RAM holds the sequences of instructions in the programs that are running and the data those programs are using.

Memory is divided into sections that hold an equal amount of data. Each section is made of eight “switches” that may be either on or off. A switch in the on position usually represents the number 1, whereas a switch in the off position usually represents the number 0. The computer stores data by setting the switches in a memory location to a pattern that represents a character or a number. Each of these switches is known as a *bit*, which stands for *binary digit*. Each section of memory, which is a collection of eight bits, is known as a *byte*. Each byte is assigned a unique number known as an *address*. The addresses are ordered from lowest to highest. A byte is identified by its address in much the same way a post office box is identified by an address. [Figure 1-3](#) shows a series of bytes with their addresses. In the illustration, sample data is stored in memory. The number 149 is stored in the byte at address 16, and the number 72 is stored in the byte at address 23.

0	1	2	3	4	5	6	7	8	9
10	11	12	13	14	15	16149	17	18	19
20	21	22	2372	24	25	26	27	28	29

Figure 1-3 Memory bytes and their addresses

RAM is usually a volatile type of memory, used only for temporary storage. When the computer is turned off, the contents of RAM are erased.

Secondary Storage

Secondary storage is a type of memory that can hold data for long periods of time—even when there is no power to the computer. Frequently used programs are stored in secondary memory and loaded into main memory as needed. Important data, such as word processing documents, payroll data, and

inventory figures, is saved to secondary storage as well.

The most common type of secondary storage device is the *disk drive*. A traditional disk drive stores data by magnetically encoding it onto a spinning circular disk. *Solid state drives*, which store data in solid-state memory, are increasingly becoming popular. A solid-state drive has no moving parts, and operates faster than a traditional disk drive. Most computers have some sort of secondary storage device, either a traditional disk drive or a solid-state drive, mounted inside their case. External drives are also available, which connect to one of the computer's communication ports. External drives can be used to create backup copies of important data or to move data to another computer.

In addition to external drives, many types of devices have been created for copying data, and for moving it to other computers. *Universal Serial Bus drives*, or *USB drives* are small devices that plug into the computer's USB (Universal Serial Bus) port, and appear to the system as a disk drive. These drives do not actually contain a disk, however. They store data in a special type of memory known as *flash memory*. USB drives are inexpensive, reliable, and small enough to be carried in your pocket.

Optical devices such as the *CD* (compact disc) and the *DVD* (digital versatile disc) are also popular for data storage. Data is not recorded magnetically on an optical disc, but is encoded as a series of pits on the disc surface. CD and DVD drives use a laser to detect the pits and thus read the encoded data. Optical discs hold large amounts of data, and because recordable CD and DVD drives are now commonplace, they make a good medium for creating backup copies of data.

Input Devices

Input is any data the computer collects from the outside world. The device that collects the data and sends it to the computer is called an *input device*. Common input devices are the keyboard, mouse, scanner, and digital camera. Disk drives, optical drives, and USB drives can also be considered input devices because programs and data are retrieved from them and loaded into the computer's memory.

Output Devices

Output is any data the computer sends to the outside world. It might be a sales report, a list of names, or a graphic image. The data is sent to an output device, which formats and presents it. Common output devices are monitors and printers. Disk drives, USB drives, and CD recorders can also be considered output devices because the CPU sends data to them to be saved.

Software

As previously mentioned, software refers to the programs that run on a computer. There are two general categories of software: operating systems and application software. An operating system is a set of programs that manages the computer's hardware devices and controls their processes. Most all

modern operating systems are multitasking, which means they are capable of running multiple programs at once. Through a technique called time sharing, a multitasking system divides the allocation of hardware resources and the attention of the CPU among all the executing programs. UNIX, Linux, Mac OS, and Windows are multitasking operating systems.

Application software refers to programs that make the computer useful to the user. These programs solve specific problems or perform general operations that satisfy the needs of the user. Word processing, spreadsheet, and database packages are all examples of application software.

Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 1.1 Why is the computer used by so many different people, in so many different professions?
2. 1.2 List the five major hardware components of a computer system.
3. 1.3 Internally, the CPU consists of what two units?
4. 1.4 Describe the steps in the fetch/decode/execute cycle.
5. 1.5 What is a memory address? What is its purpose?
6. 1.6 Explain why computers have both main memory and secondary storage.
7. 1.7 What does the term *multitasking* mean?

1.4 Programming Languages

CONCEPT: A program is a set of instructions a computer follows in order to perform a task. A programming language is a special language used to write computer programs.

What Is a Program?

Computers are designed to follow instructions. A computer program is a set of instructions that enable the computer to solve a problem or perform a task. For example, suppose we want the computer to calculate someone's gross pay. The following is a list of things the computer should do to perform this task.

1. Display a message on the screen: "How many hours did you work?"
2. Allow the user to enter the number of hours worked.
3. Once the user enters a number, store it in memory.
4. Display a message on the screen: "How much do you get paid per hour?"
5. Allow the user to enter an hourly pay rate.
6. Once the user enters a number, store it in memory.
7. Once both the number of hours worked and the hourly pay rate are entered, multiply the two numbers and store the result in memory.
8. Display a message on the screen that shows the amount of money earned. The message must include the result of the calculation performed in Step 7.

Collectively, these instructions are called an *algorithm*. An algorithm is a set of well-defined steps for performing a task or solving a problem. Notice that these steps are sequentially ordered. Step 1 should be performed before Step 2, and so forth. It is important that these instructions be performed in their proper sequence.

Although you and I might easily understand the instructions in the pay-calculating algorithm, it is not ready to be executed on a computer. A computer's CPU can only process instructions that are written in *machine language*. If you were to look at a machine language program, you would see a stream of binary numbers (numbers consisting of only 1s and 0s). The binary numbers form machine language instructions, which the CPU interprets as commands. Here is an example of what a machine language instruction might look like:

1011010000000101

As you can imagine, the process of encoding an algorithm in machine language is very tedious and difficult. In addition, each different type of CPU has its own machine language. If you wrote a machine language program for computer A and then wanted to run it on computer B, which has a different type of CPU, you would have to rewrite the program in computer B’s machine language.

Programming languages, which use words instead of numbers, were invented to ease the task of programming. A program can be written in a programming language, which is much easier to understand than machine language, and then translated into machine language. Programmers use software to perform this translation. Many programming languages have been created. [Table 1-1](#) lists a few of the well-known ones.

Table 1-1 Programming languages

Language	Description
BASIC	Beginners All-purpose Symbolic Instruction Code is a general-purpose, procedural programming language. It was originally designed to be simple enough for beginners to learn.
FORTRAN	FORmula TRANslator is a procedural language designed for programming complex mathematical algorithms.
COBOL	Common Business-Oriented Language is a procedural language designed for business applications.
Pascal	Pascal is a structured, general-purpose, procedural language designed primarily for teaching programming.
C	C is a structured, general-purpose, procedural language developed at Bell Laboratories.
C++	Based on the C language, C++ offers object-oriented features not found in C. C++ was also invented at Bell Laboratories.
C#	Pronounced “C sharp.” It is a language invented by Microsoft for developing applications based on the Microsoft .NET platform.

Java	Java is an object-oriented language invented at Sun Microsystems, and is now owned by Oracle. It may be used to develop stand-alone applications that operate on a single computer, or applications that run over the Internet from a Web server.
JavaScript	JavaScript is a programming language that can be used in a Web site to perform simple operations. Despite its name, JavaScript is not related to Java.
Perl	A general-purpose programming language used widely on Internet servers.
PHP	A programming language used primarily for developing Web server applications and dynamic Web pages.
Python	Python is an object-oriented programming language used in both business and academia. Many popular Web sites contain features developed in Python.
Ruby	Ruby is a simple but powerful object-oriented programming language. It can be used for a variety of purposes, from small utility programs to large Web applications.
Visual Basic	Visual Basic is a Microsoft programming language and software development environment that allows programmers to create Windows-based applications quickly.

A History of Java

In 1991 a team was formed at Sun Microsystems (a company that is now owned by Oracle) to speculate about the important technological trends that might emerge in the near future. The team, which was named the Green Team, concluded that computers would merge with consumer appliances. Their first project was to develop a handheld device named *7 (pronounced star seven) that could be used to control a variety of home entertainment devices. For the unit to work, it had to use a programming language that could be processed by all the devices it controlled. This presented a problem because different brands of consumer devices use different processors, each with its own machine language.

Because no such universal language existed, James Gosling, the team's lead engineer, created one.

Programs written in this language, which was originally named Oak, were not translated into the machine language of a specific processor, but were translated into an intermediate language known as *byte code*. Another program would then translate the byte code into machine language that could be executed by the processor in a specific consumer device.

Unfortunately, the technology developed by the Green Team was ahead of its time. No customers could be found, mostly because the computer-controlled consumer appliance industry was just beginning. But rather than abandoning their hard work and moving on to other projects, the team saw another opportunity: the Internet. The Internet is a perfect environment for a universal programming language such as Oak. It consists of numerous different computer platforms connected together in a single network.

To demonstrate the effectiveness of its language, which was renamed Java, the team used it to develop a Web browser. The browser, named HotJava, was able to download and run small Java programs known as applets. This gave the browser the capability to display animation and interact with the user. HotJava was demonstrated at the 1995 SunWorld conference before a wowed audience. Later the announcement was made that Netscape would incorporate Java technology into its Navigator browser. Other Internet companies rapidly followed, increasing the acceptance and the influence of the Java language. Today, Java is very popular for developing not only applets for developing Web applications, mobile apps, and desktop applications.

1.5 What Is a Program Made Of?

CONCEPT: There are certain elements that are common to all programming languages.

Language Elements

All programming languages have some things in common. [Table 1-2](#) lists the common elements you will find in almost every language.

Table 1-2 The common elements of a programming language

Language Element	Description
Key Words	These are words that have a special meaning in the programming language. They may be used for their intended purpose only. Key words are also known as reserved words.
Operators	Operators are symbols or words that perform operations on one or more operands. An operand is usually an item of data, such as a number.
Punctuation	Most programming languages require the use of punctuation characters. These characters serve specific purposes, such as marking the beginning or ending of a statement, or separating items in a list.
Programmer-Defined Names	Unlike key words, which are part of the programming language, these are words or names that are defined by the programmer. They are used to identify storage locations in memory and parts of the program that are created by the programmer. Programmer-defined names are often called identifiers.

These are rules that must be followed when writing a program. Syntax

Syntax dictates how key words and operators may be used, and where punctuation symbols must appear.

Let's look at an example Java program and identify an instance of each of these elements. [Code Listing 1-1](#) shows the code listing with each line numbered.

NOTE:

The line numbers are not part of the program. They are included to help point out specific parts of the program.

Code Listing 1-1 Payroll.java

```
1 public class Payroll
2 {
3     public static void main(String[] args)
4     {
5         int hours = 40;
6         double grossPay, payRate = 25.0;
7
8         grossPay = hours * payRate;
9         System.out.println("Your gross pay is $" + grossPay);
10    }
11 }
```

Key Words (Reserved Words)

Two of Java's key words appear in line 1: `public` and `class`. In line 3, the words `public`, `static`, and `void` are all key words. The words `int` in line 5 and `double` in line 6 are also key words. These words, which are always written in lowercase, each have a special meaning in Java and can only be used for their intended purpose. As you will see, the programmer is allowed to make up his or her own names for certain things in a program. Key words, however, are reserved and cannot be used for anything other than their designated purpose. Part of learning a programming language is learning the commonly used key words, what they mean, and how to use them.

[Table 1-3](#) shows a list of the Java key words¹.

¹Java 9 also introduces a set of *restricted words* that are treated as key words under certain circumstances. See [Appendix C](#) for the full list.

Table 1-3 The Java key words

<code>abstract</code>	<code>const</code>	<code>final</code>	<code>int</code>	<code>public</code>	<code>throw</code>
<code>assert</code>	<code>continue</code>	<code>finally</code>	<code>interface</code>	<code>return</code>	<code>throws</code>
<code>boolean</code>	<code>default</code>	<code>float</code>	<code>long</code>	<code>short</code>	<code>transient</code>
<code>break</code>	<code>do</code>	<code>for</code>	<code>native</code>	<code>static</code>	<code>true</code>
<code>byte</code>	<code>double</code>	<code>goto</code>	<code>new</code>	<code>strictfp</code>	<code>try</code>
<code>case</code>	<code>else</code>	<code>if</code>	<code>null</code>	<code>super</code>	<code>void</code>
<code>catch</code>	<code>enum</code>	<code>implements</code>	<code>package</code>	<code>switch</code>	<code>volatile</code>
<code>char</code>	<code>extends</code>	<code>import</code>	<code>private</code>	<code>synchronized</code>	<code>while</code>
<code>class</code>	<code>false</code>	<code>instanceof</code>	<code>protected</code>	<code>this</code>	

Programmer-Defined Names

The words `hours`, `payRate`, and `grossPay` that appear in the program in lines 5, 6, 8, and 9 are programmer-defined names. They are not part of the Java language but are names made up by the programmer. In this particular program, these are the names of variables. As you will learn later in this chapter, variables are the names of memory locations that may hold data.

Operators

In line 8 the following line appears:

```
grossPay = hours * payRate;
```

The `=` and `*` symbols are both operators. They perform operations on items of data, known as

operands. The `*` operator multiplies its two operands, which in this example are the variables `hours` and `payRate`. The `=` symbol is called the assignment operator. It takes the value of the expression that appears at its right and stores it in the variable whose name appears at its left. In this example, the `=` operator stores in the `grossPay` variable the result of the `hours` variable multiplied by the `payRate` variable. In other words, the statement says, “the `grossPay` variable is assigned the value of `hours` times `payRate`.”

Punctuation

Notice that lines 5, 6, 8, and 9 end with a semicolon. A semicolon in Java is similar to a period in English: It marks the end of a complete sentence (or statement, as it is called in programming jargon). Semicolons do not appear at the end of every line in a Java program, however. There are rules that govern where semicolons are required and where they are not. Part of learning Java is learning where to place semicolons and other punctuation symbols.

Lines and Statements

Often, the contents of a program are thought of in terms of lines and statements. A *line* is just that—a single line as it appears in the body of a program. [Code Listing 1-1](#) is shown with each of its lines numbered. Most of the lines contain something meaningful; however, line 7 is empty. Blank lines are only used to make a program more readable.

A statement is a complete instruction that causes the computer to perform some action. Here is the statement that appears in line 9 of [Code Listing 1-1](#):

```
System.out.println("Your gross pay is $" + grossPay);
```

This statement causes the computer to display a message on the screen. Statements can be a combination of key words, operators, and programmer-defined names. Statements often occupy only one line in a program, but sometimes they are spread out over more than one line.

Variables

The most fundamental way that a Java program stores an item of data in memory is with a variable. A *variable* is a named storage location in the computer’s memory. The data stored in a variable may change while the program is running (hence the name “variable”). Notice that in [Code Listing 1-1](#) the programmer-defined names `hours`, `payRate`, and `grossPay` appear in several places. All three of these are the names of variables. The `hours` variable is used to store the number of hours the user has worked. The `payRate` variable stores the user’s hourly pay rate. The `grossPay` variable holds the result of `hours` multiplied by `payRate`, which is the user’s gross pay.

Variables are symbolic names made up by the programmer that represent locations in the computer's RAM. When data is stored in a variable, it is actually stored in RAM. Assume that a program has a variable named `length`. [Figure 1-4](#) illustrates the way the variable name represents a memory location.




Figure 1-4 A variable name represents a location in memory

In [Figure 1-4](#), the variable `length` is holding the value 72. The number 72 is actually stored in RAM at address 23, but the name `length` symbolically represents this storage location. If it helps, you can think of a variable as a box that holds data. In [Figure 1-4](#), the number 72 is stored in the box named `length`. Only one item may be stored in the box at any given time. If the program stores another value in the box, it will take the place of the number 72.

The Compiler and the Java Virtual Machine

When a Java program is written, it must be typed into the computer and saved to a file. A *text editor*, which is similar to a word processing program, is used for this task. The Java programming statements written by the programmer are called *source code*, and the file they are saved in is called a *source file*. Java source files end with the *.java* extension.

After the programmer saves the source code to a file, he or she runs the Java compiler. A *compiler* is a program that translates source code into an executable form. During the translation process, the compiler uncovers any syntax errors that may be in the program. *Syntax errors* are mistakes that the programmer has made that violate the rules of the programming language. These errors must be corrected before the compiler can translate the source code. Once the program is free of syntax errors, the compiler creates another file that holds the translated instructions.

Most programming language compilers translate source code directly into files that contain machine language instructions. These are called *executable files* because they may be executed directly by the computer's CPU. The Java compiler, however, translates a Java source file into a file that contains byte code instructions. Byte code instructions are not machine language, and therefore cannot be directly executed by the CPU. Instead, they are executed by the *Java Virtual Machine (JVM)*. The

JVM is a program that reads Java byte code instructions and executes them as they are read. For this reason, the JVM is often called an interpreter, and Java is often referred to as an interpreted language. [Figure 1-5](#) illustrates the process of writing a Java program, compiling it to byte code, and running it.




Figure 1-5 Program development process

[Figure 1-5 Full Alternative Text](#)

Although Java byte code is not machine language for a CPU, it can be considered as machine language for the JVM. You can think of the JVM as a program that simulates a computer whose machine language is Java byte code.

Portability

The term *portable* means that a program may be written on one type of computer and then run on a wide variety of computers, with little or no modification necessary. Because Java byte code is the same on all computers, compiled Java programs are highly portable. In fact, a compiled Java program may be run on any computer that has a Java Virtual Machine. [Figure 1-6](#) illustrates the concept of a compiled Java program running on Windows, Linux, Mac, and UNIX computers.




Figure 1-6 Java byte code may be run on any computer with a Java Virtual Machine

With most other programming languages, portability is achieved by the creation of a compiler for each type of computer that the language is to run on. For example, in order for the C++ language to be supported by Windows, Linux, and Mac computers, a separate C++ compiler must be created for each of those environments. Compilers are very complex programs, and more difficult to develop than interpreters. For this reason, a JVM has been developed for many types of computers.

Java Software Editions

The software that you use to create Java programs is referred to as the *JDK* (Java Development Kit) or the *SDK* (Software Development Kit). There are the following different editions of the JDK available from Oracle:

- *Java SE*—The Java Standard Edition provides all the essential software tools necessary for writing Java applications.
- *Java EE*—The Java Enterprise Edition provides tools for creating large business applications that employ servers and provide services over the Web.
- *Java ME*—The Java Micro Edition provides a small, highly optimized runtime environment for consumer products such as cell phones, pagers, and appliances.

These editions of Java may be downloaded from Oracle by going to:

<http://java.oracle.com>



You can follow the instructions in Appendix D, which can be downloaded from the book's companion Web site, to install the JDK on your system. You can access the book's companion Web site by going to www.pearsonhighered.com/gaddis.

Compiling and Running a Java Program

Compiling a Java program is a simple process. Once you have installed the JDK, go to your operating system's command prompt.



In Windows click Start, go to All Programs, and then go to Accessories. Click Command Prompt on the Accessories menu. A command prompt window should open.



Compiling and Running a Java Program

At the operating system command prompt, make sure you are in the same directory or folder where the Java program that you want to compile is located. Then, use the `javac` command, in the following

form:

```
javac Filename
```

Filename is the name of a file that contains the Java source code. As mentioned earlier, this file has the *.java* extension. For example, if you want to compile the *Payroll.java* file, you would execute the following command:

```
javac Payroll.java
```

This command runs the compiler. If the file contains any syntax errors, you will see one or more error messages and the compiler will not translate the file to byte code. When this happens you must open the source file in a text editor and fix the error. Then you can run the compiler again. If the file has no syntax errors, the compiler will translate it to byte code. Byte code is stored in a file with the *.class* extension, so the byte code for the *Payroll.java* file will be stored in *Payroll.class*, which will be in the same directory or folder as the source file.

To run the Java program, you use the `java` command in the following form:

```
java ClassFilename
```

ClassFilename is the name of the *.class* file that you wish to execute; however, you do not type the *.class* extension. For example, to run the program that is stored in the *Payroll.class* file, you would enter the following command:

```
java Payroll
```

This command runs the Java interpreter (the JVM) and executes the program.

Integrated Development Environments

In addition to the command prompt programs, there are also several Java integrated development environments (IDEs). These environments consist of a text editor, compiler, debugger, and other utilities integrated into a package with a single set of menus. A program is compiled and executed with a single click of a button, or by selecting a single item from a menu. [Figure 1-7](#) shows a screen from the NetBeans IDE.




Figure 1-7 An integrated development environment (IDE)

(Oracle Corporate Counsel)



Using an IDE



MyProgrammingLab™ www.myprogramminglab.com

1. 1.8 Describe the difference between a key word and a programmer-defined symbol.

2. 1.9 Describe the difference between operators and punctuation symbols.
3. 1.10 Describe the difference between a program line and a statement.
4. 1.11 Why are variables called “variable”?
5. 1.12 What happens to a variable’s current contents when a new value is stored there?
6. 1.13 What is a compiler?
7. 1.14 What is a syntax error?
8. 1.15 What is byte code?
9. 1.16 What is the JVM?

1.6 The Programming Process

CONCEPT: The programming process consists of several steps, which include design, creation, testing, and debugging activities.

Now that you have been introduced to what a program is, it's time to consider the process of creating a program. Quite often when inexperienced students are given programming assignments, they have trouble getting started because they don't know what to do first. If you find yourself in this dilemma, the following steps may help.

1. Clearly define what the program is to do.
2. Visualize the program running on the computer.
3. Use design tools to create a model of the program.
4. Check the model for logical errors.
5. Enter the code and compile it.
6. Correct any errors found during compilation. Repeat Steps 5 and 6 as many times as necessary.
7. Run the program with test data for input.
8. Correct any runtime errors found while running the program. Repeat Steps 5 through 8 as many times as necessary.
9. Validate the results of the program.

These steps emphasize the importance of planning. Just as there are good ways and bad ways to paint a house, there are good ways and bad ways to create a program. A good program always begins with planning. With the pay-calculating algorithm that was presented earlier in this chapter serving as our example, let's look at each of the steps in more detail.

1. Clearly define what the program is to do

This step commonly requires you to identify the purpose of the program, the data that is to be input, the processing that is to take place, and the desired output. Let's examine each of these requirements for the pay-calculating algorithm.

Purpose To calculate the user's gross pay.


Input Number of hours worked, hourly pay rate.

Process Multiply number of hours worked by hourly pay rate. The result is the user's gross pay.

Output Display a message indicating the user's gross pay.

2. Visualize the program running on the computer

Before you create a program on the computer, you should first create it in your mind. Try to imagine what the computer screen will look like while the program is running. If it helps, draw pictures of the screen, with sample input and output, at various points in the program. For instance, [Figure 1-8](#) shows the screen we might want produced by a program that implements the pay-calculating algorithm.



```
How many hours did you work? 10
How much do you get paid per hour? 15
Your gross pay is $150.0
```

Figure 1-8 Screen produced by the pay-calculating algorithm

In this step, you must put yourself in the shoes of the user. What messages should the program display? What questions should it ask? By addressing these concerns, you can determine most of the program's output.

3. Use design tools to create a model of the program

While planning a program, the programmer uses one or more design tools to create a model of the program. For example, *pseudocode* is a cross between human language and a programming language and is especially helpful when designing an algorithm. Although the computer can't understand pseudocode, programmers often find it helpful to write an algorithm in a language that's "almost" a programming language, but still very similar to natural language. For example, here is pseudocode that describes the pay-calculating algorithm:

- *Get payroll data.*
- *Calculate gross pay.*
- *Display gross pay.*

Although this pseudocode gives a broad view of the program, it doesn't reveal all the program's details. A more detailed version of the pseudocode follows:

- *Display "How many hours did you work?"*
- *Input hours.*
- *Display "How much do you get paid per hour?"*
- *Input rate.*
- *Store the value of hours times rate in the pay variable.*
- *Display the value in the pay variable.*

Notice that the pseudocode uses statements that look more like commands than the English statements that describe the algorithm in Section 1.4. The pseudocode even names variables and describes mathematical operations.

4. Check the model for logical errors

Logical errors are mistakes that cause the program to produce erroneous results. Once a model of the program is assembled, it should be checked for these errors. For example, if pseudocode is used, the programmer should trace through it, checking the logic of each step. If an error is found, the model can be corrected before the next step is attempted.

5. Enter the code and compile it

Once a model of the program has been created, checked, and corrected, the programmer is ready to write source code on the computer. The programmer saves the source code to a file and begins the process of compiling it. During this step the compiler will find any syntax errors that may exist in the program.

6. Correct any errors found during compilation. Repeat Steps 5 and 6 as many

times as necessary

If the compiler reports any errors, they must be corrected. Steps 5 and 6 must be repeated until the program is free of compile-time errors.

7. Run the program with test data for input

Once an executable file is generated, the program is ready to be tested for runtime errors. A runtime error is an error that occurs while the program is running. These are usually logical errors, such as mathematical mistakes.

Testing for runtime errors requires that the program be executed with sample data or sample input. The sample data should be such that the correct output can be predicted. If the program does not produce the correct output, a logical error is present in the program.

8. Correct any runtime errors found while running the program. Repeat Steps 5 through 8 as many times as necessary

When runtime errors are found in a program, they must be corrected. You must identify the step where the error occurred and determine the cause. If an error is a result of incorrect logic (such as an improperly stated math formula), you must correct the statement or statements involved in the logic. If an error is due to an incomplete understanding of the program requirements, then you must restate the program purpose and modify the program model and source code. The program must then be saved, recompiled, and retested. This means Steps 5 through 8 must be repeated until the program reliably produces satisfactory results.

9. Validate the results of the program

When you believe you have corrected all the runtime errors, enter test data and determine whether the program solves the original problem.

Software Engineering

The field of software engineering encompasses the whole process of crafting computer software. It includes designing, writing, testing, debugging, documenting, modifying, and maintaining complex

software development projects. Like traditional engineers, software engineers use a number of tools in their craft. Here are a few examples:

- Program specifications
- Diagrams of screen output
- Diagrams representing the program components and the flow of data
- Pseudocode
- Examples of expected input and desired output
- Special software designed for testing programs

Most commercial software applications are large and complex. Usually a team of programmers, not a single individual, develops them. It is important that the program requirements be thoroughly analyzed and divided into subtasks that are handled by individual teams, or individuals within a team.



MyProgrammingLab™ www.myprogramminglab.com

1. 1.17 What four items should you identify when defining what a program is to do?
2. 1.18 What does it mean to “visualize a program running”? What is the value of such an activity?
3. 1.19 What is pseudocode?
4. 1.20 Describe what a compiler does with a program’s source code.
5. 1.21 What is a runtime error?
6. 1.22 Is a syntax error (such as misspelling a key word) found by the compiler or when the program is running?
7. 1.23 What is the purpose of testing a program with sample data or input?

1.7 Object-Oriented Programming

CONCEPT: Java is an object-oriented programming (OOP) language. OOP is a method of software development that has its own practices, concepts, and vocabulary.

There are primarily two methods of programming in use today: procedural and object-oriented. The earliest programming languages were procedural, meaning a program was made of one or more procedures. A *procedure* is a set of programming statements that, together, perform a specific task. The statements might gather input from the user, manipulate data stored in the computer's memory, and perform calculations or any other operation necessary to complete the procedure's task.

Procedures typically operate on data items that are separate from the procedures. In a procedural program, the data items are commonly passed from one procedure to another, as shown in [Figure 1-9](#).




Figure 1-9 Data is passed among procedures

As you might imagine, the focus of procedural programming is on the creation of procedures that operate on the program's data. The separation of data and the code that operates on the data often leads to problems, however. For example, the data is stored in a particular format, which consists of variables and more complex structures that are created from variables. The procedures that operate on the data must be designed with that format in mind. But, what happens if the format of the data is altered? Quite often, a program's specifications change, resulting in a redesigned data format. When the structure of the data changes, the code that operates on the data must also be changed to accept the new format. This results in added work for programmers and a greater opportunity for bugs to appear in the code.

This has helped influence the shift from procedural programming to object-oriented programming (OOP). Whereas procedural programming is centered on creating procedures, object-oriented programming is centered on creating objects. An object is a software entity that contains data and

procedures. The data contained in an object is known as the object's *attributes*. The procedures, or behaviors, that an object performs are known as the object's *methods*. The object is, conceptually, a self-contained unit consisting of data (attributes) and procedures (methods). This is illustrated in [Figure 1-10](#).




Figure 1-10 An object contains data and procedures

OOP addresses the problem of code/data separation through encapsulation and data hiding. *Encapsulation* refers to the combining of data and code into a single object. *Data hiding* refers to an object's ability to hide its data from code that is outside the object. Only the object's methods may then directly access and make changes to the object's data. An object typically hides its data, but allows outside code to access the methods that operate on the data. As shown in [Figure 1-11](#), the object's methods provide programming statements outside the object with indirect access to the object's data.




Figure 1-11 Code outside the object interacts with the object's methods

When an object's internal data is hidden from outside code and access to that data is restricted to the object's methods, the data is protected from accidental corruption. In addition, the programming code outside the object does not need to know about the format or internal structure of the object's data. The code only needs to interact with the object's methods. When a programmer changes the structure of an object's internal data, he or she also modifies the object's methods so they may properly operate on the data. The way in which outside code interacts with the methods, however, does not change.

These are just a few of the benefits of object-oriented programming. Because Java is fully object-oriented, you will learn much more about OOP practices, concepts, and terms as you progress through this book.

Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 1.24 In procedural programming, what two parts of a program are typically separated?
2. 1.25 What are an object's attributes?
3. 1.26 What are an object's methods?
4. 1.27 What is encapsulation?

5. 1.28 What is data hiding?

Review Questions and Exercises

Multiple Choice

1. This part of the computer fetches instructions, carries out the operations commanded by the instructions, and produces some outcome or resultant information.
 1. memory
 2. CPU
 3. secondary storage
 4. input device
2. A byte is made up of eight
 1. CPUs
 2. addresses
 3. variables
 4. bits
3. Each byte is assigned a unique
 1. address
 2. CPU
 3. bit
 4. variable
4. This type of memory can hold data for long periods of time—even when there is no power to the computer.
 1. RAM
 2. primary storage
 3. secondary storage

4. CPU storage
5. If you were to look at a machine language program, you would see .
 1. Java source code
 2. a stream of binary numbers
 3. English words
 4. circuits
6. These are words that have a special meaning in the programming language.
 1. punctuation
 2. programmer-defined names
 3. key words
 4. operators
7. These are symbols or words that perform operations on one or more operands.
 1. punctuation
 2. programmer-defined names
 3. key words
 4. operators
8. These characters serve specific purposes, such as marking the beginning or ending of a statement, or separating items in a list.
 1. punctuation
 2. programmer-defined names
 3. key words
 4. operators
9. These are words or names that are used to identify storage locations in memory and parts of the program that are created by the programmer.
 1. punctuation

2. programmer-defined names
 3. key words
 4. operators
10. These are the rules that must be followed when writing a program.
1. syntax
 2. punctuation
 3. key words
 4. operators
11. This is a named storage location in the computer's memory.
1. class
 2. key word
 3. variable
 4. operator
12. The Java compiler generates .
1. machine code
 2. byte code
 3. source code
 4. HTML
13. JVM stands for .
1. Java Variable Machine
 2. Java Variable Method
 3. Java Virtual Method
 4. Java Virtual Machine

Find the Error

1. The following pseudocode algorithm has an error. The program is supposed to ask the user for the length and width of a rectangular room, and then display the room's area. The program must multiply the width by the length to determine the area. Find the error.
 - $\text{area} = \text{width} \times \text{length}$
 - *Display "What is the room's width?"*
 - *Input width*
 - *Display "What is the room's length?"*
 - *Input length*
 - *Display area*

Algorithm Workbench

Write pseudocode algorithms for the programs described as follows:

1. Available Credit

A program that calculates a customer's available credit should ask the user for the following:

- The customer's maximum amount of credit
- The amount of credit used by the customer

Once these items have been entered, the program should calculate and display the customer's available credit. You can calculate available credit by subtracting the amount of credit used from the maximum amount of credit.

2. Sales Tax

A program that calculates the total of a retail sale should ask the user for the following:

- The retail price of the item being purchased
- The sales tax rate

Once these items have been entered, the program should calculate and display the following:

- The sales tax for the purchase
- The total of the sale

3. Account Balance

A program that calculates the current balance in a savings account must ask the user for the following:

- The starting balance
- The total dollar amount of deposits made
- The total dollar amount of withdrawals made
- The monthly interest rate

Once the program calculates the current balance, it should be displayed on the screen.

Predict the Result

The following are programs expressed as English statements. What would each display on the screen if they were actual programs?

1. The variable x starts with the value 0.

- The variable y starts with the value 5.
- Add 1 to x .
- Add 1 to y .
- Add x and y , and store the result in y .
- Display the value in y on the screen.

2. The variable a starts with the value 10.

- The variable b starts with the value 2.
- The variable c starts with the value 4.
- Store the value of a times b in a .
- Store the value of b times c in c .
- Add a and c , and store the result in b .
- Display the value in b on the screen.

Short Answer

1. Both main memory and secondary storage are types of memory. Describe the difference between the two.
2. What type of memory is usually volatile?
3. What is the difference between operating system software and application software?
4. Why must programs written in a high-level language be translated into machine language before they can be run?
5. Why is it easier to write a program in a high-level language than in machine language?
6. What is a source file?
7. What is the difference between a syntax error and a logical error?
8. What is an algorithm?
9. What is a compiler?
10. What must a computer have in order for it to execute Java programs?
11. What is the difference between machine language code and byte code?
12. Why does byte code make Java a portable language?
13. Is encapsulation a characteristic of procedural or object-oriented programming?
14. Why should an object hide its data?
15. What part of an object forms an interface through which outside code may access the object's data?
16. What type of program do you use to write Java source code?
17. Will the Java compiler translate a source file that contains syntax errors?
18. What does the Java compiler translate Java source code to?
19. Assuming you are using the JDK, what command would you type at the operating system command prompt to compile the program `LabAssignment.java`?
20. Assuming there are no syntax errors in the *LabAssignment.java* program when it is compiled, answer the following questions.

1. What file will be produced?
2. What will the file contain?
3. What command would you type at the operating system command prompt to run the program?

Programming Challenge

MyProgrammingLab™ Visit www.myprogramminglab.com to complete many of these Programming Challenges online and get instant feedback.

1. Your First Java Program

This assignment will help you get acquainted with your Java development software. Here is the Java program you will enter:



Your First Java Program

```
// This is my first Java program.
public class MyFirstProgram
{
    public static void main(String[] args)
    {
        System.out.println("Hello World!");
    }
}
```

If You Are Using the JDK at the Command Prompt:

1. Use a text editor to type the source code exactly as it is shown. Be sure to place all the punctuation characters and be careful to match the case of the letters as they are shown. Save it to a file named *MyFirstProgram.java*.
2. After saving the program, go to your operating system's command prompt and change your current directory or folder to the one that contains the Java program you just created. Then use the following command to compile the program:

```
javac MyFirstProgram.java
```

If you typed the contents of the file exactly as shown, you shouldn't have any syntax errors. If you see error messages, open the file in the editor and compare your code to that shown. Correct any mistakes you have made, save the file, and run the compiler again. If you see no error messages, the file was successfully compiled.

3. Next, enter the following command to run the program:

```
java MyFirstProgram
```

Be sure to use the capitalization of `MyFirstProgram` exactly as it is shown here. You should see the message “Hello World!” displayed on the screen.

If You Are Using an IDE:

Because there are many Java IDEs, we cannot include specific instructions for all of these. The following are general steps that should apply to most of them. You will need to consult your IDE’s documentation for specific instructions.

1. Start your Java IDE and perform any necessary setup operations, such as starting a new project and creating a new Java source file.
2. Use the IDE’s text editor to type the source code exactly as it is shown. Be sure to place all the punctuation characters and be careful to match the case of the letters as they are shown. Save it to a file named *MyFirstProgram.java*.
3. After saving the program, use your IDE’s command to compile the program. If you typed the contents of the file exactly as shown, you shouldn’t have any syntax errors. If you see error messages, compare your code to that shown. Correct any mistakes you have made, save the file, and run the compiler again. If you see no error messages, the file was successfully compiled.

Use your IDE’s command to run the program. You should see the message “Hello World!” displayed.

Chapter 2 Java Fundamentals

TOPICS

1. [2.1 The Parts of a Java Program](#)
2. [2.2 The `print` and `println` Methods, and the Java API](#)
3. [2.3 Variables and Literals](#)
4. [2.4 Primitive Data Types](#)
5. [2.5 Arithmetic Operators](#)
6. [2.6 Combined Assignment Operators](#)
7. [2.7 Conversion between Primitive Data Types](#)
8. [2.8 Creating Named Constants with `final`](#)
9. [2.9 The `String` Class](#)
10. [2.10 Scope](#)
11. [2.11 Comments](#)
12. [2.12 Programming Style](#)
13. [2.13 Reading Keyboard Input](#)
14. [2.14 Dialog Boxes](#)
15. [2.15 Common Errors to Avoid](#)

2.1 The Parts of a Java Program

CONCEPT: A Java program has parts that serve specific purposes.

Java programs are made up of different parts. Your first step in learning Java is to learn what the parts are. We will begin by looking at a simple example, shown in [Code Listing 2-1](#).

Code Listing 2-1 (*Simple.java*)

```
1  // This is a simple Java program.
2
3  public class Simple
4  {
5      public static void main(String[] args)
6      {
7          System.out.println("Programming is great fun!");
8      }
9  }
```



Remember, the line numbers shown in the program listings are not part of the program. The numbers are shown so we can refer to specific lines in the programs.

As mentioned in [Chapter 1](#), the names of Java source code files end with *.java*. The program shown in [Code Listing 2-1](#) is named *Simple.java*. Using the Java compiler, this program may be compiled with the following command:

```
javac Simple.java
```

The compiler will create another file named *Simple.class*, which contains the translated Java byte code. This file can be executed with the following command:

```
java Simple
```



Remember, you do not type the *.class* extension when using the `java` command.

The output of the program is as follows. This is what appears on the screen when the program runs.

Program Output

Programming is great fun!

Lets's examine the program line by line. Here's the statement in line 1:

```
// This is a simple Java program.
```

Other than the two slash marks that begin this line, it looks pretty much like an ordinary sentence. The `//` marks the beginning of a comment. The compiler ignores everything from the double-slash to the end of the line. That means you can type anything you want on that line and the compiler never complains. Although comments are not required, they are very important to programmers. Most programs are much more complicated than this example, and comments help explain what's going on.

Line 2 is blank. Programmers often insert blank lines in programs to make them easier to read. Line 3 reads:

```
public class Simple
```

This line is known as a *class header*, and it marks the beginning of a *class definition*. One of the uses of a class is to serve as a container for an application. As you progress through this book, you will learn more and more about classes. For now, just remember that a Java program must have at least one class definition. This line of code consists of three words: `public`, `class`, and `Simple`. Let's take a closer look at each word.

- `public` is a Java key word, and it must be written in all lowercase letters. It is known as an *access specifier*, and it controls where the class may be accessed from. The `public` specifier means access to the class is unrestricted. (In other words, the class is “open to the public.”)
- `class`, which must also be written in lowercase letters, is a Java key word that indicates the beginning of a class definition.
- `Simple` is the class name. This name was made up by the programmer. The class could have been called `Pizza`, or `Dog`, or anything else the programmer wanted. Programmer-defined names may be written in lowercase letters, uppercase letters, or a mixture of both.

In a nutshell, this line of code tells the compiler that a publicly accessible class named `Simple` is being defined. Here are two more points to know about classes:

- You may create more than one class in a file, but you may have only one `public class` per Java file.
- When a Java file has a `public class`, the name of the public class must be the same as the name of the file (without the `.java` extension). For instance, the program in [Code Listing 2-1](#) has a `public class` named `Simple`, so it is stored in a file named *Simple.java*.

 NOTE:

Java is a case-sensitive language. That means it regards uppercase letters as being entirely different characters than their lowercase counterparts. The word `Public` is not the same as `public`, and `Class` is not the same as `class`. Some words in a Java program must be entirely in lowercase, while other words may use a combination of lower and uppercase characters. Later in this chapter, you will see a list of all the Java key words, which must appear in lowercase.

Line 4 contains only a single character:

```
{
```

This is called a left brace, or an opening brace, and is associated with the beginning of the class definition. All of the programming statements that are part of the class are enclosed in a set of braces. If you glance at the last line in the program, line 9, you'll see the closing brace. Everything between the two braces is the *body* of the class named `Simple`. Here is the program code again, this time the body of the class definition is shaded.

```
// This is a simple Java program.  
public class Simple  
{  
    public static void main(String[] args)  
    {  
        System.out.println("Programming is great fun!");  
    }  
}
```

 WARNING!

Make sure you have a closing brace for every opening brace in your program!

Line 5 reads:

```
public static void main(String[] args)
```

This line is known as a *method header*. It marks the beginning of a *method*. A method can be thought of as a group of one or more programming statements that collectively has a name. When creating a method, you must tell the compiler several things about it. That is why this line contains so many words. At this point, the only thing you should be concerned about is that the name of the method is `main`, and the rest of the words are required for the method to be properly defined. This is shown in [Figure 2-1](#).

Name of the Method
↓

```
public static void main(String[] args)
```

The other parts of this line are necessary
for the method to be properly defined.

Figure 2-1 The `main` method header

Recall from [Chapter 1](#) that a stand-alone Java program that runs on your computer is known as an application. Every Java application must have a method named `main`. The `main` method is the starting point of an application.

NOTE:

For the time being, all the programs you will write will consist of a class with a `main` method whose header looks exactly like the one shown in [Code Listing 2-1](#). As you progress through this book you will learn what `public static void` and `(String[] args)` mean. For now, just assume that you are learning a “recipe” for assembling a Java program.

Line 6 has another opening brace:

```
{
```

This opening brace belongs to the `main` method. Remember that braces enclose statements, and every opening brace must have an accompanying closing brace. If you look at line 8 you will see the closing brace that corresponds with this opening brace. Everything between these braces is the *body* of the `main` method.

Line 7 appears as follows:

```
System.out.println("Programming is great fun!");
```

To put it simply, this line displays a message on the screen. The message, “Programming is great fun!” is printed without the quotation marks. In programming terms, the group of characters inside the quotation marks is called a *string literal*.

NOTE:

This is the only line in the program that causes anything to be printed on the screen. The other lines, like `public class Simple` and `public static void main(String[] args)`, are necessary for

the framework of your program, but they do not cause any screen output. Remember, a program is a set of instructions for the computer. If something is to be displayed on the screen, you must use a programming statement for that purpose.

At the end of the line is a *semicolon*. Just as a period marks the end of a sentence, a semicolon marks the end of a statement in Java. Not every line of code ends with a semicolon, however. Here is a summary of where you do not place a semicolon:

- Comments do not have to end with a semicolon because they are ignored by the compiler.
- Class headers and method headers do not end with a semicolon because they are terminated with a body of code inside braces.
- The brace characters, { and }, are not statements, so you do not place a semicolon after them.

It might seem that the rules for where to put a semicolon are not clear at all. For now, just concentrate on learning the parts of a program. You'll soon get a feel for where you should and should not use semicolons.

As has already been pointed out, lines 8 and 9 contain the closing braces for the `main` method and the class definition:

```
    }  
}
```

Before continuing, let's review the points we just covered, including some of the more elusive rules.

- Java is a case-sensitive language. It does not regard uppercase letters as being the same character as their lowercase equivalents.
- All Java programs must be stored in a file with a name that ends with *.java*.
- Comments are ignored by the compiler.
- A *.java* file may contain many classes, but may have only one `public` class. If a *.java* file has a public class, the class must have the same name as the file. For instance, if the file *Pizza.java* contains a `public class`, the class's name would be `Pizza`.
- Every Java application program must have a method named `main`.
- For every left brace, or opening brace, there must be a corresponding right brace, or closing brace.
- Statements are terminated with semicolons. This does not include comments, class headers, method headers, or braces.

In the sample program, you encountered several special characters. [Table 2-1](#) summarizes how they

were used.

Table 2-1 Special characters

Characters	Name	Meaning
//	Double slash	Marks the beginning of a comment
()	Opening and closing parentheses	Used in a method header
{ }	Opening and closing braces	Encloses a group of statements, such as the contents of a class or a method
" "	Quotation marks	Encloses a string of characters, such as a message that is to be printed on the screen
;	Semicolon	Marks the end of a complete programming statement



MyProgrammingLab™ www.myprogramminglab.com

1. 2.1 The following program will not compile because the lines have been mixed up.

```
public static void main(String[] args)
{
// A crazy mixed up program
public class Columbus
{
System.out.println("In 1492 Columbus sailed the ocean blue.");
{
}
```

When the lines are properly arranged, the program should display the following on the screen:

```
In 1492 Columbus sailed the ocean blue.
```

Rearrange the lines in the correct order. Test the program by entering it on the computer, compiling it, and running it.

2. 2.2 When the program in [Question 2.1](#) is saved to a file, what should the file be named?
3. 2.3 Complete the following program skeleton so it displays the message “Hello World” on the screen.


```
public class Hello
{
    public static void main(String[] args)
    {
        // Insert code here to complete the program
    }
}
```

4. 2.4 On paper, write a program that will display your name on the screen. Place a comment with today’s date at the top of the program. Test your program by entering, compiling, and running it.
5. 2.5 All Java source code filenames must end with .
 1. a semicolon
 2. *.class*
 3. *.java*
 4. none of the above
6. 2.6 Every Java application program must have .
 1. a method named `main`
 2. more than one class definition
 3. one or more comments

2.2 The `print` and `println` Methods, and the Java API

CONCEPT: The `print` and `println` methods are used to display text output. They are part of the Java API, which is a collection of prewritten classes and methods for performing specific operations.

In this section, you will learn how to write programs that produce output on the screen. The simplest type of output that a program can display on the screen is console output. *Console output* is merely plain text. When you display console output in a system that uses a graphical user interface, such as Windows or Mac OS, the output usually appears in a window similar to the one shown in [Figure 2-2](#).



```
C:\Users\Tony\Programs>javac Simple.java
C:\Users\Tony\Programs>java Simple
Programming is great fun!
C:\Users\Tony\Programs>
```

Figure 2-2 A console window

(Microsoft Corporation)



Displaying Console Output

The word *console* is an old computer term. It comes from the days when the operator of a large computer system interacted with the system by typing on a terminal that consisted of a simple screen and keyboard. This terminal was known as the *console*. The console screen, which displayed only text, was known as the standard output device. Today, the term *standard output device* typically refers to the device that displays console output.

Performing output in Java, as well as many other tasks, is accomplished by using the Java API. The term *API* stands for *Application Programmer Interface*. The API is a standard library of prewritten

classes for performing specific operations. These classes and their methods are available to all Java programs. The `print` and `println` methods are part of the API and provide ways for output to be displayed on the standard output device.

The program in [Code Listing 2-1](#) (`Simple.java`) uses the following statement to print a message on the screen:

```
System.out.println("Programming is great fun!");
```

SystemSystemSystemoutoutprint and printlnSystem, out, printprintln[Figure 2-3](#).




Figure 2-3 Relationship among the `System` class, the `out` object, and the `print` and `println` methods

Here is a brief summary of how it all works together:

- The `System` class is part of the Java API. It has member objects and methods for performing system-level operations, such as sending output to the console.
- The `out` object is a member of the `System` class. It provides methods for sending output to the screen.
- The `print` and `println` methods are members of the `out` object. They actually perform the work of writing characters on the screen.

This hierarchy explains why the statement that executes `println` is so long. The sequence `System.out.println` specifies that `println` is a member of `out`, which is a member of `System`.

NOTE:

The period that separates the names of the objects is pronounced “dot.” `System.out.println` is pronounced “system dot out dot print line.”

The value that is to be displayed on the screen is placed inside the parentheses. This value is known as an *argument*. For example, the following statement executes the `println` method using the string “King Arthur” as its argument. This will print “King Arthur” on the screen. (The quotation marks are not displayed.)

```
System.out.println("King Arthur");
```

An important thing to know about the `println` method is that after it displays its message, it advances the cursor to the beginning of the next line. The next item printed on the screen will begin in this position. For example, look at the program in [Code Listing 2-2](#).

Because each string is printed with separate `println` statements in [Code Listing 2-2](#), they appear on separate lines in the Program Output.

Code Listing 2-2 (`TwoLines.java`)

```
1  // This is another simple Java program.
2
3  public class TwoLines
4  {
5      public static void main(String[] args)
6      {
7          System.out.println("Programming is great fun!");
8          System.out.println("I can't get enough of it!");
9      }
10 }
```

Program Output

```
Programming is great fun!
I can't get enough of it!
```

The `print` Method

The `print` method, which is also part of the `System.out` object, serves a purpose similar to that of `println`—to display output on the screen. The `print` method, however, does not advance the cursor to the next line after its message is displayed. Look at [Code Listing 2-3](#).

Code Listing 2-3 (GreatFun.java)

```
1 // This is another simple Java program.
2
3 public class GreatFun
4 {
5     public static void main(String[] args)
6     {
7         System.out.print("Programming is ");
8         System.out.println("great fun!");
9     }
10 }
```

Program Output

```
Programming is great fun!
```

An important concept to understand about [Code Listing 2-3](#) is that, although the output is broken up into two programming statements, this program will still display the message on one line. The data that you send to the `print` method is displayed in a continuous stream. Sometimes this can produce less-than-desirable results. The program in [Code Listing 2-4](#) is an example.

Code Listing 2-4 (Unruly.java)

```
1 // An unruly printing program
2
3 public class Unruly
4 {
5     public static void main(String[] args)
6     {
7         System.out.print("These are our top sellers:");
8         System.out.print("Computer games");
9         System.out.print("Coffee");
10        System.out.println("Aspirin");
11    }
12 }
```

Program Output

```
These are our top sellers:Computer gamesCoffeeAspirin
```


The layout of the actual output looks nothing like the arrangement of the strings in the source code. First, even though the output is broken up into four lines in the source code (lines 7 through 10), it comes out on the screen as one line. Second, notice that some of the words that are displayed are not separated by spaces. The strings are displayed exactly as they are sent to the `print` method. If spaces are to be displayed, they must appear in the strings.

There are two ways to fix this program. The most obvious way is to use `println` methods instead of `print` methods. Another way is to use escape sequences to separate the output into different lines. An *escape sequence* starts with the backslash character (`\`), and is followed by one or more *control characters*. It allows you to control the way output is displayed by embedding commands within the string itself. The escape sequence that causes the output cursor to go to the next line is `\n`. [Code Listing 2-5](#) illustrates its use.

Code Listing 2-5 (Adjusted.java)

```
1  // A well adjusted printing program
2
3  public class Adjusted
4  {
5      public static void main(String[] args)
6      {
7          System.out.print("These are our top sellers:\n");
8          System.out.print("Computer games\nCoffee\n");
9          System.out.println("Aspirin");
10     }
11 }
```

Program Output

```
These are our top sellers:
Computer games
Coffee
Aspirin
```

The `\n` characters are called the newline escape sequence. When the `print` or `println` method encounters `\n` in a string, it does not print the `\n` characters on the screen, but interprets them as a special command to advance the output cursor to the next line. There are several other escape sequences as well. For instance, `\t` is the tab escape sequence. When `print` or `println` encounters it in a string, it causes the output cursor to advance to the next tab position. [Code Listing 2-6](#) shows it in use.

Code Listing 2-6 (Tabs.java)

```
1  // Another well-adjusted printing program
2
```

```

3 public class Tabs
4 {
5     public static void main(String[] args)
6     {
7         System.out.print("These are our top sellers:\n");
8         System.out.print("\tComputer games\n\tCoffee\n");
9         System.out.println("\tAspirin");
10    }
11 }

```

Program Output

```

These are our top sellers:
Computer games
Coffee
Aspirin

```

NOTE:

Although you have to type two characters to write an escape sequence, they are stored in memory as a single character.

[Table 2-2](#) lists the common escape sequences and describes them.

Table 2-2 Common escape sequences

Escape Sequence	Name	Description
\n	Newline	Advances the cursor to the next line for subsequent printing
\t	Horizontal tab	Causes the cursor to skip over to the next tab stop
\b	Backspace	Causes the cursor to back up, or move left, one position
\r	Return	Causes the cursor to go to the beginning of the current line, not the next line

\\	Backslash	Causes a backslash to be printed
\'	Single quote	Causes a single quotation mark to be printed
\"	Double quote	Causes a double quotation mark to be printed



WARNING!

Do not confuse the backslash (\) with the forward slash (/). An escape sequence will not work if you accidentally start it with a forward slash. Also, do not put a space between the backslash and the control character.



Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 2.7 The following program will not compile because the lines have been mixed up.

```
System.out.print("Success\n");
}
public class Success
{
System.out.print("Success\n");
public static void main(String[] args)
System.out.print("Success ");
}
// It's a mad, mad program.
System.out.println("\nSuccess");
{
```

When the lines are arranged properly, the program should display the following output on the screen:

Program Output

```
Success
Success Success
Success
```

Rearrange the lines in the correct order. Test the program by entering it on the computer, compiling it, and running it.

2. 2.8 Study the following program and show what it will print on the screen.

```
// The Works of Wolfgang
public class Wolfgang
{
    public static void main(String[] args)
    {
        System.out.print("The works of Wolfgang\ninclude ");
        System.out.print("the following");
        System.out.print("\nThe Turkish March ");
        System.out.print("and Symphony No. 40 ");
        System.out.println("in G minor.");
    }
}
```

3. 2.9 On paper, write a program that will display your name on the first line; your street address on the second line; your city, state, and ZIP code on the third line; and your telephone number on the fourth line. Place a comment with today's date at the top of the program. Test your program by entering, compiling, and running it.

2.3 Variables and Literals

CONCEPT: A variable is a named storage location in the computer's memory. A literal is a value that is written into the code of a program.

As you discovered in [Chapter 1](#), variables allow you to store and work with data in the computer's memory. Part of the job of programming is to determine how many variables a program will need and what types of data they will hold. The program in [Code Listing 2-7](#) is an example of a Java program with a variable.

Code Listing 2-7 (Variable.java)

```
1  // This program has a variable.
2
3  public class Variable
4  {
5      public static void main(String[] args)
6      {
7          int value;
8
9          value = 5;
10         System.out.print("The value is ");
11         System.out.println(value);
12     }
13 }
```

Program Output

The value is 5

Let's look more closely at this program. Here is line 7:

```
int value;
```



Declaring Variables

This is called a *variable declaration*. Variables must be declared before they can be used. A variable declaration tells the compiler the variable's name and the type of data it will hold. This line indicates the variable's name is value. The word `int` stands for integer, so value will only be used to hold integer numbers. Notice that variable declarations end with a semicolon. The next statement in this

program appears in line 9.

```
value = 5;
```

This is called an *assignment statement*. The equal sign is an operator that stores the value on its right (in this case 5) into the variable named on its left. After this line executes, the value variable will contain the value 5.

NOTE:

This line does not print anything on the computer screen. It runs silently behind the scenes.

Now look at lines 10 and 11:

```
System.out.print("The value is ");  
System.out.println(value);
```

The statement in line 10 sends the string literal “The value is ” to the `print` method. The statement in line 11 sends the name of the value variable to the `println` method. When you send a variable name to `print` or `println`, the variable’s contents are displayed. Notice there are no quotation marks around `value`. Look at what happens in [Code Listing 2-8](#).

Code Listing 2-8 (Variable2.java)

```
1  // This program has a variable.  
2  
3  public class Variable2  
4  {  
5      public static void main(String[] args)  
6      {  
7          int value;  
8  
9          value = 5;  
10         System.out.print("The value is ");  
11         System.out.println("value");  
12     }  
13 }
```

Program Output

```
The value is value
```

When double quotation marks are placed around the word `value` it becomes a string literal, not a variable name. When string literals are sent to `print` or `println`, they are displayed exactly as they appear inside the quotation marks.

Displaying Multiple Items with the + Operator

When the + operator is used with strings, it is known as the *string concatenation operator*. To concatenate means to append, so the string concatenation operator appends one string to another. For example, look at the following statement:

```
System.out.println("This is " + "one string.");
```

This statement will print:

```
This is one string.
```

The + operator produces a string that is the combination of the two strings used as its operands. You can also use the + operator to concatenate the contents of a variable to a string. The following code shows an example:

```
number = 5;  
System.out.println("The value is " + number);
```

The second line uses the + operator to concatenate the contents of the `number` variable with the string “The value is ”. Although `number` is not a string, the + operator converts its value to a string and then concatenates that value with the first string. The output that will be displayed is:

```
The value is 5
```

Sometimes the argument you use with `print` or `println` is too long to fit on one line in your program code. However, a string literal cannot begin on one line and end on another. For example, the following will cause an error:

```
// This is an error!  
System.out.println("Enter a value that is greater than zero  
                    and less than 10." );
```

You can remedy this problem by breaking the argument up into smaller string literals, and then using the string concatenation operator to spread them out over more than one line. Here is an example:

```
System.out.println("Enter a value that is " +  
                  "greater than zero and less " +  
                  "than 10." );
```

In this statement, the argument is broken up into three strings and joined using the + operator. The following example shows the same technique used when the contents of a variable are part of the concatenation:

```
sum = 249;  
System.out.println("The sum of the three " +  
                  "numbers is " + sum);
```

Be Careful with Quotation Marks

As shown in [Code Listing 2-8](#), placing quotation marks around a variable name changes the program's results. In fact, placing double quotation marks around anything that is not intended to be a string literal will create an error of some type. For example, in [Code Listings 2-7](#) and [2-8](#), the number 5 was assigned to the variable `value`. It would have been an error to perform the assignment this way:

```
value = "5"; // Error!
```

In this statement, 5 is no longer an integer, but a string literal. Because `value` was declared as an integer variable, you can only store integers in it. In other words, 5 and “5” are not the same thing.

The fact that numbers can be represented as strings frequently confuses students who are new to programming. Just remember that strings are intended for humans to read. They are to be printed on computer screens or paper. Numbers, however, are intended primarily for mathematical operations. You cannot perform math on strings, and before numbers can be displayed on the screen, first they must be converted to strings. (Fortunately, `print` and `println` handle the conversion automatically when you send numbers to them.) Don't fret if this still bothers you. Later in this chapter, we will shed more light on the differences among numbers, characters, and strings by discussing their internal storage.

More about Literals

A literal is a value that is written in the code of a program. Literals are commonly assigned to variables or displayed. [Code Listing 2-9](#) contains both literals and a variable.

Code Listing 2-9 (Literals.java)

[illegible]

Program Output

Today we sold 20 bushels of apples.

Of course, the variable in this program is `apples`. It is declared as an integer. [Table 2-3](#) shows a list of the literals found in the program.

Table 2-3 Literals

Literal	Type of Literal
20	Integer literal
“Today we sold ”	String literal
“ bushels of apples.”	String literal

Identifiers

An *identifier* is a programmer-defined name that represents some element of a program. Variable names and class names are examples of identifiers. You may choose your own variable names and class names in Java, as long as you do not use any of the Java key words. The *key words* make up the core of the language and each has a specific purpose. [Table 1-3](#) in [Chapter 1](#) and Appendix C (available on the book’s companion Web site) show a complete list of Java key words.

You should always choose names for your variables that give an indication of what they are used for. You may be tempted to declare variables with names like this:

```
int x;
```

The rather nondescript name, `x`, gives no clue as to what the variable’s purpose is. Here is a better example.

```
int itemsOrdered;
```

The name `itemsOrdered` gives anyone reading the program an idea of what the variable is used for. This method of coding helps produce *self-documenting programs*, which means you get an

understanding of what the program is doing just by reading its code. Because real-world programs usually have thousands of lines of code, it is important that they be as self-documenting as possible.

You have probably noticed the mixture of uppercase and lowercase letters in the name `itemsOrdered`. Although all of Java’s key words must be written in lowercase, you may use uppercase letters in variable names. The reason the `O` in `itemsOrdered` is capitalized is to improve readability. Normally “items ordered” is used as two words. Variable names cannot contain spaces, however, so the two words must be combined. When “items” and “ordered” are stuck together, you get a variable declaration like this:

```
int itemsordered;
```

Capitalization of the letter `O` makes `itemsOrdered` easier to read. Typically, variable names begin with a lowercase letter, and after that, the first letter of each individual word that makes up the variable name is capitalized.

The following are some specific rules that must be followed with all identifiers:

- The first character must be one of the letters a–z or A–Z, an underscore (`_`), or a dollar sign (`$`).
- After the first character, you may use the letters a–z or A–Z, the digits 0–9, underscores (`_`), or dollar signs (`$`).
- Uppercase and lowercase characters are distinct. This means `itemsOrdered` is not the same as `itemsordered`.
- Identifiers cannot include spaces.

NOTE:

Although the `$` is a legal identifier character, it is normally used for special purposes. So, don’t use it in your variable names.

[Table 2-4](#) shows a list of variable names and tells whether each is legal or illegal in Java.

Table 2-4 Some variable names

Variable Name	Legal or Illegal?
<code>dayOfWeek</code>	Legal

3dGraph	Illegal because identifiers cannot begin with a digit
june1997	Legal
mixture#3	Illegal because identifiers may use only alphabetic letters, digits, - underscores, or dollar signs
week day	Illegal because identifiers cannot contain spaces

Class Names

As mentioned before, it is standard practice to begin variable names with a lowercase letter, and then capitalize the first letter of each subsequent word that makes up the name. It is also a standard practice to capitalize the first letter of a class name, as well as the first letter of each subsequent word it contains. This helps differentiate the names of variables from the names of classes. For example, `payRate` would be a variable name, and `Employee` would be a class name.

Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 2.10 Examine the following program.

```
// This program uses variables and literals.
public class BigLittle
{
    public static void main(String[] args)
    {
        int little;
        int big;
        little = 2;
        big = 2000;
        System.out.println("The little number is " + little);
        System.out.println("The big number is " + big);
    }
}
```

List the variables and literals found in the program.

2. 2.11 What will the following program display on the screen?

```
public class CheckPoint
{
    public static void main(String[] args)
    {
        int number;
        number = 712;
        System.out.println("The value is " + "number");
    }
}
```

2.4 Primitive Data Types

CONCEPT: There are many different types of data. Variables are classified according to their data type, which determines the kind of data that may be stored in them.

Computer programs collect pieces of data from the real world and manipulate them in various ways. There are many different types of data. In the realm of numeric data, for example, there are whole and fractional numbers, negative and positive numbers, and numbers so large and others so small that they don't even have a name. Then there is textual information. Names and addresses, for instance, are stored as strings of characters. When you write a program you must determine what types of data it is likely to encounter.

Each variable has a *data type*, which is the type of data that the variable can hold. Selecting the proper data type is important because a variable's data type determines the amount of memory the variable uses, and the way the variable formats and stores data. It is important to select a data type that is appropriate for the type of data that your program will work with. If you are writing a program to calculate the number of miles to a distant star, you need variables that can hold very large numbers. If you are designing software to record microscopic dimensions, you need variables that store very small and precise numbers. If you are writing a program that must perform thousands of intensive calculations, you want variables that can be processed quickly. The data type of a variable determines all of these factors.

[Table 2-5](#) shows all of the Java *primitive data types* for holding numeric data.

Table 2-5 Primitive data types for numeric data

Data Type	Size	Range
byte	1 byte	Integers in the range of -128 to +127
short	2 bytes	Integers in the range of -32,768 to +32,767

int	4 bytes	Integers in the range of $-2,147,483,648$ to $+2,147,483,647$
long	8 bytes	Integers in the range of $-9,223,372,036,854,775,808$ to $+9,223,372,036,854,775,807$
float	4 bytes	Floating-point numbers in the range of $\pm 3.4 \times 10^{-38}$ to $\pm 3.4 \times 10^{38}$, with 7 digits of accuracy
double	8 bytes	Floating-point numbers in the range of $\pm 1.7 \times 10^{-308}$ to $\pm 1.7 \times 10^{308}$, with 15 digits of accuracy

The words listed in the left column of [Table 2-5](#) are the key words that you use in variable declarations. A variable declaration takes the following general format:

```
DataType VariableName;
```

```
DataTypeVariableName
```

```
byte inches;
int speed;
short month;
float salesCommission;
double distance;
```

The size column in [Table 2-5](#) shows the number of bytes that a variable of each of the data types uses. For example, an `int` variable uses 4 bytes, and a `double` variable uses 8 bytes. The range column shows the ranges of numbers that may be stored in variables of each data type. For example, an `int` variable can hold numbers from $-2,147,483,648$ up to $+2,147,483,647$. One of the appealing characteristics of the Java language is that the sizes and ranges of all the primitive data types are the same on all computers.

NOTE:

These data types are called “primitive” because you cannot use them to create objects. Recall from [Chapter 1](#)’s discussion on object-oriented programming that an object has attributes and methods. With the primitive data types, you can only create variables, and a variable can only be used to hold a single value. Such variables do not have attributes or methods.

The Integer Data Types

The first four data types listed in [Table 2-5](#), byte, int, short, and long, are all integer data types. An integer variable can hold whole numbers such as 7, 125, -14, and 6928. The program in [Code Listing 2-10](#) shows several variables of different integer data types being used.

Code Listing 2-10 (IntegerVariables.java)

```
1  // This program has variables of several of the integer types.
2
3  public class IntegerVariables
4  {
5      public static void main(String[] args)
6      {
7          int checking; // Declare an int variable named checking.
8          byte miles; // Declare a byte variable named miles.
9          short minutes; // Declare a short variable named minutes.
10         long days; // Declare a long variable named days.
11
12         checking = -20;
13         miles = 105;
14         minutes = 120;
15         days = 189000;
16         System.out.println("We have made a journey of " + miles +
17                             " miles.");
18         System.out.println("It took us " + minutes + " minutes.");
19         System.out.println("Our account balance is $" + checking);
20         System.out.println("About " + days + " days ago Columbus " +
21                             "stood on this spot.");
22     }
23 }
```

Program Output

```
We have made a journey of 105 miles.
It took us 120 minutes.
Our account balance is $-20
About 189000 days ago Columbus stood on this spot.
```

In most programs you will need more than one variable of any given data type. If a program uses three integers, length, width, and area, they could be declared separately, as follows:

```
int length;
int width;
int area;
```

It is easier, however, to combine the three variable declarations:

```
int length, width, area;
```

You can declare several variables of the same type, simply by separating their names with commas.

Integer Literals

When you write an integer literal in your program code, Java assumes it to be of the `int` data type. For example, in [Code Listing 2-10](#), the literals `-20`, `105`, `120`, and `189000` are all treated as `int` values. You can force an integer literal to be treated as a long, however, by suffixing it with the letter `L`. For example, the value `57L` would be treated as a long. Although you can use either an uppercase or a lowercase `L`, it is advisable to use the uppercase `L` because the lowercase `l` looks too much like the number `1`.



You cannot embed commas in numeric literals. For example, the following statement will cause an error:

```
number = 1,257,649;    // ERROR!
```

This statement must be written as:

```
number = 1257649;      // Correct.
```

Floating-Point Data Types

Whole numbers are not adequate for many jobs. If you are writing a program that works with dollar amounts or precise measurements, you need a data type that allows fractional values. In programming terms, these are called *floating-point* numbers. Values such as `1.7` and `-45.316` are floating-point numbers.

In Java there are two data types that can represent floating-point numbers. They are `float` and `double`. The `float` data type is considered a single precision data type. It can store a floating-point number with 7 digits of accuracy. The `double` data type is considered a double precision data type. It can store a floating-point number with 15 digits of accuracy. The `double` data type uses twice as much memory as the `float` data type, however. A `float` variable occupies 4 bytes of memory, whereas a `double` variable uses 8 bytes.

[Code Listing 2-11](#) shows a program that uses three `double` variables.

Code Listing 2-11 (Sale.java)

```
1  // This program demonstrates the double data type.
2
3  public class Sale
```



```

4  {
5      public static void main(String[] args)
6      {
7          double price, tax, total;
8
9          price = 29.75;
10         tax = 1.76;
11         total = 31.51;
12         System.out.println("The price of the item " +
13                             "is " + price);
14         System.out.println("The tax is " + tax);
15         System.out.println("The total is " + total);
16     }
17 }

```

Program Output

```

The price of the item is 29.75
The tax is 1.76
The total is 31.51

```

Floating-Point Literals

When you write a floating-point literal in your program code, Java assumes it to be of the `double` data type. For example, in [Code Listing 2-11](#), the literals 29.75, 1.76, and 31.51 are all treated as `double` values. Because of this, a problem can arise when assigning a floating-point literal to a `float` variable. Java is a *strongly typed language*, which means that it only allows you to store values of compatible data types in variables. A `double` value is not compatible with a `float` variable because a `double` can be much larger or much smaller than the allowable range for a `float`. As a result, code such as the following will cause an error:

```

float number;
number = 23.5;           // Error!

```

You can force a `double` literal to be treated as a `float`, however, by suffixing it with the letter F or f. The preceding code can be rewritten in the following manner to prevent an error:

```

float number;
number = 23.5F;          // This will work.

```



If you are working with literals that represent dollar amounts, remember that you cannot embed currency symbols (such as \$) or commas in the literal. For example, the following statement will cause an error:

```
grossPay = $1,257.00;      // ERROR!
```

This statement must be written as:

```
grossPay = 1257.00;      // Correct.
```

Scientific and E Notation

Floating-point literals can be represented in scientific notation. Take the number 47,281.97. In scientific notation this number is 4.728197×10^4 . (10⁴ is equal to 10,000, and $4.728197 \times 10,000$ is 47,281.97.)

Java uses E notation to represent values in scientific notation. In E notation, the number 4.728197×10^4 would be 4.728197E4. [Table 2-6](#) shows other numbers represented in scientific and E notation.

Table 2-6 Floating-point representations

Decimal Notation	Scientific Notation	E Notation

247.91	2.4791×10^2	2.4791E2

0.00072	7.2×10^{-4}	7.2E−4

2,900,000	2.9×10^6	2.9E6

 NOTE:

The E can be uppercase or lowercase.

[Code Listing 2-12](#) demonstrates the use of floating-point literals expressed in E notation.

Code Listing 2-12 (SunFacts.java)

```
1 // This program uses E notation.
```

```

2
3 public class SunFacts
4 {
5     public static void main(String[] args)
6     {
7         double distance, mass;
8
9         distance = 1.495979E11;
10        mass = 1.989E30;
11        System.out.println("The sun is " + distance +
12                            " meters away.");
13        System.out.println("The sun's mass is " + mass +
14                            " kilograms.");
15    }
16 }

```

Program Output

The sun is 1.495979E11 meters away.
The sun's mass is 1.989E30 kilograms.

The boolean Data Type

The `boolean` data type allows you to create variables that may hold one of two possible values: `true` or `false`. [Code Listing 2-13](#) demonstrates the declaration and assignment of a `boolean` variable.

Code Listing 2-13 (TrueFalse.java)

```

1 // A program for demonstrating boolean variables
2
3 public class TrueFalse
4 {
5     public static void main(String[] args)
6     {
7         boolean bool;
8
9         bool = true;
10        System.out.println(bool);
11        bool = false;
12        System.out.println(bool);
13    }
14 }

```

Program Output

true

false

Variables of the `boolean` data type are useful for evaluating conditions that are either true or false. You will not be using them until [Chapter 3](#), however, so for now just remember the following things:

- `boolean` variables may hold only the value `true` or `false`.
- The contents of a `boolean` variable may not be copied to a variable of any type other than `boolean`.

The char Data Type

The `char` data type is used to store characters. A variable of the `char` data type can hold one character at a time. Character literals are enclosed in *single quotation marks*. The program in [Code Listing 2-14](#) uses a `char` variable. The character literals 'A' and 'B' are assigned to the variable.

Code Listing 2-14 (Letters.java)

```
1  // This program demonstrates the char data type.
2
3  public class Letters
4  {
5      public static void main(String[] args)
6      {
7          char letter;
8
9          letter = 'A';
10         System.out.println(letter);
11         letter = 'B';
12         System.out.println(letter);
13     }
14 }
```

Program Output

A
B

It is important that you do not confuse character literals with string literals, which are enclosed in double quotation marks. String literals cannot be assigned to `char` variables.

Unicode

Characters are internally represented by numbers. Each printable character, as well as many non-printable characters, is assigned a unique number. Java uses Unicode, which is a set of numbers that are used as codes for representing characters. Each Unicode number requires two bytes of memory, so `char` variables occupy two bytes. When a character is stored in memory, it is actually the numeric code that is stored. When the computer is instructed to print the value on the screen, it displays the character that corresponds with the numeric code.

You may want to refer to Appendix A, available on the book's companion Web site (at www.pearsonhighered.com/gaddis), which shows a portion of the Unicode character set. Notice that the number 65 is the code for A, 66 is the code for B, and so on. [Code Listing 2-15](#) demonstrates that when you work with characters, you are actually working with numbers.

Code Listing 2-15 (Letters2.java)

```
1  // This program demonstrates the close relationship between
2  // characters and integers.
3
4  public class Letters2
5  {
6      public static void main(String[] args)
7      {
8          char letter;
9
10         letter = 65;
11         System.out.println(letter);
12         letter = 66;
13         System.out.println(letter);
14     }
15 }
```

Program Output

A
B

[Figure 2-4](#) illustrates that when you think of the characters A, B, and C being stored in memory, it is really the numbers 65, 66, and 67 that are stored.



These characters are stored in memory as...




Figure 2-4 Characters and how they are stored in memory

Variable Assignment and Initialization

As you have already seen in several examples, a value is put into a variable with an *assignment statement*. For example, the following statement assigns the value 12 to the variable `unitsSold`:

```
unitsSold = 12;
```

The `=` symbol is called the assignment operator. Operators perform operations on data. The data that operators work with are called operands. The assignment operator has two operands. In the statement above, the operands are `unitsSold` and `12`.

In an assignment statement, the name of the variable receiving the assignment must appear on the left side of the operator, and the value being assigned must appear on the right side. The following statement is incorrect:

```
12 = unitsSold;           // ERROR!
```

The operand on the left side of the `=` operator must be a variable name. The operand on the right side of the `=` symbol must be an expression that has a value. The assignment operator takes the value of the right operand and puts it in the variable identified by the left operand. Assuming that `length` and `width` are both `int` variables, the following code illustrates that the assignment operator's right operand may be a literal or a variable:

```
length = 20;  
width = length;
```

It is important to note that the assignment operator only changes the contents of its left operand. The second statement assigns the value of the `length` variable to the `width` variable. After the statement has executed, `length` still has the same value, 20.

You may also assign values to variables as part of the declaration statement. This is known as *initialization*. [Code Listing 2-16](#) shows how it is done.

The variable declaration statement in this program is in line 7:

```
int month = 2, days = 28;
```

Code Listing 2-16 (`Initialize.java`)

```

1 // This program shows variable initialization.
2
3 public class Initialize
4 {
5     public static void main(String[] args)
6     {
7         int month = 2, days = 28;
8
9         System.out.println("Month " + month + " has " +
10                             days + " days.");
11     }
12 }

```

Program Output

Month 2 has 28 days.

This statement declares the `month` variable and initializes it with the value 2, and declares the `days` variable and initializes it with the value 28. As you can see, this simplifies the program and reduces the number of statements that must be typed by the programmer. Here are examples of other declaration statements that perform initialization:

```

double payRate = 25.52;
float interestRate = 12.9F;
char stockCode = 'D';
int customerNum = 459;

```

Of course, there are always variations on a theme. Java allows you to declare several variables and initialize only some of them. Here is an example of such a declaration:

```
int flightNum = 89, travelTime, departure = 10, distance;
```

The variable `flightNum` is initialized to 89 and `departure` is initialized to 10. The `travelTime` and `distance` variables remain uninitialized.



WARNING!

When a variable is declared inside a method, it must have a value stored in it before it can be used. If the compiler determines that the program might be using such a variable before a value has been stored in it, an error will occur. You can avoid this type of error by initializing the variable with a value.

Variables Hold Only One Value at a Time

Remember, a variable can hold only one value at a time. When you assign a new value to a variable,

the new value takes the place of the variable's previous contents. For example, look at the following code.

```
int x = 5;
System.out.println(x);
x = 99;
System.out.println(x);
```

In this code, the variable `x` is initialized with the value 5 and its contents are displayed. Then the variable is assigned the value 99. This value overwrites the value 5 that was previously stored there. The code will produce the following output:

```
5
99
```



MyProgrammingLab™ www.myprogramminglab.com

1. 2.12 Which of the following are illegal variable names and why?

```
x
99bottles
july97
theSalesFigureForFiscalYear98
r&d
grade_report
```

2. 2.13 Is the variable name `Sales` the same as `sales`? Why or why not?

3. 2.14 Refer to the Java primitive data types listed in [Table 2-5](#) for this question.

1. If a variable needs to hold whole numbers in the range 32 to 6,000, what primitive data type would be best?

2. If a variable needs to hold whole numbers in the range -40,000 to +40,000, what primitive data type would be best?

3. Which of the following literals use more memory? `22.1` or `22.1F`?

4. 2.15 How would the number 6.31×10^{17} be represented in E notation?

5. 2.16 A program declares a `float` variable named `number`, and the following statement causes an error. What can be done to fix the error?

```
number = 7.4;
```

6. 2.17 What values can `boolean` variables hold?

7. 2.18 Write statements that do the following:

1. Declare a `char` variable named `letter`.
2. Assign the letter A to the `letter` variable.
3. Display the contents of the `letter` variable.

8. 2.19 What are the Unicode codes for the characters ‘C’, ‘F’, and ‘W’? (You may need to refer to Appendix A on the book’s companion Web site, at www.pearsonhighered.com/gaddis.)

9. 2.20 Which is a character literal, `'B'` or `"B"`?

10. 2.21 What is wrong with the following statement?

```
char letter = "Z";
```

2.5 Arithmetic Operators

CONCEPT: There are many operators for manipulating numeric values and performing arithmetic operations.

Java offers a multitude of operators for manipulating data. Generally, there are three types of operators: *unary*, *binary*, and *ternary*. These terms reflect the number of operands an operator requires.

Unary operators require only a single operand. For example, consider the following expression:

`-5`



Simple Math Expressions

Of course, we understand this represents the value negative five. We can also apply the operator to a variable, as follows:

`-number`

This expression gives the negative of the value stored in `number`. The minus sign, when used this way, is called the *negation operator*. Because it requires only one operand, it is a unary operator.

Binary operators work with two operands. The assignment operator is in this category. Ternary operators, as you may have guessed, require three operands. Java has only one ternary operator, which is discussed in [Chapter 3](#).

Arithmetic operations are very common in programming. [Table 2-7](#) shows the arithmetic operators in Java.

Table 2-7 Arithmetic operators

Operator	Meaning	Type	Example
<code>+</code>	Addition	Binary	<code>total = cost + tax;</code>

-	Subtraction	Binary <code>cost = total - tax;</code>
*	Multiplication	Binary <code>tax = cost * rate;</code>
/	Division	Binary <code>salePrice = original / 2;</code>
%	Modulus	Binary <code>remainder = value % 3;</code>

Each of these operators works as you probably expect. The addition operator returns the sum of its two operands. Here are some example statements that use the addition operator:

```
amount = 4 + 8;           // Assigns 12 to amount
total = price + tax;      // Assigns price + tax to total
number = number + 1;      // Assigns number + 1 to number
```

The subtraction operator returns the value of its right operand subtracted from its left operand. Here are some examples:

```
temperature = 112 - 14;   // Assigns 98 to temperature
sale = price - discount;  // Assigns price - discount to sale
number = number - 1;      // Assigns number - 1 to number
```

The multiplication operator returns the product of its two operands. Here are some examples:

```
markUp = 12 * 0.25;       // Assigns 3 to markUp
commission = sales * percent; // Assigns sales * percent to commission
population = population * 2; // Assigns population * 2 to population
```

The division operator returns the quotient of its left operand divided by its right operand. Here are some examples:

```
points = 100 / 20;         // Assigns 5 to points
teams = players / maxEach; // Assigns players / maxEach to teams
half = number / 2;         // Assigns number / 2 to half
```

The modulus operator returns the remainder of a division operation involving two integers. The following statement assigns 2 to `leftOver`:

```
leftOver = 17 % 3;
```

Situations arise where you need to get the remainder of a division. Computations that detect odd numbers or are required to determine how many items are left over after division use the modulus operator.

The program in [Code Listing 2-17](#) demonstrates some of these operators used in a simple payroll

calculation.

Code Listing 2-17 (Wages.java)

```
1  // This program calculates hourly wages plus overtime.
2
3  public class Wages
4  {
5      public static void main(String[] args)
6      {
7          double regularWages;           // The calculated regular wages.
8          double basePay = 25;           // The base pay rate.
9          double regularHours = 40;      // The hours worked less overtime.
10         double overtimeWages;          // Overtime wages
11         double overtimePay = 37.5;     // Overtime pay rate
12         double overtimeHours = 10;     // Overtime hours worked
13         double totalWages;             // Total wages
14
15         regularWages = basePay * regularHours;
16         overtimeWages = overtimePay * overtimeHours;
17         totalWages = regularWages + overtimeWages;
18         System.out.println("Wages for this week are $" +
19                             totalWages);
20     }
21 }
```

Program Output

Wages for this week are \$1375.0

[Code Listing 2-17](#) calculates the total wages an hourly paid worker earned in one week. As mentioned in the comments, there are variables for regular wages, base pay rate, regular hours worked, overtime wages, overtime pay rate, overtime hours worked, and total wages.

Line 15 in the program multiplies `basePay` times `regularHours` and stores the result, which is 1000, in `regularWages`:

```
regularWages = basePay * regularHours;
```

Line 16 multiplies `overtimePay` times `overtimeHours` and stores the result, which is 375, in `overtimeWages`:

```
overtimeWages = overtimePay * overtimeHours;
```

Line 17 adds the regular wages and the overtime wages and stores the result, 1375, in `totalWages`:

```
totalWages = regularWages + overtimeWages;
```

The `println` statement in lines 18 and 19 displays the message on the screen reporting the week's

wages.

Integer Division

When both operands of the division operator are integers, the operator will perform *integer division*. This means the result of the division will be an integer as well. If there is a remainder, it will be discarded. For example, look at the following code:

```
double number;  
number = 5 / 2;
```

This code divides 5 by 2 and assigns the result to the `number` variable. What value will be stored in `number`? You would probably assume that 2.5 would be stored in `number` because that is the result your calculator shows when you divide 5 by 2; however, that is not what happens when the previous Java code is executed. Because the numbers 5 and 2 are both integers, the fractional part of the result will be thrown away, or *truncated*. As a result, the value 2 will be assigned to the `number` variable.

In the previous code, it doesn't matter that `number` is declared as a `double` because the fractional part of the result is discarded before the assignment takes place. In order for a division operation to return a floating-point value, one of the operands must be of a floating-point data type. For example, the previous code could be written as follows:

```
double number;  
number = 5.0 / 2;
```

In this code, 5.0 is treated as a floating-point number, so the division operation will return a floating-point number. The result of the division is 2.5.

Operator Precedence

It is possible to build mathematical expressions with several operators. The following statement assigns the sum of 17, `x`, 21, and `y` to the variable `answer`:

```
answer = 17 + x + 21 + y;
```

Some expressions are not that straightforward, however. Consider the following statement:

```
outcome = 12 + 6 / 3;
```

What value will be stored in `outcome`? The 6 is used as an operand for both the addition and division operators. The `outcome` variable could be assigned either 6 or 14, depending on when the division takes place. The answer is 14 because the division operator has higher *precedence* than the addition operator.

Mathematical expressions are evaluated from left to right. When two operators share an operand, the

operator with the highest precedence works first. Multiplication and division have higher precedence than addition and subtraction, so the statement above works like this:

- 1. 6 is divided by 3, yielding a result of 2
- 2. 12 is added to 2, yielding a result of 14

It could be diagrammed as shown in [Figure 2-5](#).




Figure 2-5 Precedence illustrated

[Table 2-8](#) shows the precedence of the arithmetic operators. The operators at the top of the table have higher precedence than the ones below them.

Table 2-8 Precedence of arithmetic operators (highest to lowest)

Highest Precedence → - (unary negation)

* / %

Lowest Precedence → + -

The multiplication, division, and modulus operators have the same precedence. The addition and subtraction operators have the same precedence. If two operators sharing an operand have the same precedence, they work according to their *associativity*. Associativity is either *left to right* or *right to left*. [Table 2-9](#) shows the arithmetic operators and their associativity.

Table 2-9 Associativity of arithmetic operators

Operator	Associativity
- (unary negation)	Right to left
* / %	Left to right
+ -	Left to right

[Table 2-10](#) shows some expressions and their values.

Table 2-10 Some expressions and their values

Expression	Value
$5 + 2 * 4$	13
$10 / 2 - 3$	2
$8 + 12 * 2 - 4$	28
$4 + 17 \% 2 - 1$	4
$6 - 3 * 2 + 7 - 1$	6

Grouping with Parentheses

Parts of a mathematical expression may be grouped with parentheses to force some operations to be performed before others. In the statement below, the sum of a, b, c, and d is divided by 4.0.

```
average = (a + b + c + d) / 4.0;
```

Without the parentheses, however, d would be divided by 4 and the result added to a, b, and c. [Table 2-11](#) shows more expressions and their values.

Table 2-11 More expressions and their values

Expression	Value
(5 + 2) * 4	28
10 / (5 - 3)	5
8 + 12 * (6 - 2)	56
(4 + 17) % 2 - 1	0
(6 - 3) * (2 + 7) / 3	9

In the Spotlight:



Calculating Percentages and Discounts

Determining percentages is a common calculation in computer programming. Although the % symbol is used in general mathematics to indicate a percentage, most programming languages (including Java) do not use the % symbol for this purpose. In a program, you have to convert a percentage to a

floating-point number, just as you would if you were using a calculator. For example, 50 percent would be written as 0.5 and 2 percent would be written as 0.02.

Let's look at an example. Suppose you earn \$6,000 per month and you are allowed to contribute a portion of your gross monthly pay to a retirement plan. You want to determine the amount of your pay that will go into the plan if you contribute 5 percent, 8 percent, or 10 percent of your gross wages. To make this determination you write a program like the one shown in [Code Listing 2-18](#).

Code Listing 2-18 (Contribution.java)

```
1  // This program calculates the amount of pay that
2  // will be contributed to a retirement plan if 5%,
3  // 8%, or 10% of monthly pay is withheld.
4
5  public class Contribution
6  {
7      public static void main(String[] args)
8      {
9          // Variables to hold the monthly pay and
10         // the amount of contribution.
11         double monthlyPay = 6000.0;
12         double contribution;
13
14         // Calculate and display a 5% contribution.
15         contribution = monthlyPay * 0.05;
16         System.out.println("5 percent is $" +
17                             contribution +
18                             " per month.");
19
20         // Calculate and display an 8% contribution.
21         contribution = monthlyPay * 0.08;
22         System.out.println("8 percent is $" +
23                             contribution +
24                             " per month.");
25
26         // Calculate and display a 10% contribution.
27         contribution = monthlyPay * 0.1;
28         System.out.println("10 percent is $" +
29                             contribution +
30                             " per month.");
31     }
32 }
```

Program Output

```
5 percent is $300.0 per month.
8 percent is $480.0 per month.
10 percent is $600.0 per month.
```

Lines 11 and 12 declare two variables: `monthlyPay` and `contribution`. The `monthlyPay` variable,

which is initialized with the value `6000.0`, holds the amount of your monthly pay. The `contribution` variable will hold the amount of a contribution to the retirement plan.

The statements in lines 15 through 18 calculate and display 5 percent of the monthly pay. The calculation is done in line 15, where the `monthlyPay` variable is multiplied by 0.05. The result is assigned to the `contribution` variable, which is then displayed by the statement in lines 16 through 18.

Similar steps are taken in lines 21 through 24, which calculate and display 8 percent of the monthly pay, and lines 27 through 30, which calculate and display 10 percent of the monthly pay.

Calculating a Percentage Discount

Another common calculation is determining a percentage discount. For example, suppose a retail business sells an item that is regularly priced at \$59, and is planning to have a sale where the item's price will be reduced by 20 percent. You have been asked to write a program to calculate the sale price of the item.

To determine the sale price you perform two calculations:

- First, you get the amount of the discount, which is 20 percent of the item's regular price.
- Second, you subtract the discount amount from the item's regular price. This gives you the sale price.

[Code Listing 2-19](#) shows how this is done in Java.

Code Listing 2-19 (Discount.java)

```
1  // This program calculates the sale price of an
2  // item that is regularly priced at $59, with
3  // a 20 percent discount subtracted.
4
5  public class Discount
6  {
7      public static void main(String[] args)
8      {
9          // Variables to hold the regular price, the
10         // amount of a discount, and the sale price.
11         double regularPrice = 59.0;
12         double discount;
13         double salePrice;
14
15         // Calculate the amount of a 20% discount.
16         discount = regularPrice * 0.2;
17
18         // Calculate the sale price by subtracting
```

```

19      // the discount from the regular price.
20      salePrice = regularPrice - discount;
21
22      // Display the results.
23      System.out.println("Regular price: $" + regularPrice);
24      System.out.println("Discount amount $" + discount);
25      System.out.println("Sale price: $" + salePrice);
26  }
27  }

```

Program Output

```

Regular price: $59.0
Discount amount $11.8
Sale price: $47.2

```

Lines 11 through 13 declare three variables. The `regularPrice` variable holds the item's regular price, and is initialized with the value `59.0`. The `discount` variable will hold the amount of the discount once it is calculated. The `salePrice` variable will hold the item's sale price.

Line 16 calculates the amount of the 20 percent discount by multiplying `regularPrice` by `0.2`. The result is stored in the `discount` variable. Line 20 calculates the sale price by subtracting `discount` from `regularPrice`. The result is stored in the `salePrice` variable. The statements in lines 23 through 25 display the item's regular price, the amount of the discount, and the sale price.

The Math Class

The Java API provides a class named `Math`, which contains numerous methods that are useful for performing complex mathematical operations. In this section we will briefly look at the `Math.pow` and `Math.sqrt` methods.

The Math.pow Method

In Java, raising a number to a power requires the `Math.pow` method. Here is an example of how the `Math.pow` method is used:

```
result = Math.pow(4.0, 2.0);
```

The method takes two `double` arguments. It raises the first argument to the power of the second argument, and returns the result as a `double`. In this example, `4.0` is raised to the power of `2.0`. This statement is equivalent to the following algebraic statement:

$$result = 4^2$$

Here is another example of a statement using the `Math.pow` method. It assigns 3 times 63 to `x`:

```
x = 3 * Math.pow(6.0, 3.0);
```

And the following statement displays the value of 5 raised to the power of 4:

```
System.out.println(Math.pow(5.0, 4.0));
```

The `Math.sqrt` Method

The `Math.sqrt` method accepts a `double` value as its argument and returns the square root of the value. Here is an example of how the method is used:

```
result = Math.sqrt(9.0);
```

In this example the value 9.0 is passed as an argument to the `Math.sqrt` method. The method will return the square root of 9.0, which is assigned to the `result` variable. The following statement shows another example. In this statement the square root of 25.0 (which is 5.0) is displayed on the screen:

```
System.out.println(Math.sqrt(25.0));
```



For more information about the `Math` class, see Appendix F, available on the book's companion Web site at www.pearsonhighered.com/gaddis.



MyProgrammingLab™ www.myprogramminglab.com

1. 2.22 Complete the following table by writing the value of each expression in the Value column.

Expression	Value

$6 + 3 * 5$	

$12 / 2 - 4$	

$$9 + 14 * 2 - 6$$

$$5 + 19 \% 3 - 1$$

$$(6 + 2) * 3$$

$$14 / (11 - 4)$$

$$9 + 12 * (8 - 3)$$

2. 2.23 Is the division statement in the following code an example of integer division or floating-point division? What value will be stored in `portion`?

```
double portion;  
portion = 70 / 3;
```

2.6 Combined Assignment Operators

CONCEPT: The combined assignment operators combine the assignment operator with the arithmetic operators.

Quite often, programs have assignment statements of the following form:

```
x = x + 1;
```

On the right side of the assignment operator, 1 is added to `x`. The result is then assigned to `x`, replacing the value that was previously there. Effectively, this statement adds 1 to `x`. Here is another example:

```
balance = balance + deposit;
```

Assuming that `balance` and `deposit` are variables, this statement assigns the value of `balance + deposit` to `balance`. The effect of this statement is that `deposit` is added to the value stored in `balance`. Here is another example:

```
balance = balance - withdrawal;
```

Assuming that `balance` and `withdrawal` are variables, this statement assigns the value of `balance - withdrawal` to `balance`. The effect of this statement is that `withdrawal` is subtracted from the value stored in `balance`.

If you have not seen these types of statements before, they might cause some initial confusion because the same variable name appears on both sides of the assignment operator. [Table 2-12](#) shows other examples of statements written this way.

Table 2-12 Various assignment statements (assume `x` = 6 in each statement)

Statement	What It Does	Value of <code>x</code> after the Statement
<code>x = x + 4;</code>	Adds 4 to <code>x</code>	10
<code>x = x - 3;</code>	Subtracts 3 from <code>x</code>	3

`x = x * 10;` Multiplies `x` by 10 60

`x = x / 2;` Divides `x` by 2 3

`x = x % 4` Assigns the remainder of `x / 4` to `x`. 2

These types of operations are common in programming. For convenience, Java offers a special set of operators designed specifically for these jobs. [Table 2-13](#) shows the *combined assignment operators*, also known as *compound operators*.

Table 2-13 Combined assignment operators

Operator	Example	Usage	Equivalent To

<code>+=</code>	<code>x += 5;</code>		<code>x = x + 5;</code>

<code>-=</code>	<code>y -= 2;</code>		<code>y = y - 2;</code>

<code>*=</code>	<code>z *= 10;</code>		<code>z = z * 10;</code>

<code>/=</code>	<code>a /= b;</code>		<code>a = a / b;</code>

<code>%=</code>	<code>c %= 3;</code>		<code>c = c % 3;</code>

As you can see, the combined assignment operators do not require the programmer to type the variable name twice. The following statement:

```
balance = balance + deposit;
```

could be rewritten as

```
balance += deposit;
```

Similarly, the statement

```
balance = balance - withdrawal;
```

could be rewritten as

```
balance -= withdrawal;
```



MyProgrammingLab™ www.myprogramminglab.com

1. 2.24 Write statements using combined assignment operators to perform the following:

1. Add 6 to `x`
2. Subtract 4 from `amount`
3. Multiply `y` by 4
4. Divide `total` by 27
5. Store in `x` the remainder of `x` divided by 7

2.7 Conversion between Primitive Data Types

CONCEPT: Before a value can be stored in a variable, the value's data type must be compatible with the variable's data type. Java performs some conversions between data types automatically, but does not automatically perform any conversion that can result in the loss of data. Java also follows a set of rules when evaluating arithmetic expressions containing mixed data types.

Java is a *strongly typed* language. This means that before a value is assigned to a variable, Java checks the data types of the variable and the value being assigned to it to determine whether they are compatible. For example, look at the following statements:

```
int x;  
double y = 2.5;  
x = y;
```

The assignment statement is attempting to store a `double` value (2.5) in an `int` variable. When the Java compiler encounters this line of code, it will respond with an error message. (The JDK displays the message “possible loss of precision.”)

Not all assignment statements that mix data types are rejected by the compiler, however. For instance, look at the following program segment:

```
int x;  
short y = 2;  
x = y;
```

This assignment statement, which stores a `short` in an `int`, will work with no problems. So why does Java permit a `short` to be stored in an `int`, but does not permit a `double` to be stored in an `int`? The obvious reason is that a `double` can store fractional numbers and can hold values much larger than an `int` can hold. If Java were to permit a `double` to be assigned to an `int`, a loss of data would be likely.

Just like officers in the military, the primitive data types are ranked. One data type outranks another if it can hold a larger number. For example, a `float` outranks an `int`, and an `int` outranks a `short`. [Figure 2-6](#) shows the numeric data types in order of their rank. The higher a data type appears in the list, the higher is its rank.




Figure 2-6 Primitive data type ranking

In assignment statements where values of lower-ranked data types are stored in variables of higher-ranked data types, Java automatically converts the lower-ranked value to the higher-ranked type. This is called a *widening conversion*. For example, the following code demonstrates a widening conversion, which takes place when an `int` value is stored in a `double` variable:

```
double x;  
int y = 10;  
x = y;           // Performs a widening conversion
```

A *narrowing conversion* is the conversion of a value to a lower-ranked type. For example, converting a `double` to an `int` would be a narrowing conversion. Because narrowing conversions can potentially cause a loss of data, Java does not automatically perform them.

Cast Operators

The *cast operator* lets you manually convert a value, even if it means that a narrowing conversion will take place. Cast operators are unary operators that appear as a data type name enclosed in a set of parentheses. The operator precedes the value being converted. Here is an example:

```
x = (int)number;
```

The cast operator in this statement is the word `int` inside the parentheses. It returns the value in `number`, converted to an `int`. This converted value is then stored in `x`. If `number` were a floating-point variable, such as a `float` or a `double`, the value that is returned would be *truncated*, which means the fractional part of the number is lost. The original value in the `number` variable is not changed, however.

[Table 2-14](#) shows several statements using cast operators.

Table 2-14 Example uses of cast operators

Statement	Description
<pre>littleNum = (short)bigNum;</pre>	The cast operator returns the value in <code>bigNum</code> , converted to a <code>short</code> . The converted value is assigned to the variable <code>littleNum</code> .
<pre>x = (long)3.7;</pre>	The cast operator is applied to the expression <code>3.7</code> . The operator returns

the value 3, which is assigned to the variable `x`.

```
number =  
(int)72.567;
```

The cast operator is applied to the expression `72.567`. The operator returns 72, which is used to initialize the variable `number`.

```
value =  
(float)x;
```

The cast operator returns the value in `x`, converted to a `float`. The converted value is assigned to the variable `value`.

```
value =  
(byte)number;
```

The cast operator returns the value in `number`, converted to a `byte`. The converted value is assigned to the variable `value`.

Note that when a cast operator is applied to a variable, it does not change the contents of the variable. It only returns the value stored in the variable, converted to the specified data type.

Recall from our earlier discussion that when both operands of a division are integers, the operation will result in integer division. This means that the result of the division will be an integer, with any fractional part of the result thrown away. For example, look at the following code:

```
int pies = 10, people = 4;  
double piesPerPerson;  
piesPerPerson = pies / people;
```

Although 10 divided by 4 is 2.5, this code will store 2 in the `piesPerPerson` variable. Because both `pies` and `people` are `int` variables, the result will be an `int`, and the fractional part will be thrown away. We can modify the code with a cast operator, however, so it gives the correct result as a floating-point value:

```
piesPerPerson = (double)pies / people;
```

The variable `pies` is an `int` and holds the value 10. The expression `(double)pies` returns the value in `pies` converted to a `double`. This means that one of the division operator's operands is a `double`, so the result of the division will be a `double`. The statement could also have been written as follows:

```
piesPerPerson = pies / (double)people;
```

In this statement, the cast operator returns the value of the `people` variable converted to a `double`. In either statement, the result of the division is a `double`.



WARNING!

The cast operator can be applied to an entire expression enclosed in parentheses. For example, look

at the following statement:

```
piesPerPerson = (double) (pies / people);
```

This statement does not convert the value in `pies` or `people` to a `double`, but converts the result of the expression `pies / people`. If this statement were used, an integer division operation would still have been performed. Here's why: The result of the expression `pies / people` is 2 (because integer division takes place). The value 2 converted to a `double` is 2.0. To prevent the integer division from taking place, one of the operands must be converted to a `double`.

Mixed Integer Operations

One of the nuances of the Java language is the way it internally handles arithmetic operations on `int`, `byte`, and `short` variables. When values of the `byte` or `short` data types are used in arithmetic expressions, they are temporarily converted to `int` values. The result of an arithmetic operation using only a mixture of `byte`, `short`, or `int` values will always be an `int`.

For example, assume that `b` and `c` in the following expression are `short` variables:

```
b + c
```

Although both `b` and `c` are `short` variables, the result of the expression `b + c` is an `int`. This means that when the result of such an expression is stored in a variable, the variable must be an `int` or higher data type. For example, look at the following code:

```
short firstNumber = 10,
      secondNumber = 20,
      thirdNumber;
// The following statement causes an error!
thirdNumber = firstNumber + secondNumber;
```

When this code is compiled, the following statement causes an error:

```
thirdNumber = firstNumber + secondNumber;
```

The error results from the fact that `thirdNumber` is a `short`. Although `firstNumber` and `secondNumber` are also `short` variables, the expression `firstNumber + secondNumber` results in an `int` value. The program can be corrected if `thirdNumber` is declared as an `int`, or if a cast operator is used in the assignment statement, as shown here:

```
thirdNumber = (short) (firstNumber + secondNumber);
```

Other Mixed Mathematical Expressions

In situations where a mathematical expression has one or more values of the `double`, `float`, or `long`

data types, Java strives to convert all of the operands in the expression to the same data type. Let's look at the specific rules that govern evaluation of these types of expressions.

1. If one of an operator's operands is a `double`, the value of the other operand will be converted to a `double`. The result of the expression will be a `double`. For example, in the following statement assume that `b` is a `double` and `c` is an `int`:

```
a = b + c;
```

The value in `c` will be converted to a `double` prior to the addition. The result of the addition will be a `double`, so the variable `a` must also be a `double`.

2. If one of an operator's operands is a `float`, the value of the other operand will be converted to a `float`. The result of the expression will be a `float`. For example, in the following statement assume that `x` is a `short` and `y` is a `float`:

```
z = x * y;
```

The value in `x` will be converted to a `float` prior to the multiplication. The result of the multiplication will be a `float`, so the variable `z` must also be either a `double` or a `float`.

3. If one of an operator's operands is a `long`, the value of the other operand will be converted to a `long`. The result of the expression will be a `long`. For example, in the following statement assume that `a` is a `long` and `b` is a `short`:

```
c = a - b;
```

The variable `b` will be converted to a `long` prior to the subtraction. The result of the subtraction will be a `long`, so the variable `c` must also be a `long`, `float`, or `double`.



MyProgrammingLab™ www.myprogramminglab.com

1. 2.25 The following declaration appears in a program:

```
short totalPay, basePay = 500, bonus = 1000;
```

The following statement appears in the same program:

```
totalPay = basePay + bonus;
```

1. Will the statement compile properly or cause an error?
2. If the statement causes an error, why? How can you fix it?

2. 2.26 The variable `a` is a `float` and the variable `b` is a `double`. Write a statement that will assign the value of `b` to `a` without causing an error when the program is compiled.

2.8 Creating Named Constants with `final`

CONCEPT: The `final` key word can be used in a variable declaration to make the variable a named constant. Named constants are initialized with a value, and that value cannot change during the execution of the program.

Assume that the following statement appears in a banking program that calculates data pertaining to loans:

```
amount = balance * 0.069;
```

In such a program, two potential problems arise. First, it is not clear to anyone other than the original programmer what 0.069 is. It appears to be an interest rate, but in some situations there are fees associated with loan payments. How can the purpose of this statement be determined without painstakingly checking the rest of the program?

The second problem occurs if this number is used in other calculations throughout the program and must be changed periodically. Assuming the number is an interest rate, what if the rate changes from 6.9 percent to 8.2 percent? The programmer would have to search through the source code for every occurrence of the number.

Both of these problems can be addressed by using named constants. A *named constant* is a variable whose value is read only and cannot be changed during the program's execution. You can create such a variable in Java by using the `final` key word in the variable declaration. The word `final` is written just before the data type. Here is an example:

```
final double INTEREST_RATE = 0.069;
```

This statement looks just like a regular variable declaration except that the word `final` appears before the data type, and the variable name is written in all uppercase letters. It is not required that the variable name appear in all uppercase letters, but many programmers prefer to write them this way so they are easily distinguishable from regular variable names.

When you declare a variable as `final`, you must either initialize it, or assign it a value, before you can use it. For example, the following code declares and initializes a `final` variable named `PRICE` before it is used in an output operation:

```
final double PRICE = 19.99;  
System.out.println("The price is " + PRICE);
```

The following code declares a `final` variable named `PRICE`, assigns it a value in a separate statement, and then uses it in an output operation:

```
final double PRICE;  
PRICE = 19.99;
```

```
System.out.println("The price is " + PRICE);
```

The following code will cause an error when it is compiled, however, because you cannot use a `final` variable before you have given it a value:

```
final double PRICE;  
System.out.println("The price is " + PRICE); // Error!  
PRICE = 19.99;
```

Once you have given a `final` variable a value, no statement in the program can change its value. The following code will cause a compiler error:

```
final double PRICE;  
PRICE = 19.99;  
System.out.println("The price is " + PRICE);  
PRICE = 12.99; // Error! Cannot change PRICE.
```

An advantage of using named constants is that they make programs more self-documenting. The following statement:

```
amount = balance * 0.069;
```

can be changed to read

```
amount = balance * INTEREST_RATE;
```

A new programmer can read the second statement and know what is happening. It is evident that `balance` is being multiplied by the interest rate. Another advantage to this approach is that widespread changes can easily be made to the program. Let's say the interest rate appears in a dozen different statements throughout the program. When the rate changes, the initialization value in the definition of the named constant is the only value that needs to be modified. If the rate increases to 8.2 percent, the declaration can be changed to the following:

```
final double INTEREST_RATE = 0.082;
```

The program is then ready to be recompiled. Every statement that uses `INTEREST_RATE` will use the new value.

The `Math.PI` Named Constant

The `Math` class, which is part of the Java API, provides a predefined named constant, `Math.PI`. This constant is assigned the value 3.14159265358979323846, which is an approximation of the mathematical value pi. For example, look at the following statement:

```
area = Math.PI * radius * radius;
```

Assuming the `radius` variable holds the radius of a circle, this statement uses the `Math.PI` constant to calculate the area of the circle.

For more information about the Math class, see Appendix F, available on the book's companion Web site at www.pearsonhighered.com/gaddis.

2.9 The String Class

CONCEPT: The `String` class allows you to create objects for holding strings. It also has various methods that allow you to work with strings.

You have already encountered strings and examined programs that display them on the screen, but let's take a moment to make sure you understand what a string is. A string is a sequence of characters. It can be used to represent any type of data that contains text, such as names, addresses, warning messages, and so forth. String literals are enclosed in double-quotation marks, such as the following:

```
"Hello World"  
"Joe Mahoney"
```

Although programs commonly encounter strings and must perform a variety of tasks with them, Java does not have a primitive data type for storing them in memory. Instead, the Java API provides a class for handling strings. You use this class to create objects that are capable of storing strings and performing operations on them. Before discussing this class, let's briefly discuss how classes and objects are related.

Objects Are Created from Classes

[Chapter 1](#) introduced you to objects as software entities that can contain attributes and methods. An object's attributes are data values that are stored in the object. An object's methods are procedures that perform operations on the object's attributes. Before an object can be created, however, it must be designed by a programmer. The programmer determines the attributes and methods that are necessary, and then creates a class that describes the object.

You have already seen classes used as containers for applications. A class can also be used to specify the attributes and methods that a particular type of object may have. Think of a class as a “blueprint” that objects may be created from. So a class is not an object, but a description of an object. When the program is running, it can use the class to create, in memory, as many objects as needed. Each object that is created from a class is called an *instance* of the class.



Don't worry if these concepts seem a little fuzzy to you. As you progress through this book, the concepts of classes and objects will be reinforced again and again.

The String Class

The class that is provided by the Java API for handling strings is named `String`. The first step in using the `String` class is to declare a variable of the `String` class data type. Here is an example of a `String` variable declaration:

```
String name;
```



The `s` in `String` is written in an uppercase letter. By convention, the first character of a class name is always written in an uppercase letter.

This statement declares `name` as a `String` variable. Remember that `String` is a class, not a primitive data type. Let's briefly look at the difference between primitive type variables and class type variables.

Primitive Type Variables and Class Type Variables

A variable of any type can be associated with an item of data. *Primitive type variables* hold the actual data items with which they are associated. For example, assume that `number` is an `int` variable. The following statement stores the value 25 in the variable:

```
number = 25;
```

This is illustrated in [Figure 2-7](#).

The **number** variable holds the actual data with which it is associated.



Figure 2-7 A primitive type variable holds the data with which it is associated

A *class type variable* does not hold the actual data item that it is associated with, but holds the memory address of the data item it is associated with. If `name` is a `String` class variable, then `name` can hold the memory address of a `String` object. This is illustrated in [Figure 2-8](#).



Figure 2-8 A String class variable can hold the address of a String object

When a class type variable holds the address of an object, it is said that the variable references the object. For this reason, class type variables are commonly known as *reference variables*.

Creating a String Object

Any time you write a string literal in your program, Java will create a `String` object in memory to hold it. You can create a `String` object in memory and store its address in a `String` variable with a simple assignment statement. Here is an example:

```
name = "Joe Mahoney";
```

Here, the string literal causes a `String` object to be created in memory with the value “Joe Mahoney” stored in it. Then the assignment operator stores the address of that object in the `name` variable. After this statement executes, it is said that the `name` variable references a `String` object. This is illustrated in [Figure 2-9](#).




Figure 2-9 The name variable holds the address of a string object

You can also use the `=` operator to initialize a `String` variable, as shown here:

```
String name = "Joe Mahoney";
```

This statement declares `name` as a `String` variable, creates a `String` object with the value “Joe Mahoney” stored in it, and assigns the object’s memory address to the `name` variable. [Code Listing 2-20](#) shows `String` variables being declared, initialized, and then used in a `println` statement.

Code Listing 2-20 (StringDemo.java)

```
1 // A simple program demonstrating String objects.
2
3 public class StringDemo
```


Program Output

Herman has 6 characters.



The `String` class's `length` method returns the number of characters in the string, including spaces.

You will study the `String` class methods in detail in [Chapter 9](#), but let's look at a few more examples now. In addition to `length`, [Table 2-15](#) describes the `charAt`, `toLowerCase`, and `toUpperCase` methods.

Table 2-15 A few `String` class methods

Method	Description and Example
<code>charAt(index)</code>	The argument <i>index</i> is an <code>int</code> value and specifies a character position in the string. The first character is at position 0, the second character is at position 1, and so forth. The method returns the character at the specified position. The return value is of the type <code>char</code>. Example: <pre>char letter; String name = "Herman"; letter = name.charAt(3);</pre> After this code executes, the variable <code>letter</code> will hold the character 'm'.
<code>length()</code>	This method returns the number of characters in the string. The return value is of the type <code>int</code>. Example: <pre>int stringSize; String name = "Herman"; stringSize = name.length();</pre>

After this code executes, the `stringSize` variable will hold the value 6.

`toLowerCase()` This method returns a new string that is the lowercase equivalent of the string contained in the calling object.

Example:

```
String bigName = "HERMAN";
String littleName = bigName.toLowerCase();
```

After this code executes, the object referenced by `littleName` will hold the string “herman”.

`toUpperCase()` This method returns a new string that is the uppercase equivalent of the string contained in the calling object.

Example:

```
String littleName = "herman";
String bigName = littleName.toUpperCase();
```

After this code executes, the object referenced by `bigName` will hold the string “HERMAN”.

The program in [Code Listing 2-22](#) demonstrates these methods.

Code Listing 2-22 (StringMethods.java)

```
1 // This program demonstrates a few of the String methods.
2
3 public class StringMethods
4 {
5     public static void main(String[] args)
6     {
7         String message = "Java is Great Fun!";
8         String upper = message.toUpperCase();
9         String lower = message.toLowerCase();
10        char letter = message.charAt(2);
11        int stringSize = message.length();
```

```

12
13     System.out.println(message);
14     System.out.println(upper);
15     System.out.println(lower);
16     System.out.println(letter);
17     System.out.println(stringSize);
18 }
19 }

```

Program Output

```

Java is Great Fun!
JAVA IS GREAT FUN!
java is great fun!
v
18

```

Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 2.27 Write a statement that declares a `String` variable named `city`. The variable should be initialized so it references an object with the string “San Francisco”.
2. 2.28 Assume that `stringLength` is an `int` variable. Write a statement that stores the length of the string referenced by the `city` variable (declared in Checkpoint 2.27) in `stringLength`.
3. 2.29 Assume that `oneChar` is a `char` variable. Write a statement that stores the first character in the string referenced by the `city` variable (declared in Checkpoint 2.27) in `oneChar`.
4. 2.30 Assume that `upperCity` is a `String` reference variable. Write a statement that stores the uppercase equivalent of the string referenced by the `city` variable (declared in Checkpoint 2.27) in `upperCity`.
5. 2.31 Assume that `lowerCity` is a `String` reference variable. Write a statement that stores the lowercase equivalent of the string referenced by the `city` variable (declared in Checkpoint 2.27) in `lowerCity`.

2.10 Scope

CONCEPT: A variable's scope is the part of the program that has access to the variable.

Every variable has a *scope*. The scope of a variable is the part of the program where the variable may be accessed by its name. A variable is visible only to statements inside the variable's scope. The rules that define a variable's scope are complex, and you are only introduced to the concept here. In other chapters of the book we revisit this topic and expand on it.

So far, you have only seen variables declared inside the `main` method. Variables that are declared inside a method are called *local variables*. Later you will learn about variables that are declared outside a method, but for now, let's focus on the use of local variables.

A local variable's scope begins at the variable's declaration and ends at the end of the method in which the variable is declared. The variable cannot be accessed by statements that are outside this region. This means that a local variable cannot be accessed by code that is outside the method, or inside the method but before the variable's declaration. The program in [Code Listing 2-23](#) shows an example.

Code Listing 2-23 (Scope.java)

```
1  // This program can't find its variable.
2
3  public class Scope
4  {
5      public static void main(String[] args)
6      {
7          System.out.println(value); // ERROR!
8          int value = 100;
9      }
10 }
```

The program does not compile because it attempts to send the contents of the variable `value` to `println` before the variable is declared. It is important to remember that the compiler reads your program from top to bottom. If it encounters a statement that uses a variable before the variable is declared, an error will result. To correct the program, the variable declaration must be written before any statement that uses it.



If you compile this program, the compiler will display an error message such as “cannot resolve symbol.” This means that the compiler has encountered a name for which it cannot determine a

meaning.

Another rule that you must remember about local variables is that you cannot have two local variables with the same name in the same scope. For example, look at the following method.

```
public static void main(String[] args)
{
    // Declare a variable named number and
    // display its value.
    int number = 7;
    System.out.println(number);
    // Declare another variable named number and
    // display its value.
    int number = 100;                // ERROR!!!
    System.out.println(number);      // ERROR!!!
}
```

This method declares a variable named `number` and initializes it with the value 7. The variable's scope begins at the declaration statement and extends to the end of the method. Inside the variable's scope a statement appears that declares another variable named `number`. This statement will cause an error because you cannot have two local variables with the same name in the same scope.

2.11 Comments

CONCEPT: Comments are notes of explanation that document lines or sections of a program. Comments are part of the program, but the compiler ignores them. They are intended for people who may be reading the source code.

Comments are short notes that are placed in different parts of a program, explaining how those parts of the program work. Comments are not intended for the compiler. They are intended for programmers to read, to help them understand the code. The compiler skips all of the comments that appear in a program.

As a beginning programmer, you might resist the idea of writing a lot of comments in your programs. After all, it's a lot more fun to write code that actually does something! However, it's crucial that you take the extra time to write comments. They will almost certainly save you time in the future when you have to modify or debug the program. Even large and complex programs can be made easy to read and understand if they are properly commented.

In Java there are three types of comments: single-line comments, multiline comments, and documentation comments. Let's briefly discuss each type.

Single-Line Comments

You have already seen the first way to write comments in a Java program. You simply place two forward slashes (//) where you want the comment to begin. The compiler ignores everything from that point to the end of the line. [Code Listing 2-24](#) shows that comments may be placed liberally throughout a program.

Code Listing 2-24 (Comment1.java)

```
1  // PROGRAM: Comment1.java
2  // Written by Herbert Dorfmann
3  // This program calculates company payroll
4
5  public class Comment1
6  {
7      public static void main(String[] args)
8      {
9          double payRate;           // Holds the hourly pay rate
10         double hours;             // Holds the hours worked
11         int employeeNumber;       // Holds the employee number
12
13         // The Remainder of This Program is Omitted.
14     }
```

```
15 }
```

In addition to telling who wrote the program and describing the purpose of variables, comments can also be used to explain complex procedures in your code.

Multi-Line Comments

The second type of comment in Java is the multi-line comment. *Multi-line comments* start with `/*` (a forward slash followed by an asterisk) and end with `*/` (an asterisk followed by a forward slash). Everything between these markers is ignored. [Code Listing 2-25](#) illustrates how multi-line comments may be used.

Code Listing 2-25 (Comment2.java)

```
1  /*
2      PROGRAM: Comment2.java
3      Written by Herbert Dorfmann
4      This program calculates company payroll
5  */
6
7  public class Comment2
8  {
9      public static void main(String[] args)
10     {
11         double payRate;           // Holds the hourly pay rate
12         double hours;             // Holds the hours worked
13         int employeeNumber;       // Holds the employee number
14
15         // The Remainder of This Program is Omitted.
16     }
17 }
```

Unlike a comment started with `//`, a multi-line comment can span several lines. This makes it more convenient to write large blocks of comments because you do not have to mark every line. Consequently, the multi-line comment is inconvenient for writing single-line comments because you must type both a beginning and an ending comment symbol.

Remember the following advice when using multi-line comments:

- Be careful not to reverse the beginning symbol with the ending symbol.
- Be sure not to forget the ending symbol.

Many programmers use asterisks or other characters to draw borders or boxes around their comments. This helps to visually separate the comments from surrounding code. These are called block comments. [Table 2-16](#) shows four examples of block comments.

Table 2-16 Block comments

```
/*                                //*****
 *   This program demonstrates the //   This program demonstrates the *
 *   way to write comments.        //   way to write comments.          *
 */                                //*****

////////////////////////////////////// //-----
//   This program demonstrates the //   This program demonstrates the
//   way to write comments.        //   way to write comments.
////////////////////////////////////// //-----
```

Documentation Comments

The third type of comment is known as a documentation comment. *Documentation comments* can be read and processed by a program named `javadoc`, which comes with the JDK. The purpose of the `javadoc` program is to read Java source code files and generate attractively formatted HTML files that document the source code. If the source code files contain any documentation comments, the information in the comments becomes part of the HTML documentation. The HTML documentation files may be viewed in a Web browser.

Any comment that starts with `/**` and ends with `*/` is considered a documentation comment. Normally you write a documentation comment just before a class header, giving a brief description of the class. You also write a documentation comment just before each method header, giving a brief description of the method. For example, [Code Listing 2-26](#) shows a program with documentation comments. This program has a documentation comment just before the class header, and just before the `main` method header.

Code Listing 2-26 (Comment3.java)

```
1  /**
2   This class creates a program that calculates company payroll.
3  */
4
5  public class Comment3
6  {
7      /**
8       The main method is the program's starting point.
9      */
10
11     public static void main(String[] args)
12     {
13         double payRate;           // Holds the hourly pay rate
```

```
14      double hours;           // Holds the hours worked
15      int employeeNumber;     // Holds the employee number
16
17      // The Remainder of This Program is Omitted.
18  }
19 }
```

You run the `javadoc` program from the operating system command prompt. Here is the general format of the `javadoc` command:

```
javadoc SourceFile.java
```

SourceFile.java is the name of a Java source code file, including the `.java` extension. The file will be read by `javadoc` and documentation will be produced for it. For example, the following command will produce documentation for the *Comment3.java* source code file, which is shown in [Code Listing 2-26](#):

```
javadoc Comment3.java
```

After this command executes, several documentation files will be created in the same directory as the source code file. One of these files will be named *index.html*. [Figure 2-10](#) shows the *index.html* file being viewed in a Web browser. Notice that the text that was written in the documentation comments appears in the file.




Figure 2-10 Documentation generated by javadoc

(Google Inc.)

TIP:

When you write a documentation comment for a method, the HTML documentation file that is produced by `javadoc` will have two sections for the method: a summary section and a detail section. The first sentence in the method's documentation comment is used as the summary of the method. Note that `javadoc` considers the end of the sentence as a period followed by a whitespace character. For

this reason, when a method description contains more than one sentence, you should always end the first sentence with a period followed by a whitespace character. The method's detail section will contain all of the description that appears in the documentation comment.

If you look at the JDK documentation, which are HTML files that you view in a Web browser, you will see that they are formatted in the same way as the files generated by `javadoc`. A benefit of using `javadoc` to document your source code is that your documentation will have the same professional look and feel as the standard Java documentation.

From this point forward in the book, we will use documentation comments in the example source code. As we progress through various topics, you will see additional uses of documentation comments and the `javadoc` program.



Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 2.32 How do you write a single line comment? How do you write a multi-line comment? How do you write a documentation comment?
2. 2.33 How are documentation comments different from other types of comments?

2.12 Programming Style

CONCEPT: Programming style refers to the way a programmer uses spaces, indentations, blank lines, and punctuation characters to visually arrange a program's source code.

In [Chapter 1](#), you learned that syntax rules govern the way a language may be used. The syntax rules of Java dictate how and where to place key words, semicolons, commas, braces, and other elements of the language. The compiler checks for syntax errors, and if there are none, generates byte code.

When the compiler reads a program it processes it as one long stream of characters. The compiler doesn't care that each statement is on a separate line, or that spaces separate operators from operands. Humans, on the other hand, find it difficult to read programs that aren't written in a visually pleasing manner. Consider [Code Listing 2-27](#) for example.

Code Listing 2-27 (Compact.java)

```
1 public class Compact {public static void main(String [] args){int
2 shares=220; double averagePrice=14.67; System.out.println(
3 "There were "+shares+" shares sold at $" +averagePrice+
4 " per share.");}}
```

Program Output

```
There were 220 shares sold at $14.67 per share.
```

Although the program is syntactically correct (it doesn't violate any rules of Java), it is very difficult to read. The same program is shown in [Code Listing 2-28](#), written in a more understandable style.

Code Listing 2-28 (Readable.java)

```
1 /**
2  * This example is much more readable than Compact.java.
3  */
4
5 public class Readable
6 {
7     public static void main(String[] args)
8     {
9         int shares = 220;
10        double averagePrice = 14.67;
11
12        System.out.println("There were " + shares +
13                            " shares sold at $" +
```

```

14         averagePrice + " per share.");
15     }
16 }

```

Program Output

There were 220 shares sold at \$14.67 per share.

The term *programming style* usually refers to the way source code is visually arranged. It includes techniques for consistently putting spaces and indentations in a program so visual cues are created. These cues quickly tell a programmer important information about a program.

For example, notice in [Code Listing 2-28](#) that inside the class's braces each line is indented, and inside the `main` method's braces each line is indented again. It is a common programming style to indent all the lines inside a set of braces, as shown in [Figure 2-11](#).

```

/**
    This example is much more readable than Compact.java.
 */

public class Readable
{
    public static void main(String[] args)
    {
        int shares = 220;
        double averagePrice = 14.67;

        System.out.println("There were " + shares +
            " shares sold at $" +
            averagePrice + " per share.");
    }
}

```

Figure 2-11 Indentation

[Figure 2-11 Full Alternative Text](#)

Another aspect of programming style is how to handle statements that are too long to fit on one line. Notice that the `println` statement is spread out over three lines. Extra spaces are inserted at the beginning of the statement's second and third lines, which indicate that they are continuations.

When declaring multiple variables of the same type with a single statement, it is a common practice to write each variable name on a separate line with a comment explaining the variable's purpose. Here is an example:

```

int fahrenheit,    // To hold the Fahrenheit temperature
    celsius,       // To hold the Celsius temperature
    kelvin;        // To hold the Kelvin temperature

```

You may have noticed in the example programs that a blank line is inserted between the variable declarations and the statements that follow them. This is intended to separate the declarations visually from the executable statements.

There are many other issues related to programming style. They will be presented throughout the book.

2.13 Reading Keyboard Input

CONCEPT: Objects of the `Scanner` class can be used to read input from the keyboard.

Previously we discussed the `System.out` object, and how it refers to the standard output device. The Java API has another object, `System.in`, which refers to the standard input device. The *standard input device* is normally the keyboard. You can use the `System.in` object to read keystrokes that have been typed at the keyboard. However, using `System.in` is not as simple and straightforward as using `System.out` because the `System.in` object reads input only as `byte` values. This isn't very useful because programs normally require values of other data types as input. To work around this, you can use the `System.in` object in conjunction with an object of the `Scanner` class. The `Scanner` class is designed to read input from a source (such as `System.in`), and it provides methods that you can use to retrieve the input formatted as primitive values or strings.

First, you create a `Scanner` object and connect it to the `System.in` object. Here is an example of a statement that does just that:

```
Scanner keyboard = new Scanner(System.in);
```

Let's dissect the statement into two parts. The first part of the statement,

```
Scanner keyboard
```

declares a variable named `keyboard`. The data type of the variable is `Scanner`. Because `Scanner` is a class, the `keyboard` variable is a class type variable. Recall from our discussion on `String` objects that a class type variable holds the memory address of an object. Therefore, the `keyboard` variable will be used to hold the address of a `Scanner` object. The second part of the statement is as follows:

```
= new Scanner(System.in);
```

The first thing we see in this part of the statement is the assignment operator (`=`). The assignment operator will assign something to the `keyboard` variable. After the assignment operator we see the word `new`, which is a Java key word. The purpose of the `new` key word is to create an object in memory. The type of object that will be created is listed next. In this case, we see `Scanner(System.in)` listed after the `new` key word. This specifies that a `Scanner` object should be created, and it should be connected to the `System.in` object. The memory address of the object is assigned (by the `=` operator) to the variable `keyboard`. After the statement executes, the `keyboard` variable will reference the `Scanner` object that was created in memory.

[Figure 2-12](#) points out the purpose of each part of this statement. [Figure 2-13](#) illustrates how the `keyboard` variable references an object of the `Scanner` class.

This declares a variable named `keyboard`. The variable can reference an object of the `Scanner` class.

This creates a `Scanner` object in memory. The object will read input from `System.in`

```
Scanner keyboard = new Scanner(System.in);
```

The `=` operator assigns the address of the `Scanner` object to the `keyboard` variable.

Figure 2-12 The parts of the statement

[Figure 2-12 Full Alternative Text](#)




Figure 2-13 The keyboard variable references a scanner object

NOTE:

In the preceding code, we chose `keyboard` as the variable name. There is nothing special about the name `keyboard`. We simply chose that name because we will use the variable to read input from the keyboard.

The `Scanner` class has methods for reading strings, bytes, integers, long integers, short integers, floats, and doubles. For example, the following code uses an object of the `Scanner` class to read an `int` value from the keyboard and assign the value to the `number` variable.

```
int number;  
Scanner keyboard = new Scanner(System.in);  
System.out.print("Enter an integer value: ");  
number = keyboard.nextInt();
```

The last statement shown here calls the `Scanner` class's `nextInt` method. The `nextInt` method

formats an input value as an `int`, and then returns that value. Therefore, this statement formats the input that was entered at the keyboard as an `int`, and then returns it. The value is assigned to the `number` variable.

[Table 2-17](#) lists several of the `Scanner` class's methods and describes their use.

Table 2-17 Some of the `Scanner` class methods

Method	Example and Description
Example Usage:	
<code>nextByte</code>	<pre>byte x; Scanner keyboard = new Scanner(System.in); System.out.print("Enter a byte value: "); x = keyboard.nextByte();</pre>
Description: Returns input as a <code>byte</code> .	
Example Usage:	
<code>nextDouble</code>	<pre>double number; Scanner keyboard = new Scanner(System.in); System.out.print("Enter a double value: "); number = keyboard.nextDouble();</pre>
Description: Returns input as a <code>double</code> .	
Example Usage:	
<code>nextFloat</code>	<pre>float number; Scanner keyboard = new Scanner(System.in); System.out.print("Enter a float value: "); number = keyboard.nextFloat();</pre>
Description: Returns input as a <code>float</code> .	

Example Usage:

```
nextInt      int number;
             Scanner keyboard = new Scanner(System.in);
             System.out.print("Enter an integer value: ");
             number = keyboard.nextInt();
```

Description: Returns input as an `int`.

Example Usage:

```
nextLine     String name;
             Scanner keyboard = new Scanner(System.in);
             System.out.print("Enter your name: ");
             name = keyboard.nextLine();
```

Description: Returns input as a `String`.

Example Usage:

```
nextLong     long number;
             Scanner keyboard = new Scanner(System.in);
             System.out.print("Enter a long value: ");
             number = keyboard.nextLong();
```

Description: Returns input as a `long`.

Example Usage:

```
nextShort    short number;
             Scanner keyboard = new Scanner(System.in);
             System.out.print("Enter a short value: ");
             number = keyboard.nextShort();
```

Description: Returns input as a `short`.

Using the `import` Statement

There is one last detail about the `Scanner` class that you must know before you will be ready to use it. The `Scanner` class is not automatically available to your Java programs. Any program that uses the

Scanner class should have the following statement near the beginning of the file, before any class definition:

```
import java.util.Scanner;
```

This statement tells the Java compiler where in the Java library to find the `Scanner` class, and makes it available to your program.

[Code Listing 2-29](#) shows the `Scanner` class being used to read a `String`, an `int`, and a `double`.

Code Listing 2-29 (Payroll.java)

```
1  import java.util.Scanner; // Needed for the Scanner class
2
3  /**
4   * This program demonstrates the Scanner class.
5   */
6
7  public class Payroll
8  {
9      public static void main(String[] args)
10     {
11         String name;           // To hold a name
12         int hours;             // Hours worked
13         double payRate;        // Hourly pay rate
14         double grossPay;       // Gross pay
15
16         // Create a Scanner object to read input.
17         Scanner keyboard = new Scanner(System.in);
18
19         // Get the user's name.
20         System.out.print("What is your name? ");
21         name = keyboard.nextLine();
22
23         // Get the number of hours worked this week.
24         System.out.print("How many hours did you work this week? ");
25         hours = keyboard.nextInt();
26
27         // Get the user's hourly pay rate.
28         System.out.print("What is your hourly pay rate? ");
29         payRate = keyboard.nextDouble();
30
31         // Calculate the gross pay.
32         grossPay = hours * payRate;
33
34         // Display the resulting information.
35         System.out.println("Hello, " + name);
36         System.out.println("Your gross pay is $" + grossPay);
37     }
38 }
```


Program Output with Example Input Shown in Bold

```
What is your name? Joe Mahoney [Enter]
How many hours did you work this week? 40 [Enter]
What is your hourly pay rate? 20 [Enter]
Hello, Joe Mahoney
Your gross pay is $800.0
```

NOTE:

Notice that each `Scanner` class method that we used waits for the user to press the user to press the  key before it returns a value. When the  key is pressed, the cursor automatically moves to the next line for subsequent output operations.

Reading a Character

Sometimes you will want to read a single character from the keyboard. For example, your program might ask the user a yes/no question, and specify that he or she type Y for yes or N for no. The `Scanner` class does not have a method for reading a single character, however. The approach that we will use in this book for reading a character is to use the `Scanner` class's `nextLine` method to read a string from the keyboard, and then use the `String` class's `charAt` method to extract the first character of the string. This will be the character that the user entered at the keyboard. Here is an example:

```
String input;    // To hold a line of input
char answer;     // To hold a single character
// Create a Scanner object for keyboard input.
Scanner keyboard = new Scanner(System.in);
// Ask the user a question.
System.out.print("Are you having fun? (Y=yes, N=no) ");
input = keyboard.nextLine(); // Get a line of input.
answer = input.charAt(0);    // Get the first character.
```

The `input` variable references a `String` object. The last statement in this code calls the `String` class's `charAt` method to retrieve the character at position 0, which is the first character in the string. After this statement executes, the `answer` variable will hold the character that the user typed at the keyboard.

Mixing Calls to `nextLine` with Calls to Other

Scanner Methods

When you call one of the `Scanner` class's methods to read a primitive value, such as `nextInt` or `nextDouble`, and then call the `nextLine` method to read a string, an annoying and hard-to-find problem can occur. For example, look at the program in [Code Listing 2-30](#).

Code Listing 2-30 (InputProblem.java)

```
1  import java.util.Scanner; // Needed for the Scanner class
2
3  /**
4   * This program has a problem reading input.
5   */
6
7  public class InputProblem
8  {
9      public static void main(String[] args)
10     {
11         String name;    // To hold the user's name
12         int age;         // To hold the user's age
13         double income;  // To hold the user's income
14
15         // Create a Scanner object to read input.
16         Scanner keyboard = new Scanner(System.in);
17
18         // Get the user's age.
19         System.out.print("What is your age? ");
20         age = keyboard.nextInt();
21
22         // Get the user's income
23         System.out.print("What is your annual income? ");
24         income = keyboard.nextDouble();
25
26         // Get the user's name.
27         System.out.print("What is your name? ");
28         name = keyboard.nextLine();
29
30         // Display the information back to the user.
31         System.out.println("Hello, " + name + ". Your age is " +
32                             age + " and your income is $" +
33                             income);
34     }
35 }
```

Program Output with Example Input Shown in Bold

What is your age? **24** [Enter]

What is your annual income? 50000.00 [Enter]

What is your name? Hello, . Your age is 24 and your income is \$50000.0

Notice in the example output that the program first allows the user to enter his or her age. The statement in line 20 reads an `int` from the keyboard and stores the value in the `age` variable. Next, the user enters his or her income. The statement in line 24 reads a `double` from the keyboard and stores the value in the `income` variable. Then the user is asked to enter his or her name, but it appears that the statement in line 28 is skipped. The name is never read from the keyboard. This happens because of a slight difference in behavior between the `nextLine` method and the other `Scanner` class methods.

When the user types keystrokes at the keyboard, those keystrokes are stored in an area of memory that is sometimes called the *keyboard buffer*. Pressing the  key causes a newline character to be stored in the keyboard buffer. In the example running of the program in [Code Listing 2-30](#), the user was asked to enter his or her age, and the statement in line 20 called the `nextInt` method to read an integer from the keyboard buffer. Notice that the user typed 24 and then pressed the  key. The `nextInt` method read the value 24 from the keyboard buffer, and then stopped when it encountered the newline character. So the value 24 was read from the keyboard buffer, but the newline character was not read. The newline character remained in the keyboard buffer.

Next, the user was asked to enter his or her annual income. The user typed 50000.00 and then pressed the  key. When the `nextDouble` method in line 24 executed, it first encountered the newline character that was left behind by the `nextInt` method. This does not cause a problem because the `nextDouble` method is designed to skip any leading newline characters it encounters. It skips over the initial newline, reads the value 50000.00 from the keyboard buffer, and stops reading when it encounters the next newline character. This newline character is then left in the keyboard buffer.

Next, the user is asked to enter his or her name. In line 28 the `nextLine` method is called. The `nextLine` method, however, is not designed to skip over an initial newline character. If a newline character is the first character that the `nextLine` method encounters, then nothing will be read. Because the `nextDouble` method, back in line 24, left a newline character in the keyboard buffer, the `nextLine` method will not read any input. Instead, it will immediately terminate and the user will not be given a chance to enter his or her name.

Although the details of this problem might seem confusing, the solution is easy. The program in [Code Listing 2-31](#) is a modification of [Code Listing 2-30](#), with the input problem fixed.

Code Listing 2-31

(CorrectedInputProblem.java)

```
1  import java.util.Scanner; // Needed for the Scanner class
2
3  /**
```

```

4      This program correctly reads numeric and string input.
5      */
6
7      public class CorrectedInputProblem
8      {
9          public static void main(String[] args)
10         {
11             String name;        // To hold the user's name
12             int age;            // To hold the user's age
13             double income;      // To hold the user's income
14
15             // Create a Scanner object to read input.
16             Scanner keyboard = new Scanner(System.in);
17
18             // Get the user's age.
19             System.out.print("What is your age? ");
20             age = keyboard.nextInt();
21
22             // Get the user's income
23             System.out.print("What is your annual income? ");
24             income = keyboard.nextDouble();
25
26             // Consume the remaining newline.
27             keyboard.nextLine();
28
29             // Get the user's name.
30             System.out.print("What is your name? ");
31             name = keyboard.nextLine();
32
33             // Display the information back to the user.
34             System.out.println("Hello, " + name + ". Your age is " +
35                               age + " and your income is $" +
36                               income);
37         }
38     }

```

Program Output with Example Input Shown in Bold

```

What is your age? 24 [Enter]
What is your annual income? 50000.00 [Enter]
What is your name? Mary Simpson [Enter]
Hello, Mary Simpson. Your age is 24 and your income is $50000.0

```

Notice that after the user's income is read by the `nextDouble` method in line 24, the `nextLine` method is called in line 27. The purpose of this call is to consume, or remove, the newline character that remains in the keyboard buffer. Then, in line 31, the `nextLine` method is called again. This time it correctly reads the user's name.



NOTE:

Notice that in line 27, where we consume the remaining newline character, we do not assign the method's return value to any variable. This is because we are simply calling the method to remove the newline character, and we do not need to keep the method's return value.

2.14 Dialog Boxes

CONCEPT: The `JOptionPane` class allows you to quickly display a dialog box, which is a small graphical window displaying a message or requesting input.

A *dialog box* is a small graphical window that displays a message to the user or requests input. You can quickly display dialog boxes with the `JOptionPane` class. In this section we will discuss the following types of dialog boxes and how you can display them using `JOptionPane`:

- **Message Dialog** A dialog box that displays a message; an OK button is also displayed
- **Input Dialog** A dialog box that prompts the user for input and provides a text field where input is typed; an OK button and a Cancel button are also displayed

[Figure 2-14](#) shows an example of each type of dialog box.




Figure 2-14 A message dialog and an input dialog

(Oracle Corporate Counsel)

[Figure 2-14 Full Alternative Text](#)

The `JOptionPane` class is not automatically available to your Java programs. Any program that uses the `JOptionPane` class must have the following statement near the beginning of the file:

```
import javax.swing.JOptionPane;
```

This statement tells the compiler where to find the `JOptionPane` class and makes it available to your program.

Displaying Message Dialogs

The `showMessageDialog` method is used to display a message dialog. Here is a statement that calls the method:

```
JOptionPane.showMessageDialog(null, "Hello World");
```

The first argument is only important in programs that display other graphical windows. In our examples, we will always pass the key word `null` as the first argument. This causes the dialog box to be displayed in the center of the screen. The second argument is the message that we wish to display in the dialog box. This code will cause the dialog box in [Figure 2-15](#) to appear. When the user clicks the OK button, the dialog box will close.




Figure 2-15 Message dialog

(Oracle Corporate Counsel)

Displaying Input Dialogs

An input dialog is a quick and simple way to ask the user to enter data. You use the `JOptionPane` class's `showInputDialog` method to display an input dialog. The following code calls the method:

```
String name;  
name = JOptionPane.showInputDialog("Enter your name.");
```

The argument passed to the method is a message to display in the dialog box. This statement will cause the dialog box shown in [Figure 2-16](#) to be displayed in the center of the screen. If the user clicks the OK button, `name` will reference the string value entered by the user into the text field. If the user clicks the Cancel button, `name` will reference the special value `null`.




Figure 2-16 Input dialog

(Oracle Corporate Counsel)

An Example Program

The program in [Code Listing 2-32](#) demonstrates how to use both types of dialog boxes. This program uses input dialogs to ask the user to enter his or her first, middle, and last names, and then displays a greeting with a message dialog. When this program executes, the dialog boxes shown in [Figure 2-17](#) will be displayed, one at a time.




Figure 2-17 Dialog boxes displayed by the NamesDialog program

[Figure 2-17 Full Alternative Text](#)

Code Listing 2-32 (NamesDialog.java)

```
1  import javax.swing.JOptionPane;
2
3  /**
4   * This program demonstrates using dialogs
5   * with JOptionPane.
6   */
7
8  public class NamesDialog
9  {
10     public static void main(String[] args)
11     {
12         String firstName, lastName;
13
14         // Get the user's first name.
15         firstName =
16             JOptionPane.showInputDialog("What is your first name?");
17
18         // Get the user's last name.
19         lastName =
20             JOptionPane.showInputDialog("What is your last name?");
21
22         // Display a greeting
23         JOptionPane.showMessageDialog(null, "Hello " + firstName +
24                                     " " + lastName);
25         System.exit(0);
26     }
27 }
```

Notice the last statement in the `main` method:

```
System.exit(0);
```

This statement causes the program to end, and is necessary when you use the `JOptionPane` class in a console program. When a console program uses `JOptionPane` to display a dialog box, the program does not automatically stop executing when the end of the `main` method is reached. This is because the `JOptionPane` class causes an additional task to run in the JVM. If the `System.exit` method is not called, this task, also known as a *thread*, will continue to execute, even after the end of the `main` method has been reached.

The `System.exit` method requires an integer argument. This argument is an exit code that is passed back to the operating system. Although this code is usually ignored, it can be used outside the program to indicate whether the program ended successfully or as the result of a failure. The value 0 traditionally indicates that the program ended successfully.



WARNING!

Use `System.exit` only when absolutely necessary. In a large program, `System.exit` can cause problems because it unconditionally shuts down the JVM, bypassing the program's normal, logical flow. In this book we use `System.exit` only because it is necessary to end a console program that uses `JOptionPane`.

Converting String Input to Numbers

Unlike the `Scanner` class, the `JOptionPane` class does not have different methods for reading values of different data types as input. The `showInputDialog` method always returns the user's input as a `String`, even if the user enters numeric data. For example, if the user enters the number 72 into an input dialog, the `showInputDialog` method will return the string "72". This can be a problem if you wish to use the user's input in a math operation because, as you know, you cannot perform math on strings. In such a case, you must convert the input to a numeric value. To convert a string value to a numeric value, you use one of the methods listed in [Table 2-18](#).

Table 2-18 Methods for converting strings to numbers

Method	Use This Method To ...	Example Code
<code>Byte.parseByte</code>	Convert a string to a byte.	<pre>byte num; num = Byte.parseByte(str);</pre>
<code>Double.parseDouble</code>	Convert a string to a double.	<pre>double num; num = Double.parseDouble(str);</pre>
<code>Float.parseFloat</code>	Convert a string to a float.	<pre>float num; num = Float.parseFloat(str);</pre>
<code>Integer.parseInt</code>	Convert a string to an int.	<pre>int num; num = Integer.parseInt(str);</pre>

`Long.parseLong`

Convert a string to a long.

```
long num;
```

```
num = Long.parseLong(str);
```

`Short.parseShort`

Convert a string to a short.

```
short num;
```

```
num = Short.parseShort(str);
```



The methods in [Table 2-18](#) are part of Java's wrapper classes, which you will learn more about in [Chapter 9](#).

Here is an example of how you would use the `Integer.parseInt` method to convert the value returned from the `JOptionPane.showInputDialog` method to an `int`:

```
int number;
String str;
str = JOptionPane.showInputDialog("Enter a number.");
number = Integer.parseInt(str);
```

After this code executes, the `number` variable will hold the value entered by the user, converted to an `int`. Here is an example of how you would use the `Double.parseDouble` method to convert the user's input to a `double`:


```
double price;
String str;
str = JOptionPane.showInputDialog("Enter the retail price.");
price = Double.parseDouble(str);
```

After this code executes, the `price` variable will hold the value entered by the user, converted to a `double`. [Code Listing 2-33](#) shows a complete program. This is a modification of the *Payroll.java* program in [Code Listing 2-29](#). When this program executes, the dialog boxes shown in [Figure 2-18](#) will be displayed, one at a time.

The first dialog box appears as shown here. The user enters his or her name and then clicks OK.



The second dialog box appears, as shown here. The user enters the number of hours worked and then clicks OK.



The third dialog box appears, as shown here. The user enters his or her hourly pay rate and then clicks OK.



The fourth dialog box appears, as shown here.




Figure 2-18 Dialog boxes displayed by PayrollDialog.java

(Oracle Corporate Counsel)

[Figure 2-18 Full Alternative Text](#)

Code Listing 2-33 (PayrollDialog.java)

```
1  import javax.swing.JOptionPane;
2
3  /**
4   * This program demonstrates using dialogs with
5   * JOptionPane.
6   */
7
8  public class PayrollDialog
```

```

9      {
10         public static void main(String[] args)
11         {
12             String inputString;           // For reading input
13             String name;                  // The user's name
14             int hours;                    // The number of hours worked
15             double payRate;              // The user's hourly pay rate
16             double grossPay;             // The user's gross pay
17
18             // Get the user's name.
19             name = JOptionPane.showInputDialog("What is " +
20                                                "your name? ");
21
22             // Get the hours worked.
23             inputString =
24                 JOptionPane.showInputDialog("How many hours " +
25                                             "did you work this week? ");
26
27             // Convert the input to an int.
28             hours = Integer.parseInt(inputString);
29
30             // Get the hourly pay rate.
31             inputString =
32                 JOptionPane.showInputDialog("What is your " +
33                                             "hourly pay rate? ");
34
35             // Convert the input to a double.
36             payRate = Double.parseDouble(inputString);
37
38             // Calculate the gross pay.
39             grossPay = hours * payRate;
40
41             // Display the results.
42             JOptionPane.showMessageDialog(null, "Hello " +
43                                                name + ". Your gross pay is $" +
44                                                grossPay);
45
46             // End the program.
47             System.exit(0);
48         }
49     }

```



MyProgrammingLab™ www.myprogramminglab.com

1. 2.34 What is the purpose of the following types of dialog boxes?
 - Message dialog
 - Input dialog

2. 2.35 Write code that will display each of the dialog boxes shown in [Figure 2-19](#).




Figure 2-19 Dialog boxes

(Oracle Corporate Counsel)

3. 2.36 Write code that displays an input dialog asking the user to enter his or her age. Convert the input value to an `int` and store it in an `int` variable named `age`.
4. 2.37 What `import` statement do you write in a program that uses the `JOptionPane` class?

2.15 Common Errors to Avoid

- Mismatched braces, quotation marks, or parentheses. In this chapter you saw that the statements making up a class definition are enclosed in a set of braces. Also, you saw that the statements in a method are also enclosed in a set of braces. For every opening brace, there must be a closing brace in the proper location. The same is true of double-quotation marks that enclose string literals and single-quotation marks that enclose character literals. Also, in a statement that uses parentheses, such as a mathematical expression, you must have a closing parenthesis for every opening parenthesis.
- Misspelling key words. Java will not recognize a key word that has been misspelled.
- Using capital letters in key words. Remember that Java is a case-sensitive language, and all key words are written in lowercase. Using an uppercase letter in a key word is the same as misspelling the key word.
- Using a key word as a variable name. The key words are reserved for special uses; they cannot be used for any other purpose.
- Using inconsistent spelling of variable names. Each time you use a variable name, it must be spelled exactly as it appears in its declaration statement.
- Using inconsistent case of letters in variable names. Because Java is a case-sensitive language, it distinguishes between uppercase and lowercase letters. Java will not recognize a variable name that is not written exactly as it appears in its declaration statement.
- Inserting a space in a variable name. Spaces are not allowed in variable names. Instead of using a two-word name such as `gross pay`, use one word, such as `grossPay`.
- Forgetting the semicolon at the end of a statement. A semicolon appears at the end of each complete statement in Java.
- Assigning a `double` literal to a `float` variable. Java is a strongly typed language, which means that it only allows you to store values of compatible data types in variables. All floating-point literals are treated as `doubles`, and a `double` value is not compatible with a `float` variable. A floating-point literal must end with the letter `f` or `F` in order to be stored in a `float` variable.
- Using commas or other currency symbols in numeric literals. Numeric literals cannot contain commas or currency symbols, such as the dollar sign.
- Unintentionally performing integer division. When both operands of a division statement are integers, the statement will result in an integer. If there is a remainder, it will be discarded.
- Forgetting to group parts of a mathematical expression. If you use more than one operator in a

mathematical expression, the expression will be evaluated according to the order of operations. If you wish to change the order in which the operators are used, you must use parentheses to group part of the expression.

- Inserting a space in a combined assignment operator. A space cannot appear between the two operators that make a combined assignment operator.
- Using a variable to receive the result of a calculation when the variable's data type is incompatible with the data type of the result. A variable that receives the result of a calculation must be of a data type that is compatible with the data type of the result.
- Incorrectly terminating a multi-line comment or a documentation comment. Multi-line comments and documentation comments are terminated by the `*/` characters. Forgetting to place these characters at a comment's desired ending point, or accidentally switching the `*` and the `/`, will cause the comment not to have an ending point.
- Forgetting to use the correct `import` statement in a program that uses the `Scanner` class or the `JOptionPane` class. In order for the `Scanner` class to be available to your program, you must have the `import java.util.Scanner;` statement near the top of your program file. In order for the `JOptionPane` class to be available to your program, you must have the `import javax.swing.JOptionPane;` statement near the top of the program file.
- When using an input dialog to read numeric input, not converting the `showInputDialog` method's return value to a number. The `showInputDialog` method always returns the user's input as a string. If the user enters a numeric value, it must be converted to a number before it can be used in a math statement.

Review Questions and Exercises

Multiple Choice and True/False

1. Every complete statement ends with a _____.

1. period
2. parenthesis
3. semicolon
4. ending brace

2. The following data

```
72
'A'
"Hello World"
2.8712
```

are all examples of _____.

1. variables
2. literals
3. strings
4. none of these

3. A group of statements, such as the contents of a class or a method, are enclosed in _____.

1. braces { }
2. parentheses ()
3. brackets []
4. any of these will do

4. Which of the following are *not* valid assignment statements? (Indicate all that apply.)

1. `total = 9;`

2. `72 = amount;`
 3. `profit = 129`
 4. `letter = 'W';`
5. Which of the following are not valid `println` statements? (Indicate all that apply.)
1. `System.out.println + "Hello World";`
 2. `System.out.println("Have a nice day");`
 3. `out.System.println(value);`
 4. `println.out(Programming is great fun);`
6. The negation operator is .
1. unary
 2. binary
 3. ternary
 4. none of these
7. This key word is used to declare a named constant.
1. `constant`
 2. `namedConstant`
 3. `final`
 4. `concrete`
8. These characters mark the beginning of a multi-line comment.
1. `//`
 2. `/*`
 3. `*/`
 4. `/**`
9. These characters mark the beginning of a single-line comment.

1. `//`
2. `/*`
3. `*/`
4. `/**`

10. These characters mark the beginning of a documentation comment.

1. `//`
2. `/*`
3. `*/`
4. `/**`

11. Which `Scanner` class method would you use to read a string as input?

1. `nextString`
2. `nextLine`
3. `readString`
4. `getLine`

12. Which `Scanner` class method would you use to read a `double` as input?

1. `nextDouble`
2. `getDouble`
3. `readDouble`
4. None of these; you cannot read a `double` with the `Scanner` class

13. You can use this class to display dialog boxes.

1. `JOptionPane`
2. `BufferedReader`
3. `InputStreamReader`
4. `DialogBox`

14. When Java converts a lower-ranked value to a higher-ranked type, it is called a(n) .
1. 4-bit conversion
 2. escalating conversion
 3. widening conversion
 4. narrowing conversion
15. This type of operator lets you manually convert a value, even if it means that a narrowing conversion will take place.
1. cast
 2. binary
 3. uploading
 4. dot
16. **True or False:** A left brace in a Java program is always followed by a right brace later in the program.
17. **True or False:** A variable must be declared before it can be used.
18. **True or False:** Variable names may begin with a number.
19. **True or False:** You cannot change the value of a variable whose declaration uses the `final` key word.
20. **True or False:** Comments that begin with `//` can be processed by `javadoc`.
21. **True or False:** If one of an operator's operands is a `double`, and the other operand is an `int`, Java will automatically convert the value of the `double` to an `int`.

Predict the Output

What will the following code segments print on the screen?

```
1. int freeze = 32, boil = 212;

   ◦ freeze = 0;

   ◦ boil = 100;
```

- `System.out.println(freeze + "\n"+boil + "\n");`

2. `int x = 0, y = 2;`

- `x = y * 4;`

- `System.out.println(x + "\n" + y + "\n");`

3. `System.out.print("I am the incredible");`

- `System.out.print("computing\nmachine");`

- `System.out.print("\nand I will\namaze\n");`

- `System.out.println("you.");`

4. `System.out.print("Be careful\n");`

- `System.out.print("This might/n be a trick ");`

- `System.out.println("question.");`

5. `int a, x = 23;`

- `a = x % 2;`

- `System.out.println(x + "\n" + a);`

Find the Error

There are a number of syntax errors in the following program. Locate as many as you can.

```
*/ What's wrong with this program? /*
public MyProgram
{
    public static void main(String[] args);
    }
    int a, b, c  \\ Three integers
    a = 3
    b = 4
    c = a + b
    System.out.println('The value of c is' + C);
}
```

Algorithm Workbench

1. Show how the `double` variables `temp`, `weight`, and `age` can be declared in one statement.
2. Show how the `int` variables `months`, `days`, and `years` may be declared in one statement, with `months` initialized to 2 and `years` initialized to 3.
3. Write assignment statements that perform the following operations with the variables `a`, `b`, and `c`.
 1. Adds 2 to `a` and stores the result in `b`
 2. Multiplies `b` times 4 and stores the result in `a`
 3. Divides `a` by 3.14 and stores the result in `b`
 4. Subtracts 8 from `b` and stores the result in `a`
 5. Stores the character 'K' in `c`
 6. Stores the Unicode code for 'B' in `c`
4. Assume the variables `result`, `w`, `x`, `y`, and `z` are all integers, and that `w = 5`, `x = 4`, `y = 8`, and `z = 2`. What value will be stored in `result` in each of the following statements?
 1. `result = x + y;`
 2. `result = z * 2;`
 3. `result = y / x;`
 4. `result = y - z;`
 5. `result = w % 2;`
5. How would each of the following numbers be represented in E notation?
 1. 3.287×10^6
 2. -9.7865×10^{12}
 3. 7.65491×10^{-3}
6. Modify the following program so it prints two blank lines between each line of text.

```
public class
{
    public static void main(String[] args)
    {
        System.out.print("Hearing in the distance");
        System.out.print("Two mandolins like creatures in the");
        System.out.print("dark");
    }
}
```

```

        System.out.print("Creating the agony of ecstasy.");
        System.out.println("                - George Barker");
    }
}

```

7. What will the following code output?

```

int apples = 0, bananas = 2, pears = 10;
apples += 10;
bananas *= 10;
pears /= 10;
System.out.println(apples + " " +
                    bananas + " " +
                    pears);

```

8. What will the following code output?

```

double d = 12.9;
int i = (int)d;
System.out.println(i);

```

9. What will the following code output?

```

String message = "Have a great day!";
System.out.println(message.charAt(5));

```

10. What will the following code output?

```

String message = "Have a great day!";
System.out.println(message.toUpperCase());
System.out.println(message);

```

11. Convert the following pseudocode to Java code. Be sure to declare the appropriate variables.

- *Store 20 in the speed variable.*
- *Store 10 in the time variable.*
- *Multiply speed by time and store the result in the distance variable.*
- *Display the contents of the distance variable.*

12. Convert the following pseudocode to Java code. Be sure to declare the appropriate variables.

- *Store 172.5 in the force variable.*
- *Store 27.5 in the area variable.*
- *Divide area by force and store the result in the pressure variable.*
- *Display the contents of the pressure variable.*

13. Write the code to set up all the necessary objects for reading keyboard input. Then write code that asks the user to enter his or her desired annual income. Store the input in a `double` variable.
14. Write the code to display a dialog box that asks the user to enter his or her desired annual income. Store the input in a `double` variable.
15. A program has a `float` variable named `total` and a `double` variable named `number`. Write a statement that assigns `number` to `total` without causing an error when compiled.

Short Answer

1. Is the following comment a single-line style comment or a multi-line style comment?

```
/* This program was written by M. A. Codewriter */
```
2. Is the following comment a single-line style comment or a multi-line style comment?

```
// This program was written by M. A. Codewriter
```
3. Describe what the phrase “self-documenting program” means.
4. What is meant by “case-sensitive”? Why is it important for a programmer to know that Java is a case-sensitive language?
5. Briefly explain how the `print` and `println` methods are related to the `System` class and the `out` object.
6. What does a variable declaration tell the Java compiler about a variable?
7. Why are variable names like `x` not recommended?
8. What things must be considered when deciding on a data type to use for a variable?
9. Briefly describe the difference between variable assignment and variable initialization.
10. What is the difference between comments that start with the `//` characters and comments that start with the `/*` characters?
11. Briefly describe what programming style means. Why should your programming style be consistent?
12. Assume that a program uses the named constant `PI` to represent the value 3.14. The program uses the named constant in several statements. What is the advantage of using the named constant instead of the actual value 3.14 in each statement?
13. Assume the file *SalesAverage.java* is a Java source file that contains documentation comments.

Assuming you are in the same folder or directory as the source code file, what command would you enter at the operating system command prompt to generate the HTML documentation files?

14. An expression adds a `byte` variable and a `short` variable. Of what data type will the result be?

Programming Challenges

MyProgrammingLab™ Visit www.myprogramminglab.com to complete many of these Programming Challenges online and get instant feedback.

1. Name, Age, and Annual Income

Write a program that declares the following:

- a `String` variable named `name`
- an `int` variable named `age`
- a `double` variable named `annualPay`

Store your age, name, and desired annual income as literals in these variables. The program should display these values on the screen in a manner similar to the following:

```
My name is Joe Mahoney, my age is 26 and  
I hope to earn $100000.0 per year.
```

2. Name and Initials

Write a program that has the following `String` variables: `firstName`, `middleName`, and `lastName`. Initialize these with your first, middle, and last names. The program should also have the following `char` variables: `firstInitial`, `middleInitial`, and `lastInitial`. Store your first, middle, and last initials in these variables. The program should display the contents of these variables on the screen.

3. Personal Information

Write a program that displays the following information, each on a separate line:

- Your name
- Your address, with city, state, and ZIP
- Your telephone number
- Your college major

Although these items should be displayed on separate output lines, use only a single `println` statement in your program.

4. Star Pattern

Write a program that displays the following pattern:

```
  *
 * * *
* * * * *
* * * * * *
 * * * * *
  * * *
    *
```

5. Sales Prediction

The East Coast sales division of a company generates 62 percent of total sales. Based on that percentage, write a program that will predict how much the East Coast division will generate if the company has \$4.6 million in sales this year. *Hint: Use the value 0.62 to represent 62 percent.*

6. Land Calculation

One acre of land is equivalent to 43,560 square feet. Write a program that calculates the number of acres in a tract of land with 389,767 square feet. *Hint: Divide the size of the tract of land by the size of an acre to get the number of acres.*

7. Sales Tax

Write a program that will ask the user to enter the amount of a purchase. The program should then compute the state and county sales tax. Assume the state sales tax is 4 percent and the county sales tax is 2 percent. The program should display the amount of the purchase, the state sales tax, the county sales tax, the total sales tax, and the total of the sale (which is the sum of the amount of purchase plus the total sales tax). *Hint: Use the value 0.02 to represent 2 percent, and 0.04 to represent 4 percent.*

8. Cookie Calories

A bag of cookies holds 40 cookies. The calorie information on the bag claims that there are 10 servings in the bag and that a serving equals 300 calories. Write a program that lets the user enter the

number of cookies he or she actually ate and then reports the number of total calories consumed.



The Miles-per-Gallon Problem

9. Miles-per-Gallon

A car's miles-per-gallon (MPG) can be calculated with the following formula:

$$\text{MPG} = \text{Miles driven} / \text{Gallons of gas used}$$

Write a program that asks the user for the number of miles driven and the gallons of gas used. It should calculate the car's miles-per-gallon and display the result on the screen.

10. Test Average

Write a program that asks the user to enter three test scores. The program should display each test score, as well as the average of the scores.

11. Circuit Board Profit

An electronics company sells circuit boards at a 40 percent profit. If you know the retail price of a circuit board, you can calculate its profit with the following formula:

$$\text{Profit} = \text{Retail price} \times 0.4$$

Write a program that asks the user for the retail price of a circuit board, calculates the amount of profit earned for that product, and displays the results on the screen.

12. String Manipulator

Write a program that asks the user to enter the name of his or her favorite city. Use a `String` variable to store the input. The program should display the following:

- The number of characters in the city name
- The name of the city in all uppercase letters

- The name of the city in all lowercase letters
- The first character in the name of the city

13. Restaurant Bill

Write a program that computes the tax and tip on a restaurant bill. The program should ask the user to enter the charge for the meal. The tax should be 6.75 percent of the meal charge. The tip should be 20 percent of the total after adding the tax. Display the meal charge, tax amount, tip amount, and total bill on the screen.

14. Male and Female Percentages

Write a program that asks the user for the number of males and the number of females registered in a class. The program should display the percentage of males and females in the class.

Hint: Suppose there are 8 males and 12 females in a class. There are 20 students in the class. The percentage of males can be calculated as $8 \div 20 = 0.4$, or 40%. The percentage of females can be calculated as $12 \div 20 = 0.6$, or 60%.

15. Stock Commission

Kathryn bought 600 shares of stock at a price of \$21.77 per share. She must pay her stock-broker a 2 percent commission for the transaction. Write a program that calculates and displays the following:

- The amount paid for the stock alone (without the commission)
- The amount of the commission
- The total amount paid (for the stock plus the commission)

16. Energy Drink Consumption

A soft drink company recently surveyed 12,467 of its customers and found that approximately 14 percent of those surveyed purchase one or more energy drinks per week. Of those customers who purchase energy drinks, approximately 64 percent of them prefer citrus-flavored energy drinks. Write a program that displays the following:

- The approximate number of customers in the survey who purchase one or more energy drinks per week

- The approximate number of customers in the survey who prefer citrus-flavored energy drinks

17. Ingredient Adjuster

A cookie recipe calls for the following ingredients:

- 1.5 cups of sugar
- 1 cup of butter
- 2.75 cups of flour

The recipe produces 48 cookies with these amounts of the ingredients. Write a program that asks the user how many cookies he or she wants to make, and then displays the number of cups of each ingredient needed for the specified number of cookies.

18. Word Game

Write a program that plays a word game with the user. The program should ask the user to enter the following:

- His or her name
- His or her age
- The name of a city
- The name of a college
- A profession
- A type of animal
- A pet's name

After the user has entered these items, the program should display the following story, inserting the user's input into the appropriate locations:

There once was a person named **NAME** who lived in **CITY**. At the age of **AGE**, **NAME** went to college at **COLLEGE**. **NAME** graduated and went to work as a **PROFESSION**. Then, **NAME** adopted a(n) **ANIMAL** named **PETNAME**. They both lived happily ever after!

19. Stock Transaction Program

Last month Joe purchased some stock in Acme Software, Inc. Here are the details of the purchase:

- The number of shares that Joe purchased was 1,000.
- When Joe purchased the stock, he paid \$32.87 per share.
- Joe paid his stockbroker a commission that amounted to 2% of the amount he paid for the stock.

Two weeks later Joe sold the stock. Here are the details of the sale:

- The number of shares that Joe sold was 1,000.
- He sold the stock for \$33.92 per share.
- He paid his stockbroker another commission that amounted to 2% of the amount he received for the stock.

Write a program that displays the following information:

- The amount of money Joe paid for the stock.
- The amount of commission Joe paid his broker when he bought the stock.
- The amount that Joe sold the stock for.
- The amount of commission Joe paid his broker when he sold the stock.
- Display the amount of profit that Joe made after selling his stock and paying the two commissions to his broker. (If the amount of profit that your program displays is a negative number, then Joe lost money on the transaction.)

20. Planting Grapevines

A vineyard owner is planting several new rows of grapevines and needs to know how many grapevines to plant in each row. She has determined that after measuring the length of a future row, she can use the following formula to calculate the number of vines that will fit in the row, along with the trellis end-post assemblies that will need to be constructed at each end of the row:

$$V = R - 2ES$$

The terms in the formula are:

- V is the number of grapevines that will fit in the row.
- R is the length of the row, in feet.
- E is the amount of space used by an end-post assembly.
- S is the space between vines, in feet.

Write a program that makes the calculation for the vineyard owner. The program should ask the user to input the following:

- The length of the row, in feet
- The amount of space used by an end-post assembly, in feet
- The amount of space between the vines, in feet

Once the input data has been entered, the program should calculate and display the number of grapevines that will fit in the row.

21. Compound Interest

When a bank account pays compound interest, it pays interest not only on the principal amount that was deposited into the account, but also on the interest that has accumulated over time. Suppose you want to deposit some money into a savings account, and let the account earn compound interest for a certain number of years. The formula for calculating the balance of the account after a specified number of years is:

$$A = P (1 + r/n)^{nt}$$

The terms in the formula are:

- A is the amount of money in the account after the specified number of years.
- P is the principal amount that was originally deposited into the account.
- r is the annual interest rate.
- n is the number of times per year that the interest is compounded.
- t is the specified number of years.

Write a program that makes the calculation for you. The program should ask the user to input the following:

- The amount of principal originally deposited into the account

- The annual interest rate paid by the account
- The number of times per year that the interest is compounded (For example, if interest is compounded monthly, enter 12. If interest is compounded quarterly, enter 4.)
- The number of years the account will be left to earn interest

Once the input data has been entered, the program should calculate and display the amount of money that will be in the account after the specified number of years.

 NOTE:

The user should enter the interest rate as a percentage. For example, 2 percent would be entered as 2, not as .02. The program will then have to divide the input by 100 to move the decimal point to the correct position.

Chapter 3 Decision Structures

TOPICS

1. [3.1 The `if` Statement](#)
2. [3.2 The `if-else` Statement](#)
3. [3.3 Nested `if` Statements](#)
4. [3.4 The `if-else-if` Statement](#)
5. [3.5 Logical Operators](#)
6. [3.6 Comparing `String` Objects](#)
7. [3.7 More about Variable Declaration and Scope](#)
8. [3.8 The Conditional Operator \(Optional\)](#)
9. [3.9 The `switch` Statement](#)
10. [3.10 Displaying Formatted Output with `System.out.printf` and `String.format`](#)
11. [3.11 Common Errors to Avoid](#)

3.1 The `if` Statement

CONCEPT: The `if` statement is used to create a decision structure, which allows a program to have more than one path of execution. The `if` statement causes one or more statements to execute only when a `boolean` expression is `true`.



The `if` Statement

In all the programs you have written so far, the statements are executed one after the other, in the order they appear. You might think of sequentially executed statements as the steps you take as you walk down a road. To complete the journey, you must start at the beginning and take each step, one after the other, until you reach your destination. This is illustrated in [Figure 3-1](#).




Figure 3-1 Sequence structure

[Figure 3-1 Full Alternative Text](#)

The type of code shown in [Figure 3-1](#) is called a *sequence structure*, because the statements are executed in sequence, without branching off in another direction. Programs often need more than one path of execution, however. Many algorithms require a program to execute some statements only under certain circumstances. This can be accomplished with a *decision structure*.

In a decision structure's simplest form, a specific action is taken only when a condition exists. If the

condition does not exist, the action is not performed. The flowchart in [Figure 3-2](#) shows the logic of a decision structure. The diamond symbol represents a yes/no question or a true/false condition. If the answer to the question is yes (or if the condition is true), the program flow follows one path, which leads to an action being performed. If the answer to the question is no (or the condition is false), the program flow follows another path, which skips the action.




Figure 3-2 Simple decision structure logic

[Figure 3-2 Full Alternative Text](#)

In the flowchart, the action “Wear a coat” is performed only when it is cold outside. If it is not cold outside, the action is skipped. The action is *conditionally executed* because it is performed only when a certain condition (cold outside) exists. [Figure 3-3](#) shows a more elaborate flowchart, where three actions are taken only when it is cold outside.




Figure 3-3 Three-action decision structure logic

[Figure 3-3 Full Alternative Text](#)

One way to code a decision structure in Java is with the `if` statement. Here is the general format of the `if` statement:

```
if (BooleanExpression)
    statement;
```

The `if` statement is simple in the way it works: The *BooleanExpression* that appears inside the parentheses must be a `boolean` expression. A *boolean expression* is one that is either `true` or `false`. If the `boolean` expression is `true`, the very next statement is executed. Otherwise, it is skipped. The `statement` is conditionally executed because it executes only under the condition that the expression in the parentheses is `true`.

Using Relational Operators to Form Conditions

Typically, the `boolean` expression that is tested by an `if` statement is formed with a relational operator. A *relational operator* determines whether a specific relationship exists between two values. For example, the greater than operator (`>`) determines whether one value is greater than another. The equal to operator (`==`) determines whether two values are equal. [Table 3-1](#) lists all of the Java relational operators.

Table 3-1 Relational operators

Relational Operators (in Order of Precedence)	Meaning
>	Greater than
<	Less than
>=	Greater than or equal to
<=	Less than or equal to
==	Equal to

`!=`

Not equal to

All of the relational operators are binary, which means they use two operands. Here is an example of an expression using the greater than operator:

```
length > width
```

This expression determines whether `length` is greater than `width`. If `length` is greater than `width`, the value of the expression is `true`. Otherwise, the value of the expression is `false`. Because the expression can be only `true` or `false`, it is a `boolean` expression. The following expression uses the less than operator to determine whether `length` is less than `width`:

```
length < width
```

[Table 3-2](#) shows examples of several `boolean` expressions that compare the variables `x` and `y`.

Table 3-2 `boolean` expressions using relational operators

Expression	Meaning
<code>x > y</code>	Is <code>x</code> greater than <code>y</code> ?
<code>x < y</code>	Is <code>x</code> less than <code>y</code> ?
<code>x >= y</code>	Is <code>x</code> greater than or equal to <code>y</code> ?
<code>x <= y</code>	Is <code>x</code> less than or equal to <code>y</code> ?
<code>x == y</code>	Is <code>x</code> equal to <code>y</code> ?
<code>x != y</code>	Is <code>x</code> not equal to <code>y</code> ?

Two of the operators, `>=` and `<=`, test for more than one relationship. The `>=` operator determines

whether the operand on its left is greater than or equal to the operand on its right. Assuming that `a` is 4, `b` is 6, and `c` is 4, both of the expressions `b >= a` and `a >= c` are `true`, but `a >= 5` is `false`. When using this operator, the `>` symbol must precede the `=` symbol, and there is no space between them. The `<=` operator determines whether the operand on its left is less than or equal to the operand on its right. Once again, assuming that `a` is 4, `b` is 6, and `c` is 4, both `a <= c` and `b <= 10` are `true`, but `b <= a` is `false`. When using this operator, the `<` symbol must precede the `=` symbol, and there is no space between them.

The `==` operator determines whether the operand on its left is equal to the operand on its right. If both operands have the same value, the expression is `true`. Assuming that `a` is 4, the expression `a == 4` is `true` and the expression `a == 2` is `false`. Notice the equality operator is two `=` symbols together. Don't confuse this operator with the assignment operator, which is one `=` symbol.

The `!=` operator is the not equal operator. It determines whether the operand on its left is not equal to the operand on its right, which is the opposite of the `==` operator. As before, assuming `a` is 4, `b` is 6, and `c` is 4, both `a != b` and `b != c` are `true` because `a` is not equal to `b` and `b` is not equal to `c`. However, `a != c` is `false` because `a` is equal to `c`.

Putting It All Together

Let's look at an example of the `if` statement:

```
if (sales > 50000)
    bonus = 500.0;
```

This statement uses the `>` operator to determine whether `sales` is greater than 50,000. If the expression `sales > 50000` is `true`, the variable `bonus` is assigned 500.0. If the expression is `false`, however, the assignment statement is skipped. The program in [Code Listing 3-1](#) shows another example. The user enters three test scores and the program calculates their average. If the average is greater than 95, the program congratulates the user on obtaining a high score.

Code Listing 3-1 (AverageScore.java)

```
1 import java.util.Scanner;
2
3 /**
4  * This program demonstrates the if statement.
5  */
6
7 public class AverageScore
8 {
9     public static void main(String[] args)
10    {
11        double score1,    // Score #1
12              score2,    // Score #2
13              score3,    // Score #3
```

```

14         average;    // Average score
15
16     // Create a Scanner object to read input.
17     Scanner keyboard = new Scanner(System.in);
18
19     System.out.println("This program averages 3 test scores.");
20
21     // Get the first score.
22     System.out.print("Enter score #1: ");
23     score1 = keyboard.nextDouble();
24
25     // Get the second score.
26     System.out.print("Enter score #2: ");
27     score2 = keyboard.nextDouble ();
28
29     // Get the third score.
30     System.out.print("Enter score #3: ");
31     score3 = keyboard.nextDouble ();
32
33     // Calculate and display the average score.
34     average = (score1 + score2 + score3) / 3.0;
35     System.out.println("The average is " + average);
36
37     // If the average is higher than 95, congratulate the user.
38     if (average > 95)
39         System.out.println("That's a great score!");
40     }
41 }

```

Program Output with Example Input Shown in Bold

```

This program averages 3 test scores.
Enter score #1: 82 [Enter]
Enter score #2: 76 [Enter]
Enter score #3: 91 [Enter]
The average is 83.0

```

Program Output with Example Input Shown in Bold

```

This program averages 3 test scores.
Enter score #1: 97 [Enter]
Enter score #2: 94 [Enter]
Enter score #3: 100 [Enter]
The average is 97.0
That's a great score!

```


The code in lines 38 and 39 causes the congratulatory message to be printed:

```
if (average > 95)
    System.out.println("That's a great score!");
```

[Figure 3-4](#) shows the logic of this `if` statement.




Figure 3-4 Logic of the `if` statement

[Table 3-3](#) shows other examples of `if` statements and their outcomes.

Table 3-3 Other examples of `if` statements

Statement	Outcome
<pre>if (hours > 40) overTime = true;</pre>	If <code>hours</code> is greater than 40, assigns <code>true</code> to the boolean variable <code>overTime</code> .
<pre>if (value < 32) System.out.println("Invalid number");</pre>	If <code>value</code> is less than 32, displays the message "Invalid number".

Programming Style and the `if` Statement

Even though an `if` statement usually spans more than one line, it is really one long statement. For instance, the following `if` statements are identical except for the style in which they are written:

```
if (value > 32)
    System.out.println("Invalid number.");
if (value > 32) System.out.println("Invalid number.");
```

In both of these examples, the compiler considers the `if` statement and the conditionally executed statement as one unit, with a semicolon properly placed at the end. Indentations and spacing are for the human readers of a program, not the compiler. Here are two important style rules you should adopt for writing `if` statements:

- The conditionally executed statement should appear on the line after the `if` statement.
- The conditionally executed statement should be indented one level from the `if` statement.

In most editors, each time you press the tab key, you are indenting one level. By indenting the conditionally executed statement, you are causing it to stand out visually. This is so you can tell at a glance what part of the program the `if` statement executes. This is a standard way of writing `if` statements and is the method you should use.

Be Careful with Semicolons

You do not put a semicolon after the `if (expression)` portion of an `if` statement, as illustrated in [Figure 3-5](#). This is because the `if` statement isn't complete without its conditionally executed statement.

```
if (expression)
    statement;
```

No semicolon here

Semicolon goes here

Figure 3-5 Do not prematurely terminate an `if` statement with a semicolon

If you prematurely terminate an `if` statement with a semicolon, the compiler will not display an error message, but will assume that you are placing a *null statement* there. The null statement, which is an empty statement that does nothing, will become the conditionally executed statement. The statement that you intended to be conditionally executed will be disconnected from the `if` statement and will always execute.

For example, look at the following code:

```
int x = 0, y = 10;

// The following if statement is prematurely
```

```
// terminated with a semicolon.  
if (x > y);  
    System.out.println(x + " is greater than " + y);
```

The `if` statement in this code is prematurely terminated with a semicolon. Because the `println` statement is not connected to the `if` statement, it will always execute.

Having Multiple Conditionally Executed Statements

The previous examples of the `if` statement conditionally execute a single statement. The `if` statement can also conditionally execute a group of statements, as long as they are enclosed in a set of braces. Enclosing a group of statements inside braces creates a *block* of statements. Here is an example:

```
if (sales > 50000)  
{  
    bonus = 500.0;  
    commissionRate = 0.12;  
    daysOff += 1;  
}
```

If `sales` is greater than 50,000, this code will execute all three of the statements inside the braces, in the order they appear. If the braces are accidentally left out, however, the `if` statement conditionally executes only the very next statement. [Figure 3-6](#) illustrates this.




Figure 3-6 An `if` statement missing its braces

[Figure 3-6 Full Alternative Text](#)

Flags

A flag is a `boolean` variable that signals when some condition exists in the program. When the flag variable is set to `false`, it indicates the condition does not yet exist. When the flag variable is set to `true`, it means the condition does exist.

For example, suppose a program similar to the previous test averaging program has a `boolean` variable named `highScore`. The variable might be used to signal that a high score has been achieved by the following code:

```
if (average > 95)
    highScore = true;
```

Later, the same program might use code similar to the following to test the `highScore` variable, in order to determine whether a high score has been achieved:

```
if (highScore)
    System.out.println("That's a high score!");
```

You will find flag variables useful in many circumstances, and we will come back to them in future chapters.

Comparing Characters

You can use the relational operators to test character data as well as numbers. For example, assuming that `ch` is a `char` variable, the following code segment uses the `==` operator to compare it to the character 'A':

```
if (ch == 'A')
    System.out.println("The letter is A.");
```

The `!=` operator can also be used with characters to test for inequality. For example, the following statement determines whether the `char` variable `ch` is not equal to the letter 'A':

```
if (ch != 'A')
    System.out.println("Not the letter A.");
```

You can also use the `>`, `<`, `>=`, and `<=` operators to compare characters. Computers do not actually store characters, such as A, B, C, and so forth, in memory. Instead, they store numeric codes that represent the characters. Recall from [Chapter 2](#) that Java uses Unicode, which is a set of numbers that represents all the letters of the alphabet (both lowercase and uppercase), the printable digits 0 through 9, punctuation symbols, and special characters. When a character is stored in memory, it is actually the Unicode number that is stored. When the computer is instructed to print the value on the screen, it displays the character that corresponds with the numeric code.

NOTE:

Unicode is an international encoding system that is extensive enough to represent all the characters of all the world's alphabets.

In Unicode, letters are arranged in alphabetic order. Because 'A' comes before 'B', the numeric code for the character 'A' is less than the code for the character 'B'. (The code for 'A' is 65 and the code for 'B' is 66. Appendix A, available for download from this book's companion Web site, lists the codes for all of the printable English characters.) In the following `if` statement, the `boolean` expression `'A' < 'B'` is `true`:

```
if ('A' < 'B')
    System.out.println("A is less than B.");
```

In Unicode, the uppercase letters come before the lowercase letters, so the numeric code for 'A' (65) is less than the numeric code for 'a' (97). In addition, the space character (code 32) comes before all the alphabetic characters.



MyProgrammingLab™ www.myprogramminglab.com

1. 3.1 Write an `if` statement that assigns 0 to `x` when `y` is equal to 20.
2. 3.2 Write an `if` statement that multiplies `payRate` by 1.5 if `hours` is greater than 40.
3. 3.3 Write an `if` statement that assigns 0.2 to `commission` if `sales` is greater than or equal to 10000.
4. 3.4 Write an `if` statement that sets the variable `fees` to 50 if the `boolean` variable `max` is `true`.
5. 3.5 Write an `if` statement that assigns 20 to the variable `y` and assigns 40 to the variable `z` if the variable `x` is greater than 100.
6. 3.6 Write an `if` statement that assigns 0 to the variable `b` and assigns 1 to the variable `c` if the variable `a` is less than 10.
7. 3.7 Write an `if` statement that displays "Goodbye" if the variable `myCharacter` contains the character 'D'.

3.2 The `if-else` Statement

CONCEPT: The `if-else` statement will execute one group of statements if its `boolean` expression is `true`, or another group if its `boolean` expression is `false`.

The `if-else` statement is an expansion of the `if` statement. Here is its format:



The `if-else` Statement

```
if (BooleanExpression)
    statement or block
else
    statement or block
```

Like the `if` statement, a `boolean` expression is evaluated. If the expression is `true`, a statement or block of statements is executed. If the expression is `false`, however, a separate group of statements is executed. The program in [Code Listing 3-2](#) uses the `if-else` statement to handle a classic programming problem: division by zero. In Java, a program crashes when it divides an integer by 0. When a floating-point value is divided by 0, the program doesn't crash. Instead, the special value `Infinity` is produced as the result of the division.

Code Listing 3-2 (`Division.java`)

```
1 import java.util.Scanner; // Needed for the Scanner class
2
3 /**
4     This program demonstrates the if-else statement.
5 */
6
7 public class Division
8 {
9     public static void main(String[] args)
10    {
11        double number1, number2; // Division operands
12        double quotient;         // Result of division
13
14        // Create a Scanner object for keyboard input.
15        Scanner keyboard = new Scanner(System.in);
16
17        // Get the first number.
18        System.out.print("Enter a number: ");
19        number1 = keyboard.nextDouble();
20
21        // Get the second number.
```

```

22     System.out.print("Enter another number: ");
23     number2 = keyboard.nextDouble();
24
25     if (number2 == 0)
26     {
27         System.out.println("Division by zero is not possible.");
28         System.out.println("Please run the program again and ");
29         System.out.println("enter a number other than zero.");
30     }
31     else
32     {
33         quotient = number1 / number2;
34         System.out.print("The quotient of " + number1);
35         System.out.print(" divided by " + number2);
36         System.out.println(" is " + quotient);
37     }
38 }
39 }

```

Program Output with Example Input Shown in Bold

```

Enter a number: 10 [Enter]
Enter another number: 0 [Enter]
Division by zero is not possible.
Please run the program again and
enter a number other than zero.

```

Program Output with Example Input Shown in Bold

```

Enter a number: 10 [Enter]
Enter another number: 5 [Enter]
The quotient of 10 divided by 5 is 2.0

```

The value of `number2` is tested before the division is performed. If the user enters 0, the block of statements controlled by the `if` clause executes, displaying a message that indicates the program cannot perform division by zero. Otherwise, the `else` clause takes control, which divides `number1` by `number2` and displays the result. [Figure 3-7](#) shows the logic of the `if-else` statement.




Figure 3-7 Logic of the `if-else` statement

[Figure 3-7 Full Alternative Text](#)



MyProgrammingLab™ www.myprogramminglab.com

1. 3.8 Write an `if-else` statement that assigns 20 to the variable `y` if the variable `x` is greater than 100. Otherwise, it should assign 0 to the variable `y`.
2. 3.9 Write an `if-else` statement that assigns 1 to `x` when `y` is equal to 100. Otherwise, it should assign 0 to `x`.
3. 3.10 Write an `if-else` statement that assigns 0.10 to `commission` unless `sales` is greater than or equal to 50000.0, in which case it assigns 0.2 to `commission`.
4. 3.11 Write an `if-else` statement that assigns 0 to the variable `b` and assigns 1 to the variable `c` if the variable `a` is less than 10. Otherwise, it should assign -99 to the variable `b` and assign 0 to the variable `c`.

3.3 Nested if Statements

CONCEPT: To test more than one condition, an `if` statement can be nested inside another `if` statement.

Sometimes an `if` statement must be nested inside another `if` statement. For example, consider a banking program that determines whether a bank customer qualifies for a special, low interest rate on a loan. To qualify, two conditions must exist: (1) the customer's salary must be at least \$50,000, and (2) the customer must have held his or her current job for at least two years. [Figure 3-8](#) shows a flowchart for an algorithm that could be used in such a program.




Figure 3-8 Logic of nested if statements

[Figure 3-8 Full Alternative Text](#)

If we follow the flow of execution in the flowchart, we see that the expression `salary >= 30000` is tested. If this expression is false, there is no need to perform further tests; we know that the customer does not qualify for the special interest rate. If the expression is true, however, we need to test the second condition. This is done with a nested decision structure that tests the expression `yearsOnJob`

>= 2. If this expression is true, then the customer qualifies for the special interest rate. If this expression is false, then the customer does not qualify. [Code Listing 3-3](#) shows the complete program.

Code Listing 3-3 (LoanQualifier.java)

```
1 import java.util.Scanner;
2
3 /**
4  * This program demonstrates a nested if statement.
5  */
6
7 public class LoanQualifier
8 {
9     public static void main(String[] args)
10    {
11        double salary;        // Annual salary
12        double yearsOnJob;    // Years at current job
13
14        // Create a Scanner object for keyboard input.
15        Scanner keyboard = new Scanner(System.in);
16
17        // Get the user's annual salary.
18        System.out.print("Enter your annual salary: ");
19        salary = keyboard.nextDouble();
20
21        // Get the number of years at the current job.
22        System.out.print("Enter the number of years " +
23                        "at your current job: ");
24        yearsOnJob = keyboard.nextDouble();
25
26        // Determine whether the user qualifies for the loan.
27        if (salary >= 50000)
28        {
29            if (yearsOnJob >= 2)
30            {
31                System.out.println("You qualify for the loan.");
32            }
33            else
34            {
35                System.out.println("You must have been on your " +
36                                "current job for at least " +
37                                "two years to qualify.");
38            }
39        }
40        else
41        {
42            System.out.println("You must earn at least " +
43                                "$50,000 per year to qualify.");
44        }
45    }
46 }
```


Not only do the braces make the statements easier to read, but they also help in debugging code. When debugging a program with nested `if-else` statements, it's important to know which `if` clause each `else` clause belongs to. The rule for matching `else` clauses with `if` clauses is this: An `else` clause goes with the closest previous `if` clause that doesn't already have its own `else` clause. This is easy to see when the conditionally executed statements are enclosed in braces and are properly indented, as shown in [Figure 3-9](#). Each `else` clause lines up with the `if` clause it belongs to. These visual cues are important because nested `if` statements can be very long and complex.




Figure 3-9 Alignment of `if` and `else` clauses

[Figure 3-9 Full Alternative Text](#)

Testing a Series of Conditions

In the previous example, you saw how a program can use nested decision structures to test more than one condition. It is not uncommon for a program to have a series of conditions to test, and then perform an action depending on which condition is true. One way to accomplish this is to have a decision structure with numerous other decision structures nested inside it. For example, consider the program presented in the following *In the Spotlight* section.



Multiple Nested Decision Structures

Suppose one of your professors uses the following 10-point grading scale for exams:

Test Score Grade

90 and above A

80–89 B

70–79 C

60–69 D

Below 60 F

Your professor has asked you to write a program that will allow a student to enter a test score and then display the grade for that score. Here is the algorithm that you will use:

- *Ask the user to enter a test score.*
- *Determine the grade in the following manner:*
- *If the score is less than 60, then the grade is F.*
- *Otherwise, if the score is less than 70, then the grade is D.*
- *Otherwise, if the score is less than 80, then the grade is C.*
- *Otherwise, if the score is less than 90, then the grade is B.*
- *Otherwise, the grade is A.*

You decide that the process of determining the grade will require several nested decision structures, as shown in [Figure 3-10](#). [Code Listing 3-4](#) shows the complete program. The code for the nested decision structures is in lines 23 through 51. [Figures 3-11](#) and [3-12](#) show what happens in two different sessions with the program.




Figure 3-10 Nested decision structure to determine a grade

[Figure 3-10 Full Alternative Text](#)

Code Listing 3-4 (NestedDecision.java)

```

1 import javax.swing.JOptionPane; // Needed for JOptionPane
2
3 /**
4  This program asks the user to enter a numeric test
5  score and displays a letter grade for the score. The
6  program uses nested decision structures
7  to determine the grade.


```

```

8  */
9
10 public class NestedDecision
11 {
12     public static void main(String[] args)
13     {
14         int testScore; // Numeric test score
15         String input; // To hold the user's input
16
17         // Get the numeric test score.
18         input = JOptionPane.showInputDialog("Enter your numeric " +
19             "test score and I will tell you the grade: ");
20         testScore = Integer.parseInt(input);
21
22         // Display the grade.
23         if (testScore < 60)
24         {
25             JOptionPane.showMessageDialog(null, "Your grade is F.");
26         }
27         else
28         {
29             if (testScore < 70)
30             {
31                 JOptionPane.showMessageDialog(null, "Your grade is D.");
32             }
33             else
34             {
35                 if (testScore < 80)
36                 {
37                     JOptionPane.showMessageDialog(null, "Your grade is C.");
38                 }
39                 else
40                 {
41                     if (testScore < 90)
42                     {
43                         JOptionPane.showMessageDialog(null, "Your grade is B.");
44                     }
45                     else
46                     {
47                         JOptionPane.showMessageDialog(null, "Your grade is A.");
48                     }
49                 }
50             }
51         }
52         System.exit(0);
53     }
54 }

```

This input dialog box appears first. The user enters 80 and then clicks the OK button.



This message dialog box appears next.





Figure 3-11 Interaction with the NestedDecision program

(Oracle Corporate Counsel)

[Figure 3-11 Full Alternative Text](#)

This input dialog box appears first. The user enters 72 and then clicks the OK button.



This message dialog box appears next.




Figure 3-12 Interaction with the NestedDecision program

(Oracle Corporate Counsel)

[Figure 3-12 Full Alternative Text](#)



Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 3.12 Write nested `if` statements that perform the following test: If `amount1` is greater than 10 and `amount2` is less than 100, display the greater of the two.
2. 3.13 Write code that tests the variable `x` to determine whether it is greater than 0. If `x` is greater than 0, the code should test the variable `y` to determine whether it is less than 20. If `y` is less than 20, the code should assign 1 to the variable `z`. If `y` is not less than 20, the code should assign 0 to the variable `z`.

3.4 The if-else-if Statement

CONCEPT: The `if-else-if` statement tests a series of conditions. It is often simpler to test a series of conditions with the `if-else-if` statement than with a set of nested `if-else` statements.

Even though [Code Listing 3-4](#) is a simple example, the logic of the nested decision structure is fairly complex. You can alternatively test a series of conditions using the `if-else-if` statement. The `if-else-if` statement makes certain types of nested decision logic simpler to write. Here is the general format of the `if-else-if` statement:



The `if- else-if` Statement

```
if (expression_1)
{
    statement
    statement
    etc.
}
else if (expression_2)
{
    statement
    statement
    etc.
}
Insert as many else if clauses as necessary
else
{
    statement
    statement
    etc.
}
```

If expression_1 is true these statements are executed, and the rest of the structure is ignored.

Otherwise, if expression_2 is true these statements are executed, and the rest of the structure is ignored.

These statements are executed if none of the expressions above are true.

[3.4-4 Full Alternative Text](#)

When the statement executes, `expression_1` is tested. If `expression_1` is true, the block of statements that immediately follows is executed, and the rest of the structure is ignored. If `expression_1` is false, however, the program jumps to the next `else if` clause and tests `expression_2`. If it is true, the block of statements that immediately follows is executed, and then the

rest of the structure is ignored. This process continues, from the top of the structure to the bottom, until one of the expressions is found to be true. If none of the expressions are true, the last `else` clause takes over and the block of statements immediately following it is executed.

The last `else` clause, which does not have an `if` statement following it, is referred to as the *trailing else*. The trailing `else` is optional, but in most cases you will use it.

NOTE:

The general format shows braces surrounding each block of conditionally executed statements. As with other forms of the `if` statement, the braces are required only when more than one statement is conditionally executed.

[Code Listing 3-5](#) shows an example of the `if-else-if` statement. This program is a modification of [Code Listing 3-4](#), which appears in the previous *In the Spotlight* section. The output of this program is the same as [Code Listing 3-4](#).

Code Listing 3-5 (TestResults.Java)

```
1 import javax.swing.JOptionPane;      // Needed for JOptionPane
2
3 /**
4     This program asks the user to enter a numeric test
5     score and displays a letter grade for the score. The
6     program uses an if-else-if statement to determine
7     the letter grade.
8 */
9
10 public class TestResults
11 {
12     public static void main(String[] args)
13     {
14         int testScore;      // Numeric test score
15         String input;       // To hold the user's input
16
17         // Get the numeric test score.
18         input = JOptionPane.showInputDialog("Enter your numeric " +
19             "test score and I will tell you the grade: ");
20         testScore = Integer.parseInt(input);
21
22         // Display the grade.
23         if (testScore < 60)
24             JOptionPane.showMessageDialog(null, "Your grade is F.");
25         else if (testScore < 70)
26             JOptionPane.showMessageDialog(null, "Your grade is D.");
27         else if (testScore < 80)
28             JOptionPane.showMessageDialog(null, "Your grade is C.");
29         else if (testScore < 90)
```

```

30         JOptionPane.showMessageDialog(null, "Your grade is B.");
31     else
32         JOptionPane.showMessageDialog(null, "Your grade is A.");
33
34     System.exit(0);
35 }
36 }

```

Let's analyze how the `if-else-if` statement in lines 23 through 32 works. First, the expression `testScore < 60` is tested in line 23:

```

→ if (testScore < 60)
    JOptionPane.showMessageDialog(null, "Your grade is F.");
else if (testScore < 70)
    JOptionPane.showMessageDialog(null, "Your grade is D.");
else if (testScore < 80)
    JOptionPane.showMessageDialog(null, "Your grade is C.");
else if (testScore < 90)
    JOptionPane.showMessageDialog(null, "Your grade is B.");
else
    JOptionPane.showMessageDialog(null, "Your grade is A.");

```

If `testScore` is less than 60, the message "Your grade is F." is displayed and the rest of the `if-else-if` statement is skipped. If `testScore` is not less than 60, the `else` clause in line 25 takes over and causes the next `if` statement to be executed:

```

if (testScore < 60)
    JOptionPane.showMessageDialog(null, "Your grade is F.");
→ else if (testScore < 70)
    JOptionPane.showMessageDialog(null, "Your grade is D.");
else if (testScore < 80)
    JOptionPane.showMessageDialog(null, "Your grade is C.");
else if (testScore < 90)
    JOptionPane.showMessageDialog(null, "Your grade is B.");
else
    JOptionPane.showMessageDialog(null, "Your grade is A.");

```

The first `if` statement handled all of the grades less than 60, so when this `if` statement executes, `testScore` will have a value of 60 or greater. If `testScore` is less than 70, the message "Your grade is D." is displayed and the rest of the `if-else-if` statement is skipped. This chain of events continues until one of the expressions is found to be true, or the last `else` clause at the end of the statement is encountered.


Notice the alignment and indentation that are used with the `if-else-if` statement: The starting `if` clause, the `else if` clauses, and the trailing `else` clause are all aligned, and the conditionally executed statements are indented.

Using the Trailing `else` to Catch Errors

The trailing `else` clause, which appears at the end of the `if-else-if` statement, is optional, but in

many situations you will use it to catch errors. For example, [Code Listing 3-5](#) will assign the grade 'A' to any test score that is 90 or greater. What if the highest possible test score is 100? We can modify the code as shown in [Code Listing 3-6](#) so the trailing `else` clause catches any value greater than 100 and displays an error message. [Figure 3-13](#) shows what happens when the user enters a test score that is greater than 100.

This input dialog box appears first. The user enters 105 and then clicks the OK button.



This message dialog box appears next.




Figure 3-13 Interaction with the NestedDecision program

(Oracle Corporate Counsel)

[Figure 3-13 Full Alternative Text](#)

Code Listing 3-6 (TrailingElse.java)

```
1 import javax.swing.JOptionPane; // Needed for JOptionPane
2
3 /**
4  * This program asks the user to enter a numeric test
5  * score and displays a letter grade for the score. The
6  * program displays an error message if an invalid
7  * numeric score is entered.
8  */
9
10 public class TrailingElse
11 {
12     public static void main(String[] args)
13     {
14         int testScore;        // Numeric test score
15         String input;          // To hold the user's input
16
17         // Get the numeric test score.
18         input = JOptionPane.showInputDialog("Enter your numeric " +
```

```

19         "test score and I will tell you the grade:");
20     testScore = Integer.parseInt(input);
21
22     // Display the grade.
23     if (testScore < 60)
24         JOptionPane.showMessageDialog(null, "Your grade is F.");
25     else if (testScore < 70)
26         JOptionPane.showMessageDialog(null, "Your grade is D.");
27     else if (testScore < 80)
28         JOptionPane.showMessageDialog(null, "Your grade is C.");
29     else if (testScore < 90)
30         JOptionPane.showMessageDialog(null, "Your grade is B.");
31     else if (testScore <= 100)
32         JOptionPane.showMessageDialog(null, "Your grade is A.");
33     else
34         JOptionPane.showMessageDialog(null, "Invalid score.");
35
36     System.exit(0);
37 }
38 }

```

The `if-else-if` Statement Compared to a Nested Decision Structure

You never have to use the `if-else-if` statement because its logic can be coded with nested `if-else` statements. However, a long series of nested `if-else` statements has two particular disadvantages when you are debugging code:

- The code can grow complex and become difficult to understand.
- Because indenting is important in nested statements, a long series of nested `if-else` statements can become too long to be displayed on the computer screen without horizontal scrolling. Also, long statements tend to “wrap around” when printed on paper, making the code even more difficult to read.

The logic of an `if-else-if` statement is usually easier to follow than that of a long series of nested `if-else` statements. And, because all of the clauses are aligned in an `if-else-if` statement, the lengths of the lines in the statement tend to be shorter.



MyProgrammingLab™ www.myprogramminglab.com

1. 3.14 What will the following program display?

```
public class CheckPoint
```

```

{
    public static void main(String[] args)
    {
        int funny = 7, serious = 15;
        funny = serious % 2;
        if (funny != 1)
        {
            funny = 0;
            serious = 0;
        }
        else if (funny == 2)
        {
            funny = 10;
            serious = 10;
        }
        else
        {
            funny = 1;
            serious = 1;
        }
        System.out.println(funny + " " + serious);
    }
}

```

2. 3.15 The following program is used in a bookstore to determine how many discount coupons a customer gets. Complete the table that appears after the program.

```

import javax.swing.JOptionPane;
public class CheckPoint
{
    public static void main(String[] args)
    {
        int books, coupons;
        String input;
        input = JOptionPane.showInputDialog("How many books " +
                                           "are being purchased? ");
        books = Integer.parseInt(input);
        if (books < 1)
            coupons = 0;
        else if (books < 3)
            coupons = 1;
        else if (books < 5)
            coupons = 2;
        else
            coupons = 3;
        JOptionPane.showMessageDialog(null,
                                     "The number of coupons to give is " +
                                     coupons);
        System.exit(0);
    }
}

```

If the customer purchases this many books . . . this many coupons are given.

1

2

3

4

5

10

3.5 Logical Operators

CONCEPT: Logical operators connect two or more relational expressions into one or reverse the logic of an expression.

Java provides two binary logical operators, `&&` and `||`, which are used to combine two `boolean` expressions into a single expression. It also provides the unary `!` operator, which reverses the truth of a `boolean` expression. [Table 3-4](#) describes these logical operators.

Table 3-4 Logical operators

Operator Meaning		Effect
<code>&&</code>	AND	Connects two <code>boolean</code> expressions into one. Both expressions must be <code>true</code> for the overall expression to be <code>true</code> .
<code> </code>	OR	Connects two <code>boolean</code> expressions into one. One or both expressions must be <code>true</code> for the overall expression to be <code>true</code> . It is only necessary for one to be <code>true</code> , and it does not matter which one.
<code>!</code>	NOT	The <code>!</code> operator reverses the truth of a <code>boolean</code> expression. If it is applied to an expression that is <code>true</code> , the operator returns <code>false</code> . If it is applied to an expression that is <code>false</code> , the operator returns <code>true</code> .

[Table 3-5](#) shows examples of several `boolean` expressions that use logical operators.

Table 3-5 `boolean` expressions using logical operators

Expression	Meaning
<code>x > y && a < b</code>	Is <code>x</code> greater than <code>y</code> AND is <code>a</code> less than <code>b</code> ?

`x == y || x == z` Is `x` equal to `y` OR is `x` equal to `z`?

`!(x > y)` Is the expression `x > y` NOT true?

Let's take a close look at each of these operators.

The && Operator

The `&&` operator is known as the logical AND operator. It takes two `boolean` expressions as operands and creates a `boolean` expression that is `true` only when both subexpressions are `true`. Here is an example of an `if` statement that uses the `&&` operator:

```
if (temperature < 20 && minutes > 12)
{
    System.out.println("The temperature is in the " +
                       "danger zone.");
}
```

In this statement the two `boolean` expressions `temperature < 20` and `minutes > 12` are combined into a single expression. The message will be displayed only if `temperature` is less than 20 AND `minutes` is greater than 12. If either `boolean` expression is `false`, the entire expression is `false` and the message is not displayed.

[Table 3-6](#) shows a truth table for the `&&` operator. The truth table lists expressions showing all the possible combinations of `true` and `false` connected with the `&&` operator. The resulting values of the expressions are also shown.

Table 3-6 Truth table for the && operator

Expression	Value of the Expression

<code>true && false</code>	<code>false</code>

<code>false && true</code>	<code>false</code>

<code>false && false</code>	<code>false</code>

```
true && true    true
```

As the table shows, both sides of the `&&` operator must be `true` for the operator to return a `true` value.

The `&&` operator performs *short-circuit evaluation*. Here's how it works: If the expression on the left side of the `&&` operator is `false`, the expression on the right side will not be checked. Because the entire expression is `false` if only one of the subexpressions is `false`, it would waste CPU time to check the remaining expression. So, when the `&&` operator finds that the expression on its left is `false`, it short-circuits and does not evaluate the expression on its right.

The `&&` operator can be used to simplify programs that otherwise would use nested `if` statements. The program in [Code Listing 3-7](#) is a different version of the `LoanQualifier` program in [Code Listing 3-3](#), written to use the `&&` operator.

Code Listing 3-7 (LogicalAnd.java)

```
1 import java.util.Scanner;
2
3 /**
4  * This program demonstrates the logical && operator.
5  */
6
7 public class LogicalAnd
8 {
9     public static void main(String[] args)
10    {
11        double salary;           // Annual salary
12        double yearsOnJob;       // Years at current job
13
14        // Create a Scanner object for keyboard input.
15        Scanner keyboard = new Scanner(System.in);
16
17        // Get the user's annual salary.
18        System.out.print("Enter your annual salary: ");
19        salary = keyboard.nextDouble();
20
21        // Get the number of years at the current job.
22        System.out.print("Enter the number of years " +
23                        "at your current job: ");
24        yearsOnJob = keyboard.nextDouble();
25
26        // Determine whether the user qualifies for the loan.
27        if (salary >= 50000 && yearsOnJob >= 2)
28        {
29            System.out.println("You qualify for the loan.");
30        }
31        else
32        {
33            System.out.println("You do not qualify.");
```

```
34     }
35 }
36 }
```

Program Output with Example Input Shown in Bold

```
Enter your annual salary: 55000.00 [Enter]
Enter the number of years at your current job: 1 [Enter]
You do not qualify.
```

Program Output with Example Input Shown in Bold

```
Enter your annual salary: 55000.00 [Enter]
Enter the number of years at your current job: 4 [Enter]
You qualify for the loan.
```

The message "You qualify for the loan." is displayed only when both the expressions `salary >= 30000` and `yearsOnJob >= 2` are true. If either of these expressions is false, the message "You do not qualify for the loan." is displayed.

You can also use logical operators with `boolean` variables. For example, assuming that `isValid` is a `boolean` variable, the following `if` statement determines whether `isValid` is true and `x` is greater than 90.

```
if (isValid && x > 90)
```

The || Operator

The `||` operator is known as the logical OR operator. It takes two `boolean` expressions as operands and creates a `boolean` expression that is true when either of the subexpressions is true. Here is an example of an `if` statement that uses the `||` operator:

```
if (temperature < 20 || temperature > 100)
{
    System.out.println("The temperature is in the " +
                       "danger zone.");
}
```

The message will be displayed if `temperature` is less than 20 OR `temperature` is greater than 100. If either relational test is true, the entire expression is true.

[Table 3-7](#) shows a truth table for the `||` operator.

Table 3-7 Truth table for the `||` operator

Expression	Value
<code>true false</code>	<code>true</code>
<code>false true</code>	<code>true</code>
<code>false false</code>	<code>false</code>
<code>true true</code>	<code>true</code>

All it takes for an OR expression to be `true` is for one side of the `||` operator to be `true`. It doesn't matter if the other side is `false` or `true`. Like the `&&` operator, the `||` operator performs short-circuit evaluation. If the expression on the left side of the `||` operator is `true`, the expression on the right side will not be checked. Because it is necessary for only one of the expressions to be `true`, it would waste CPU time to check the remaining expression.

The program in [Code Listing 3-8](#) is a different version of the previous program, shown in [Code Listing 3-7](#). This version uses the `||` operator to determine whether `salary >= 30000` is `true` OR `yearsOnJob >= 2` is `true`. If either expression is `true`, then the person qualifies for the loan.

Code Listing 3-8 (`LogicalOr.java`)

```
1 import java.util.Scanner;
2
3 /**
4  * This program demonstrates the logical || operator.
5  */
6
7 public class LogicalOr
8 {
9     public static void main(String[] args)
10    {
11        double salary;           // Annual salary
12        double yearsOnJob;       // Years at current job
13
14        // Create a Scanner object for keyboard input.
15        Scanner keyboard = new Scanner(System.in);
```

```

16
17 // Get the user's annual salary.
18 System.out.print("Enter your annual salary: ");
19 salary = keyboard.nextDouble();
20
21 // Get the number of years at the current job.
22 System.out.print("Enter the number of years " +
23                 "at your current job: ");
24 yearsOnJob = keyboard.nextDouble();
25
26 // Determine whether the user qualifies for the loan.
27 if (salary >= 50000 || yearsOnJob >= 2)
28 {
29     System.out.println("You qualify for the loan.");
30 }
31 else
32 {
33     System.out.println("You do not qualify.");
34 }
35 }
36 }

```

Program Output with Example Input Shown in Bold

```

Enter your annual salary: 20000.00 [Enter]
Enter the number of years at your current job: 7 [Enter]
You qualify for the loan.

```

Program Output with Example Input Shown in Bold

```

Enter your annual salary: 55000.00 [Enter]
Enter the number of years at your current job: 1 [Enter]
You qualify for the loan.

```

Program Output with Example Input Shown in Bold

```

Enter your annual salary: 20000.00 [Enter]
Enter the number of years at your current job: 1 [Enter]
You do not qualify.

```

The ! Operator

The `!` operator performs a logical NOT operation. It is a unary operator that takes a `boolean` expression as its operand and reverses its logical value. In other words, if the expression is `true`, the `!` operator returns `false`, and if the expression is `false`, it returns `true`. Here is an `if` statement using the `!` operator:

```
if (!(temperature > 100))
    System.out.println("This is below the maximum temperature.");
```

First, the expression `(temperature > 100)` is tested and a value of either `true` or `false` is the result. Then the `!` operator is applied to that value. If the expression `(temperature > 100)` is `true`, the `!` operator returns `false`. If the expression `(temperature > 100)` is `false`, the `!` operator returns `true`. The previous code is equivalent to asking: “Is the temperature not greater than 100?”

[Table 3-8](#) shows a truth table for the `!` operator.

Table 3-8 Truth table for the ! operator

Expression Value	

<code>!true</code>	<code>false</code>
<code>!false</code>	<code>true</code>

The Precedence of Logical Operators

Like other operators, the logical operators have orders of precedence and associativity. [Table 3-9](#) shows the precedence of the logical operators, from highest to lowest.

Table 3-9 Logical operators in order of precedence

!

&&

||

The `!` operator has a higher precedence than many of Java's other operators. You should always enclose its operand in parentheses unless you intend to apply it to a variable or a simple expression with no other operators. For example, consider the following expressions (assume `x` is an `int` variable with a value stored in it):

```
!(x > 2)
!x > 2
```

The first expression applies the `!` operator to the expression `x > 2`. It is asking "is `x` not greater than 2?" The second expression, however, attempts to apply the `!` operator to `x` only. It is asking "is the logical complement of `x` greater than 2?" Because the `!` operator can only be applied to `boolean` expressions, this statement would cause a compiler error.

The `&&` and `||` operators rank lower in precedence than the relational operators, so precedence problems are less likely to occur. If you are unsure, however, it doesn't hurt to use parentheses anyway.

```
(a > b) && (x < y) is the same as a > b && x < y
(x == y) || (b > a) is the same as x == y || b > a
```

The logical operators evaluate their expressions from left to right. In the following expression, `a < b` is evaluated before `y == z`.

```
a < b || y == z
```

In the following expression, `y == z` is evaluated first, however, because the `&&` operator has higher precedence than `||`.

```
a < b || y == z && m > j
```

This expression is equivalent to the following:

```
(a < b) || ((y == z) && (m > j))
```

[Table 3-10](#) shows the precedence of all the operators we have discussed so far. This table includes the assignment, arithmetic, relational, and logical operators.

Table 3-10 Precedence of all operators discussed so far

Order of Precedence	Operators	Description
1	- (unary negation) !	Unary negation, logical NOT
2	* / %	Multiplication, division, modulus
3	+ -	Addition, subtraction
4	< > <= >=	Less than, greater than, less than or equal to, greater than or equal to
5	== !=	Equal to, not equal to
6	&&	Logical AND
7		Logical OR
8	= += -= *= /= %=	Assignment and combined assignment

Checking Numeric Ranges with Logical Operators

Sometimes you will need to write code that determines whether a numeric value is within a specific range of values or outside a specific range of values. When determining whether a number is inside a range, it's best to use the `&&` operator. For example, the following `if` statement checks the value in `x` to determine whether it is in the range of 20 through 40:

```
if (x >= 20 && x <= 40)
    System.out.println(x + " is in the acceptable range.");
```

The `boolean` expression in the `if` statement will be `true` only when `x` is greater than or equal to 20 AND less than or equal to 40. The value in `x` must be within the range of 20 through 40 for this

expression to be `true`.

When determining whether a number is outside a range, it's best to use the `||` operator. The following statement determines whether `x` is outside the range of 20 through 40:

```
if (x < 20 || x > 40)
    System.out.println(x + " is outside the acceptable range.");
```

It's important not to get the logic of these logical operators confused. For example, the `boolean` expression in the following `if` statement would never test `true`:

```
if (x < 20 && x > 40)
    System.out.println(x + " is outside the acceptable range.");
```

Obviously, `x` cannot be less than 20 and at the same time be greater than 40.



MyProgrammingLab™ www.myprogramminglab.com

1. 3.16 The following truth table shows various combinations of the values `true` and `false` connected by a logical operator. Complete the table by circling `T` or `F` to indicate whether the result of such a combination is `true` or `false`.

Logical Expression	Result (true or false)	
<code>true && false</code>	T	F
<code>true && true</code>	T	F
<code>false && true</code>	T	F
<code>false && false</code>	T	F
<code>true false</code>	T	F
<code>true true</code>	T	F
<code>false true</code>	T	F
<code>false false</code>	T	F
<code>!true</code>	T	F
<code>!false</code>	T	F

2. 3.17 Assume the variables `a = 2`, `b = 4`, and `c = 6`. Circle the `T` or `F` for each of the following conditions to indicate whether it is `true` or `false`.

<code>a == 4 b > 2</code>	T	F
<code>6 <= c && a > 3</code>	T	F
<code>1 != b && c != 3</code>	T	F
<code>a >= -1 a <= b</code>	T	F
<code>!(a > 2)</code>	T	F

3. 3.18 Write an `if` statement that displays the message "The number is valid" if the variable `speed` is within the range 0 through 200.
4. 3.19 Write an `if` statement that displays the message "The number is not valid" if the

variable `speed` is outside the range 0 through 200.

3.6 Comparing String Objects

CONCEPT: You cannot use relational operators to compare `String` objects. Instead you must use a `String` method.

You saw in the preceding sections how numeric values can be compared using the relational operators. You should not use the relational operators to compare `String` objects, however. Remember that a `String` object is referenced by a variable that contains the object's memory address. When you use a relational operator with the reference variable, the operator works on the memory address that the variable contains, not the contents of the `String` object. For example, suppose a program has the following declarations:

```
String name1 = "Mark";  
String name2 = "Mary";
```

And later, the same program has the following `if` statement:

```
if (name1 == name2)
```

The expression `name1 == name2` will be `false`, but not because the strings “Mark” and “Mary” are different. The expression will be `false` because the variables `name1` and `name2` reference different objects. [Figure 3-14](#) illustrates how the variables reference the `String` objects.




Figure 3-14 The `name1` and `name2` variables reference different `String` objects

[Figure 3-14 Full Alternative Text](#)

To compare the contents of two `String` objects correctly, you should use the `String` class's `equals` method. The general form of the method is as follows:

```
StringReference1.equals(StringReference2)
```

StringReference1 is a variable that references a `String` object, and *StringReference2* is another variable that references a `String` object. The method returns `true` if the two strings are equal, or `false` if they are not equal. Here is an example:

```
if (name1.equals(name2))
```

Assuming that `name1` and `name2` reference `String` objects, the expression in the `if` statement will return `true` if they are the same, or `false` if they are not the same. The program in [Code Listing 3-9](#) demonstrates.

Code Listing 3-9 (`StringCompare.java`)

```
1  /**
2   * This program correctly compares two String objects using
3   * the equals method.
4   */
5
6  public class StringCompare
7  {
8      public static void main(String[] args)
9      {
10         String name1 = "Mark",
11             name2 = "Mark",
12             name3 = "Mary";
13
14         // Compare "Mark" and "Mark"
15
16         if (name1.equals(name2))
17         {
18             System.out.println(name1 + " and " + name2 +
19                               " are the same.");
20         }
21         else
22         {
23             System.out.println(name1 + " and " + name2 +
24                               " are NOT the same.");
25         }
26
27         // Compare "Mark" and "Mary"
28
29         if (name1.equals(name3))
30         {
31             System.out.println(name1 + " and " + name3 +
32                               " are the same.");
33         }
34         else
35         {
36             System.out.println(name1 + " and " + name3 +
37                               " are NOT the same.");
38         }
39     }
```

```
40    }
```

Program Output

```
Mark and Mark are the same.  
Mark and Mary are NOT the same.
```

You can also compare `String` objects to string literals. Simply pass the string literal as the argument to the `equals` method, as follows:

```
if (name1.equals("Mark"))
```

To determine whether two strings are not equal, simply apply the `!` operator to the `equals` method's return value. Here is an example:

```
if (!name1.equals("Mark"))
```

The `boolean` expression in this `if` statement performs a not-equal-to operation. It determines whether the object referenced by `name1` is not equal to “Mark”.

The `String` class also provides the `compareTo` method, which can be used to determine whether one string is greater than, equal to, or less than another string. The general form of the method is as follows:

```
StringReference.compareTo(OtherString)
```

StringReference is a variable that references a `String` object, and *OtherString* is either another variable that references a `String` object or a string literal. The method returns an integer value that can be used in the following manner:

- If the method's return value is negative, the string referenced by *StringReference* (the calling object) is less than the *OtherString* argument.
- If the method's return value is 0, the two strings are equal.
- If the method's return value is positive, the string referenced by *StringReference* (the calling object) is greater than the *OtherString* argument.

For example, assume that `name1` and `name2` are variables that reference `String` objects. The following `if` statement uses the `compareTo` method to compare the strings:

```
if (name1.compareTo(name2) == 0)  
    System.out.println("The names are the same.");
```

Also, the following expression compares the string referenced by `name1` to the string literal “Joe”:

```
if (name1.compareTo("Joe") == 0)  
    System.out.println("The names are the same.");
```

The program in [Code Listing 3-10](#) more fully demonstrates the `compareTo` method.

Code Listing 3-10 (StringCompareTo.java)

```
1  /**
2   * This program compares two String objects using
3   * the compareTo method.
4   */
5
6  public class StringCompareTo
7  {
8      public static void main(String[] args)
9      {
10         String name1 = "Mary",
11             name2 = "Mark";
12
13         // Compare "Mary" and "Mark"
14
15         if (name1.compareTo(name2) < 0)
16         {
17             System.out.println(name1 + " is less than " + name2);
18         }
19         else if (name1.compareTo(name2) == 0)
20         {
21             System.out.println(name1 + " is equal to " + name2);
22         }
23         else if (name1.compareTo(name2) > 0)
24         {
25             System.out.println(name1 + " is greater than " + name2);
26         }
27     }
28 }
```

Program Output

```
Mary is greater than Mark
```

Let's take a closer look at this program. When you use the `compareTo` method to compare two strings, the strings are compared character by character. This is often called a *lexicographical comparison*. The program uses the `compareTo` method to compare the strings “Mary” and “Mark”, beginning with the first, or leftmost, characters. This is illustrated in [Figure 3-15](#).




Figure 3-15 String comparison of “Mary” and “Mark”

[Figure 3-15 Full Alternative Text](#)

Here is how the comparison takes place:

1. The “M” in “Mary” is compared with the “M” in “Mark.” Because these are the same, the next characters are compared.
2. The “a” in “Mary” is compared with the “a” in “Mark.” Because these are the same, the next characters are compared.
3. The “r” in “Mary” is compared with the “r” in “Mark.” Because these are the same, the next characters are compared.
4. The “y” in “Mary” is compared with the “k” in “Mark.” Because these are not the same, the two strings are not equal. The character “y” is greater than “k”, so it is determined that “Mary” is greater than “Mark.”

If one of the strings in a comparison is shorter than the other, Java can only compare the corresponding characters. If the corresponding characters are identical, then the shorter string is considered less than the longer string. For example, suppose the strings “High” and “Hi” were being compared. The string “Hi” would be considered less than “High” because it is shorter.

Ignoring Case in String Comparisons

The `equals` and `compareTo` methods perform case-sensitive comparisons, which means that uppercase letters are not considered the same as their lowercase counterparts. In other words, “A” is not the same as “a”. This can obviously lead to problems when you want to perform case-insensitive comparisons.

The `String` class provides the `equalsIgnoreCase` and `compareToIgnoreCase` methods. These methods work like the `equals` and `compareTo` methods, except the case of the characters in the strings is ignored. For example, the program in [Code Listing 3-11](#) asks the user to enter the secret word, which is similar to a password. The secret word is "PROSPERO", and the program performs a case-insensitive string comparison to determine whether the user has entered it.

Code Listing 3-11 (SecretWord.java)

```
1 import java.util.Scanner;
2
3 /**
4  * This program demonstrates a case insensitive string comparison.
5  */
6
7 public class SecretWord
8 {
9     public static void main(String[] args)
10    {
11        String input;        // To hold the user's input
12
13        // Create a Scanner object for keyboard input.
14        Scanner keyboard = new Scanner(System.in);
15
16        // Prompt the user to enter the secret word.
17        System.out.print("Enter the secret word: ");
18        input = keyboard.nextLine();
19
20        // Determine whether the user entered the secret word.
21        if (input.equalsIgnoreCase("PROSPERO"))
22        {
23            System.out.println("Congratulations! You know the " +
24                               "secret word!");
25        }
26        else
27        {
28            System.out.println("Sorry, that is NOT the " +
29                               "secret word!");
30        }
31    }
32 }
```

Program Output with Example Input Shown in Bold

```
Enter the secret word: Ferdinand [Enter]
Sorry, that is NOT the secret word!
```

Program Output with Example Input Shown in Bold

Enter the secret word: **Prospero** [Enter]
Congratulations! You know the secret word!

The `compareToIgnoreCase` method works exactly like the `compareTo` method, except the case of the characters in the strings being compared is ignored.



MyProgrammingLab™ www.myprogramminglab.com

1. 3.20 Assume the variable `name` references a `String` object. Write an `if` statement that displays “Do I know you?” if the `String` object contains “Timothy”.
2. 3.21 Assume the variables `name1` and `name2` reference two different `String` objects, containing different strings. Write code that displays the strings referenced by these variables in alphabetical order.
3. 3.22 Modify the statement you wrote in response to Checkpoint 3.20 so it performs a case-insensitive comparison.

3.7 More about Variable Declaration and Scope

CONCEPT: The scope of a variable is limited to the block in which it is declared.

Recall from [Chapter 2](#) that a local variable is a variable that is declared inside a method. Java allows you to create local variables just about anywhere in a method. For example, look at the program in [Code Listing 3-12](#). The `main` method declares two `String` reference variables: `firstName` and `lastName`. Notice that the declarations of these variables appear near the code that first uses the variables.

Code Listing 3-12 (VariableScope.java)

```
1 import java.util.Scanner;
2
3 /**
4  * This program demonstrates how variables may be declared
5  * in various locations throughout a program.
6  */
7
8 public class VariableScope
9 {
10     public static void main(String[] args)
11     {
12         // Create a Scanner object for keyboard input.
13         Scanner keyboard = new Scanner(System.in);
14
15         // Get the user's first name.
16         System.out.print("Enter your first name: ");
17         String firstName;
18         firstName = keyboard.nextLine();
19
20         // Get the user's last name.
21         System.out.print("Enter your last name: ");
22         String lastName;
23         lastName = keyboard.nextLine();
24
25         // Display a message.
26         System.out.println("Hello " + firstName + " " + lastName);
27     }
28 }
```

Program Output with Example Input Shown in Bold

Enter your first name: **Amy [Enter]**

```
Enter your last name: Jones [Enter]  
Hello Amy Jones
```

Although it is a common practice to declare all of a method's local variables at the beginning of the method, it is possible to declare them at later points. Sometimes programmers declare certain variables near the part of the program where they are used in order to make their purpose more evident.

Recall from [Chapter 2](#) that a variable's scope is the part of the program where the variable's name may be used. A local variable's scope always starts at the variable's declaration, and ends at the closing brace of the block of code in which it is declared. In [Code Listing 3-12](#), the `firstName` variable is visible only to the code in lines 18 through 26. The `lastName` variable is visible only to the code in lines 23 through 26.



NOTE:

When a program is running and it enters the section of code that constitutes a variable's scope, it is said that the variable “comes into scope.” This simply means the variable is now visible and the program may reference it. Likewise, when a variable “leaves scope” it may not be used.

3.8 The Conditional Operator (Optional)

CONCEPT: You can use the conditional operator to create short expressions that work like `if-else` statements.

The *conditional operator* is powerful and unique. Because it takes three operands, it is considered a ternary operator. The conditional operator provides a shorthand method of expressing a simple `if-else` statement. The operator consists of the question mark (?) and the colon (:). You use the operator to write a conditional expression, in the following format:

```
BooleanExpression ? Value1 : Value2;
```

The *BooleanExpression* is like the `boolean` expression in the parentheses of an `if` statement. If the *BooleanExpression* is true, then the value of the conditional expression is *Value1*. Otherwise, the value of the conditional expression is *Value2*. Here is an example of a statement using the conditional operator:

```
y = x < 0 ? 10 : 20;
```

This preceding statement performs the same operation as the following `if-else` statement:

```
if (x < 0)
    y = 10;
else
    y = 20;
```

The conditional operator gives you the ability to pack decision-making power into a concise line of code. With a little imagination it can be applied to many other programming problems. For instance, consider the following statement:

```
System.out.println("Your grade is: " +
    (score < 60 ? "Fail." : "Pass."));
```

Converted to an `if-else` statement, it would be written as follows:

```
if (score < 60)
    System.out.println("Your grade is: Fail.");
else
    System.out.println("Your grade is: Pass.");
```

NOTE:

The parentheses are placed around the conditional expression because the `+` operator has higher precedence than the `?:` operator. Without the parentheses, the `+` operator would concatenate the value

in score with the string "Your grade is: ".

For a complete example using the conditional operator, see the program named *ConsultantCharges.java* in this chapter's source code folder, available for download from the book's companion Web site (www.pearsonhighered.com/gaddis).

Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 3.23 Rewrite the following `if-else` statements as statements that use the conditional operator.

```
1. if (x > y)
    z = 1;
else
    z = 20;
```

```
2. if (temp > 45)
    population = base * 10;
else
    population = base * 2;
```

```
3. if (hours > 40)
    wages *= 1.5;
else
    wages *= 1;
```

```
4. if (result >= 0)
    System.out.println("The result is positive.");
else
    System.out.println("The result is negative.");
```

3.9 The switch Statement

CONCEPT: The switch statement lets the value of a variable or expression determine where the program will branch to.

The *switch statement* is a *multiple alternative decision structure*. It allows you to test the value of a variable or an expression and then use that value to determine which statement or set of statements to execute. [Figure 3-16](#) shows an example of how a multiple alternative decision structure looks in a flowchart.




Figure 3-16 A multiple alternative decision structure

[Figure 3-16 Full Alternative Text](#)

In the flowchart, the diamond symbol shows `month`, which is the name of a variable. If the `month` variable contains the value 1, the program displays *January*. If the `month` variable contains the value 2, the program displays *February*. If the `month` variable contains the value 3, the program displays *March*. If the `month` variable contains none of these values, the action that is labeled **Default** is executed. In this case, the program displays *Error: Invalid month*.

Here is the general format of a `switch` statement in Java:

The testExpression is a variable or expression.

```

switch (testExpression)
{
    case value_1:
        statement;
        statement;
        etc.
        break;
    case value_2:
        statement;
        statement;
        etc.
        break;
    case value_N:
        statement;
        statement;
        etc.
        break;
    default:
        statement;
        statement;
        etc.
        break;
}

```

These statements are executed if the testExpression is equal to value_1.

These statements are executed if the testExpression is equal to value_2.

Insert as many case sections as necessary.

These statements are executed if the testExpression is equal to value_N.

These statements are executed if the testExpression is not equal to any of the case values.

[3.9-12 Full Alternative Text](#)

The first line of the statement starts with the word `switch`, followed by a *testExpression*, which is enclosed in parentheses. The *testExpression* is a variable or an expression that gives a `char`, `byte`, `short`, `int`, or `string` value.

Beginning at the next line is a block of code enclosed in curly braces. Inside this block of code is one or more `case` sections. A `case` section begins with the word `case`, followed by a value, followed by a colon. Each `case` section contains one or more statements, followed by a `break` statement. After all of the `case` sections, an optional `default` section appears.

When the `switch` statement executes, it compares the value of the *testExpression* with the values that follow each of the `case` statements (from top to bottom). When it finds a `case` value that matches the *testExpression*'s value, the program branches to the `case` statement. The statements that follow the `case` statement are executed until a `break` statement is encountered. At that point, the program jumps out of the `switch` statement. If the *testExpression* does not match any of the `case` values, the program branches to the `default` statement and executes the statements that immediately follow it.



NOTE:

Each of the `case` values must be unique.

For example, the following code performs the same operation as the flowchart in [Figure 3-16](#). Assume `month` is an `int` variable.

```
switch (month)
{
    case 1:
        System.out.println("January");
        break;
    case 2:
        System.out.println("February");
        break;
    case 3:
        System.out.println("March");
        break;
    default:
        System.out.println("Error: Invalid month");
        break;
}
```

In this example the *testExpression* is the `month` variable. The `month` variable will be evaluated and one of the following actions will take place:

- If the value in the `month` variable is 1, the program will branch to the `case 1:` section and execute the `System.out.println("January")` statement that immediately follows it. The `break` statement then causes the program to exit the `switch` statement.
- If the value in the `month` variable is 2, the program will branch to the `case 2:` section and execute the `System.out.println("February")` statement that immediately follows it. The `break` statement then causes the program to exit the `switch` statement.
- If the value in the `month` variable is 3, the program will branch to the `case 3:` section and execute the `System.out.println("March")` statement that immediately follows it. The `break` statement then causes the program to exit the `switch` statement.
- If the value in the `month` variable is not 1, 2, or 3, the program will branch to the `default:` section and execute the `System.out.println("Error: Invalid month")` statement that immediately follows it.

The `switch` statement can be used as an alternative to an `if-else-if` statement that compares the same variable or expression to several different values. For example, the previously shown `switch` statement works like this `if-else-if` statement:

```
if (month == 1)
{
```

```

        System.out.println("January");
    }
    else if (month == 2)
    {
        System.out.println("February");
    }
    else if (month == 3)
    {
        System.out.println("March");
    }
    else
    {
        System.out.println("Error: Invalid month");
    }
}

```

NOTE:

The `default` section is optional. If you leave it out, however, the program will have nowhere to branch to if the *testExpression* doesn't match any of the case values.

The program in [Code Listing 3-13](#) shows how a simple `switch` statement works.

Code Listing 3-13 (SwitchDemo.java)

```

1 import java.util.Scanner;
2
3 /**
4  * This program demonstrates the switch statement.
5  */
6
7 public class SwitchDemo
8 {
9     public static void main(String[] args)
10    {
11        int number;    // A number entered by the user
12
13        // Create a Scanner object for keyboard input.
14        Scanner keyboard = new Scanner(System.in);
15
16        // Get one of the numbers 1, 2, or 3 from the user.
17        System.out.print("Enter 1, 2, or 3: ");
18        number = keyboard.nextInt();
19
20        // Determine the number entered.
21        switch (number)
22        {
23            case 1:
24                System.out.println("You entered 1.");
25                break;
26            case 2:
27                System.out.println("You entered 2.");

```

```

28         break;
29     case 3:
30         System.out.println("You entered 3.");
31         break;
32     default:
33         System.out.println("That's not 1, 2, or 3!");
34     }
35 }
36 }

```

Program Output with Example Input Shown in Bold

Enter 1, 2, or 3: **2** [Enter]
 You entered 2.

Program Output with Example Input Shown in Bold

Enter 1, 2, or 3: **5** [Enter]
 That's not 1, 2, or 3!

Notice the `break` statements that are in the `case 1`, `case 2`, and `case 3` sections.

```

switch (number)
{
    case 1:
        System.out.println("You entered 1.");
        break;
    case 2:
        System.out.println("You entered 2.");
        break;
    case 3:
        System.out.println("You entered 3.");
        break;
    default:
        System.out.println("That's not 1, 2, or 3!");
}

```

The `case` statements show the program where to start executing in the block and the `break` statements show the program where to stop. Without the `break` statements, the program would execute all of the lines from the matching `case` statement to the end of the block.

 **NOTE:**

The default section (or the last case section if there is no default) does not need a break statement. Some programmers prefer to put one there anyway for consistency.

The program in [Code Listing 3-14](#) is a modification of [Code Listing 3-13](#), without the break statements.

Code Listing 3-14 (NoBreaks.java)

```
1 import java.util.Scanner;
2
3 /**
4  * This program demonstrates the switch statement.
5  */
6
7 public class NoBreaks
8 {
9     public static void main(String[] args)
10    {
11        int number;    // A number entered by the user
12
13        // Create a Scanner object for keyboard input.
14        Scanner keyboard = new Scanner(System.in);
15
16        // Get one of the numbers 1, 2, or 3 from the user.
17        System.out.print("Enter 1, 2, or 3: ");
18        number = keyboard.nextInt();
19
20        // Determine the number entered.
21        switch (number)
22        {
23            case 1:
24                System.out.println("You entered 1.");
25            case 2:
26                System.out.println("You entered 2.");
27            case 3:
28                System.out.println("You entered 3.");
29            default:
30                System.out.println("That's not 1, 2, or 3!");
31        }
32    }
33 }
```

Program Output with Example Input Shown in Bold

```
Enter 1, 2, or 3:  1 [Enter]
You entered 1.
You entered 2.
You entered 3.
```

That's not 1, 2, or 3!

Program Output with Example Input Shown in Bold

```
Enter 1, 2, or 3:  3 [Enter]
```

```
You entered 3.
```

```
That's not 1, 2, or 3!
```

Without the `break` statement, the program “falls through” all of the statements below the one with the matching case expression. Sometimes this is what you want. For instance, the program in [Code Listing 3-15](#) asks the user to select a grade of pet food. The available choices are A, B, and C. The `switch` statement will recognize either uppercase or lowercase letters.

Code Listing 3-15 (PetFood.java)

```
1 import java.util.Scanner;
2
3 /**
4  * This program demonstrates a switch statement.
5  */
6
7 public class PetFood
8 {
9     public static void main(String[] args)
10    {
11        String input;    // To hold the user's input
12        char foodGrade;  // Grade of pet food
13
14        // Create a Scanner object for keyboard input.
15        Scanner keyboard = new Scanner(System.in);
16
17        // Prompt the user for a grade of pet food.
18        System.out.println("Our pet food is available in " +
19                            "three grades:");
20        System.out.print("A, B, and C. Which do you want " +
21                          "pricing for? ");
22        input = keyboard.nextLine();
23        foodGrade = input.charAt(0);
24
25        // Display pricing for the selected grade.
26        switch(foodGrade)
27        {
28            case 'a':
29            case 'A':
30                System.out.println("30 cents per lb.");
31                break;
32            case 'b':
33            case 'B':
```

```

34         System.out.println("20 cents per lb.");
35         break;
36     case 'c':
37     case 'C':
38         System.out.println("15 cents per lb.");
39         break;
40     default:
41         System.out.println("Invalid choice.");
42     }
43 }
44 }

```

Program Output with Example Input Shown in Bold

Our pet food is available in three grades:
A, B, and C. Which do you want pricing for? **b** [Enter]
20 cents per lb.

Program Output with Example Input Shown in Bold

Our pet food is available in three grades:
A, B, and C. Which do you want pricing for? **B** [Enter]
20 cents per lb.

When the user enters ‘a’ the corresponding `case` has no statements associated with it, so the program falls through to the next `case`, which corresponds with ‘A’.

```

case 'a':
case 'A':
    System.out.println("30 cents per lb.");
    break;

```

The same technique is used for ‘b’ and ‘c’.

If you are using a version of Java prior to Java 7, a `switch` statement’s *testExpression* can be a `char`, `byte`, `short`, or `int` value. Beginning with Java 7, however, the *testExpression* can also be a string. The program in [Code Listing 3-16](#) demonstrates.

Code Listing 3-16 (Seasons.java)

```

1 import java.util.Scanner;
2

```

```

3  /**
4      This program translates the English names of
5      the seasons into Spanish.
6  */
7
8  public class Seasons
9  {
10     public static void main(String[] args)
11     {
12         String input;
13
14         // Create a Scanner object for keyboard input.
15         Scanner keyboard = new Scanner(System.in);
16
17         // Get a day from the user.
18         System.out.print("Enter the name of a season: ");
19         input = keyboard.nextLine();
20
21         // Translate the season to Spanish.
22         switch (input)
23         {
24             case "spring":
25                 System.out.println("la primavera");
26                 break;
27             case "summer":
28                 System.out.println("el verano");
29                 break;
30             case "autumn":
31             case "fall":
32                 System.out.println("el otono");
33                 break;
34             case "winter":
35                 System.out.println("el invierno");
36                 break;
37             default:
38                 System.out.println("Please enter one of these words:\n"+
39                                     "spring, summer, autumn, fall, or winter.");
40         }
41     }
42 }

```

Program Output with Example Input Shown in Bold

```

Enter the name of a season: summer [Enter]
el verano

```

Program Output with Example Input Shown in Bold

Enter the name of a season: **fall** [Enter]
el otono



MyProgrammingLab™ www.myprogramminglab.com

1. 3.24 Complete the following program skeleton by writing a `switch` statement that displays “one” if the user has entered 1, “two” if the user has entered 2, and “three” if the user has entered 3. If a number other than 1, 2, or 3 is entered, the program should display an error message.

```
import java.util.Scanner;
public class CheckPoint
{
    public static void main(String[] args)
    {
        int userNum;
        Scanner keyboard = new Scanner(System.in);
        System.out.print("Enter one of the numbers " +
                        "1, 2, or 3: ");
        userNum = keyboard.nextInt();
        //
        // Write the switch statement here.
        //
    }
}
```

2. 3.25 Rewrite the following `if-else-if` statement as a `switch` statement.

```
if (selection == 'A')
    System.out.println("You selected A.");
else if (selection == 'B')
    System.out.println("You selected B.");
else if (selection == 'C')
    System.out.println("You selected C.");
else if (selection == 'D')
    System.out.println("You selected D.");
else
    System.out.println("Not good with letters, eh?");
```

3. 3.26 Explain why you cannot convert the following `if-else-if` statement into a `switch` statement.

```
if (temp == 100)
    x = 0;
else if (population > 1000)
    x = 1;
else if (rate < .1)
    x = -1;
```


4. 3.27 What is wrong with the following `switch` statement?

```
// This code has errors!!!
switch (temp)
{
    case temp < 0 :
        System.out.println("Temp is negative.");
        break;
    case temp = 0:
        System.out.println("Temp is zero.");
        break;
    case temp > 0 :
        System.out.println("Temp is positive.");
        break;
}
```

5. 3.28 What will the following code display?

```
int funny = 7, serious = 15;
funny = serious * 2;
switch (funny)
{ case 0 :
    System.out.println("That is funny.");
    break;
  case 30:
    System.out.println("That is serious.");
    break;
  case 32:
    System.out.println("That is seriously funny.");
    break;
  default:
    System.out.println(funny);
}
```

3.10 Displaying Formatted Output with `System.out.printf` and `String.format`

CONCEPT: The `System.out.printf` method allows you to format output in a variety of ways. The `String.format` method allows you to format a string, without displaying it. The string can be displayed at a later time.

When you display numbers with the `System.out.println` or `System.out.print` method, you have little control over the way the numbers appear. For example, a value of the `double` data type can be displayed with as many as 15 decimal places, as demonstrated by the following code:

```
double number = 10.0 / 6.0;
System.out.println(number);
```

This code will display:

```
1.6666666666666667
```

Quite often, you want to format numbers so they are displayed in a particular way. For example, you might want to round a floating-point number to a specific number of decimal places, or insert comma separators to make a number easier to read. Fortunately, Java gives us a way to do just that, and more, with the `System.out.printf` method. The method's general format is as follows:

```
System.out.printf(FormatString, ArgumentList)
```

In the general format, *FormatString* is a string that contains text, special formatting specifiers, or both. *ArgumentList* is a list of zero or more additional arguments, which will be formatted according to the format specifiers listed in the format string.

The simplest way you can use the `System.out.printf` method is with only a format string, and no additional arguments. Here is an example:

```
System.out.printf("I love Java programming.\n");
```

The format string in this example is `"I love Java programming.\n"`. This method call does not perform any special formatting, however. It simply prints the string `"I love Java programming.\n"`. Using the method in this fashion is exactly like using the `System.out.print` method.

In most cases you will call the `System.out.printf` method in the following manner:

- The format string will contain one or more format specifiers. A *format specifier* is a placeholder for a value that will be inserted into the string when it is displayed.

- After the format string, one or more additional arguments will appear. Each of the additional arguments will correspond to a format specifier that appears inside the format string.

The following code shows an example:

```
double sales = 12345.67;  
System.out.printf("Our sales are %f for the day.\n", sales);
```


Notice the following characteristics of the `System.out.printf` method call:

- Inside the format string, the `%f` is a format specifier. The letter `f` indicates that a floating-point value will be inserted into the string when it is displayed.
- Following the format string, the `sales` variable is passed as an argument. This argument corresponds to the `%f` format specifier that appears inside the format string.

When the `System.out.printf` method executes, the `%f` will not be displayed on the screen. In its place, the value of the `sales` argument will be displayed. Here is the output of the code:

```
Our sales are 12345.670000 for the day.
```

The diagram in [Figure 3-17](#) shows how the `sales` variable corresponds to the `%f` format specifier.



The diagram shows the code `System.out.printf("Our sales is %f for the day.\n", sales);`. A curved arrow originates from the `sales` variable at the end of the argument list and points to the `%f` format specifier within the format string.

Figure 3-17 The value of the `sales` variable is displayed in the place of the `%f` format specifier

Here is another example:

```
double temp1 = 72.5, temp2 = 83.7;  
System.out.printf("The temperatures are %f and %f degrees.\n", temp1, temp2);
```

First, notice that this example uses two `%f` format specifiers in the format string. Also notice that two additional arguments appear after the format string. The value of the first argument, `temp1`, will be printed in place of the first `%f`, and the value of the second argument, `temp2`, will be printed in place of the second `%f`. The code will produce the following output:

```
The temperatures are 72.500000 and 83.700000 degrees.
```

There is a one-to-one correspondence between the format specifiers and the arguments that appear after the format string. The diagram in [Figure 3-18](#) shows how the first format specifier corresponds

to the first argument after the format string (the `temp1` variable), and the second format specifier corresponds to the second argument after the format string (the `temp2` variable).




Figure 3-18 The format specifiers and their corresponding arguments

[Figure 3-18 Full Alternative Text](#)

The following code shows another example:

```
double value1 = 3.0;
double value2 = 6.0;
double value3 = 9.0;
System.out.printf("%f %f %f\n", value1, value2, value3);
```

In the `System.out.printf` method call, there are three format specifiers and three additional arguments after the format string. This code will produce the following output:

```
3.000000 6.000000 9.000000
```

The diagram in [Figure 3-19](#) shows how the format specifiers correspond to the arguments that appear after the format string.




Figure 3-19 The format specifiers and

their corresponding arguments

[Figure 3-19 Full Alternative Text](#)

The previous examples demonstrated how to format floating-point numbers with the `%f` format specifier. The letter `f` in the format specifier is a *conversion character* that indicates the data type of the argument that is being formatted. You use the `f` conversion character with any argument that is a float or a double.

If you want to format an integer value, you must use the `%d` format specifier. The `d` conversion character stands for decimal integer, and it can be used with arguments of the `int`, `short`, and `long` data types. Here is an example that displays an `int`:

```
int hours = 40;
System.out.printf("I worked %d hours this week.\n", hours);
```

In this example, the `%d` format specifier corresponds with the `hours` argument. This code will display the following:

```
I worked 40 hours this week.
```

Here is an example that displays two `int` values:

```
int dogs = 2;
int cats = 4;
System.out.printf("We have %d dogs and %d cats.\n", dogs, cats);
```

This code will display the following:

```
We have 2 dogs and 4 cats.
```

Keep in mind that `%f` must be used with floating-point values, and `%d` must be used with integer values. Otherwise, an error will occur at runtime.

Format Specifier Syntax

In the previous examples you saw how format specifiers correspond to the arguments that appear after the format string. Now you can learn how to use format specifiers to actually format the values that they correspond to. When displaying numbers, the general syntax for writing a format specifier is:

```
%[flags][width][.precision]conversion
```

The items that appear inside brackets are optional. Here is a summary of each item:

- `%`—All format specifiers begin with a `%` character.

- *flags*—After the `%` character, one or more optional flags may appear. Flags cause the value to be formatted in a variety of ways.
- *width*—After any flags, you can optionally specify the minimum field width for the value.
- *precision*—If the value is a floating-point number, after the minimum field width, you can optionally specify the precision. This is the number of decimal places that the number should be rounded to.
- *conversion*—All format specifiers must end with a conversion character, such as `f` for floating-point, or `d` for decimal integer.

Let's take a closer look at each of the optional items, beginning with precision.

Precision

You probably noticed in the previous examples that the `%f` format specifier causes floating-point values to be displayed with six decimal places. You can change the number of decimal points that are displayed, as shown in the following example:

```
double temp = 78.42819;
System.out.printf("The temperature is %.2f degrees.\n", temp);
```

Notice that this example doesn't use the regular `%f` format specifier, but uses `%.2f` instead. The `.2` that appears between the `%` and the `f` specifies the *precision* of the displayed value. It will cause the value of the `temp` variable to be rounded to two decimal places. This code will produce the following output:

```
The temperature is 78.43 degrees.
```

The following example displays the same value, rounded to one decimal place:

```
double temp = 78.42819;
System.out.printf("The temperature is %.1f degrees.\n", temp);
```

This code will produce the following output:

```
The temperature is 78.4 degrees.
```

The following code shows another example:

```
double value1 = 123.45678;
double value2 = 123.45678;
double value3 = 123.45678;
System.out.printf("%.1f %.2f %.3f\n", value1, value2, value3);
```

In this example, `value1` is rounded to one decimal place, `value2` is rounded to two decimal places,

and `value3` is rounded to three decimal places. This code will produce the following output:

```
123.5 123.46 123.457
```

Keep in mind that you can specify precision only with floating-point values. If you specify a precision with the `%d` format specifier, an error will occur at runtime.

Specifying a Minimum Field Width

A format specifier can also include a *minimum field width*, which is the minimum number of spaces that should be used to display the value. The following example prints a floating-point number in a field that is 20 spaces wide:

```
double number = 12345.6789;  
System.out.printf("The number is:%20f\n", number);
```

Notice that the number 20 appears in the format specifier, between the `%` and the `f`. This code will produce the following output:

```
The number is:      12345.678900
```

In this example, the 20 that appears inside the `%f` format specifier indicates that the number should be displayed in a field that is a minimum of 20 spaces wide. This is illustrated in [Figure 3-20](#).

The number is: 12345.678900

The number is displayed in a field that is 20 spaces wide.

Figure 3-20 The number is displayed in a field that is 20 spaces wide

In this case, the number that is displayed is shorter than the field in which it is displayed. The number `12345.678900` uses only 12 spaces on the screen, but it is displayed in a field that is 20 spaces wide. When this is the case, the number will be right-justified in the field. If a value is too large to fit in the specified field width, the field is automatically enlarged to accommodate it. The following example prints a floating-point number in a field that is only one space wide:

```
double number = 12345.6789;  
System.out.printf("The number is:%1f\n", number);
```

The value of the number variable requires more than one space, however, so the field width is expanded to accommodate the entire number. This code will produce the following output:

```
The number is:12345.678900
```

You can specify a minimum field width for integers, as well as for floating-point values. The following example displays an integer with a minimum field width of six characters:

```
int number = 200;
System.out.printf("The number is:%6d", number);
```

This code will display the following:

```
The number is:   200
```

Combining Minimum Field Width and Precision in the Same Format Specifier

When specifying the minimum field width and the precision of a floating-point number in the same format specifier, remember that the field width must appear first, followed by the precision. For example, the following code displays a number in a field of 12 spaces, rounded to two decimal places:

```
double number = 12345.6789;
System.out.printf("The number is:%12.2f\n", number);
```

This code will produce the following output:

```
The number is:   12345.68
```

Field widths can help when you need to print numbers aligned in columns. For example, look at [Code Listing 3-17](#). Each of the variables is displayed in a field that is eight spaces wide, and rounded to two decimal places. The numbers appear aligned in a column.

Code Listing 3-17 (Columns.java)

```
1  /**
2      This program displays a variety of
3      floating-point numbers in a column
4      with their decimal points aligned.
5  */
6
7  public class Columns
8  {
9      public static void main(String[] args)
10     {
11         // Declare a variety of double variables.
12         double num1 = 127.899;
13         double num2 = 3465.148;
14         double num3 = 3.776;
```



```

15         double num4 = 264.821;
16         double num5 = 88.081;
17         double num6 = 1799.999;
18
19         // Display each variable in a field of
20         // 8 spaces with 2 decimal places.
21         System.out.printf("%8.2f\n", num1);
22         System.out.printf("%8.2f\n", num2);
23         System.out.printf("%8.2f\n", num3);
24         System.out.printf("%8.2f\n", num4);
25         System.out.printf("%8.2f\n", num5);
26         System.out.printf("%8.2f\n", num6);
27     }
28 }

```

Program Output

```

127.90
3465.15
  3.78
264.82
 88.08
1800.00

```

Flags

There are several optional flags that you can insert into a format specifier to cause a value to be formatted in a particular way. In this book, we will use flags for the following purposes:

- To display numbers with comma separators
- To pad numbers with leading zeros
- To left-justify numbers

If you use a flag in a format specifier, you must write the flag before the field width and the precision.

Comma Separators

Large numbers are easier to read if they are displayed with comma separators. You can format a number with comma separators by inserting a comma (,) flag into the format specifier. Here is an example:

```

double amount = 1234567.89;
System.out.printf("%,f\n", amount);

```

This code will produce the following output:

```
1,234,567.890000
```

Quite often, you will want to format a number with comma separators, and round the number to a specific number of decimal places. You can accomplish this by inserting a comma, followed by the precision value, into the `%f` format specifier, as shown in the

following example:

```
double sales = 28756.89;
System.out.printf("Sales for the month are %, .2f\n", sales);
```

This code will produce the following output:

```
Sales for the month are 28,756.89
```

[Code Listing 3-18](#) demonstrates how the comma separator and a precision of two decimal places can be used to format a number as a currency amount.

Code Listing 3-18 (CurrencyFormat.java)

```
1  /**
2   * This program demonstrates how to use the System.out.printf
3   * method to format a number as currency.
4   */
5
6  public class CurrencyFormat
7  {
8      public static void main(String[] args)
9      {
10         double monthlyPay = 5000.0;
11         double annualPay = monthlyPay * 12;
12         System.out.printf("Your annual pay is $%,.2f\n", annualPay);
13     }
14 }
```

Program Output

```
Your annual pay is $60,000.00
```

The following example displays a floating-point number with comma separators, in a field of 15 spaces, rounded to two decimal places:

```
double amount = 1234567.8901;
System.out.printf("% ,15.2f\n", amount);
```

This code will produce the following output:

1,234,567.89

The following example displays an `int` with a minimum field width of six characters:

```
int number = 200;
System.out.printf("The number is:%6d", number);
```

This code will display the following:

The number is: 200

The following example displays an `int` with comma separators, with a minimum field width of 10 characters:

```
int number = 20000;
System.out.printf("The number is:%,10d", number);
```

This code will display the following:

The number is: 20,000

Padding Numbers with Leading Zeros

Sometimes, when a number is shorter than the field in which it is displayed, you want to pad the number with leading zeros. If you insert a `0` flag into a format specifier, the resulting number will be padded with leading zeros, if it is shorter than the field width. The following code shows an example:

```
double number = 123.4;
System.out.printf("The number is:%08.1f\n", number);
```

This code will produce the following output:

The number is:000123.4

The diagram in [Figure 3-21](#) shows the purpose of each part of the format specifier in the previous example.




Figure 3-21 Format specifier that pads

with leading zeros

[Figure 3-21 Full Alternative Text](#)

The following example displays an `int` padded with leading zeros, with a minimum field width of seven characters:

```
int number = 1234;
System.out.printf("The number is:%07d", number);
```

This code will display the following:

```
The number is:0001234
```

The program in [Code Listing 3-19](#) shows another example. This program displays a variety of floating-point numbers with leading zeros, in a field of nine spaces, rounded to two decimal places.

Code Listing 3-19 (LeadingZeros.java)

```
1  /**
2   * This program displays numbers padded with leading zeros.
3   */
4
5  public class LeadingZeros
6  {
7      public static void main(String[] args)
8      {
9          // Declare a variety of double variables.
10         double number1 = 1.234;
11         double number2 = 12.345;
12         double number3 = 123.456;
13
14         // Display each variable with leading
15         // zeros, in a field of 9 spaces, rounded
16         // to 2 decimal places.
17         System.out.printf("%09.2f\n", number1);
18         System.out.printf("%09.2f\n", number2);
19         System.out.printf("%09.2f\n", number3);
20     }
21 }
```

Program Output

```
000001.23
000012.35
000123.46
```

Left-Justifying Numbers

By default, when a number is shorter than the field in which it is displayed, the number is right-justified within that field. If you want a number to be left-justified within its field, you insert a minus sign (-) flag into the format specifier. [Code Listing 3-20](#) shows an example.

Code Listing 3-20 (LeftJustified.java)

```
1  /**
2   * This program displays a variety of
3   * numbers left-justified in columns.
4   */
5
6  public class LeftJustified
7  {
8      public static void main(String[] args)
9      {
10         // Declare a variety of int variables.
11         int num1 = 123;
12         int num2 = 12;
13         int num3 = 45678;
14         int num4 = 456;
15         int num5 = 1234567;
16         int num6 = 1234;
17
18         // Display each variable left-justified
19         // in a field of 8 spaces.
20         System.out.printf("%-8d%-8d\n", num1, num2);
21         System.out.printf("%-8d%-8d\n", num3, num4);
22         System.out.printf("%-8d%-8d\n", num5, num6);
23     }
24 }
```

Program Output

```
123      12
45678    456
1234567 1234
```

Formatting String Arguments

If you wish to print a string argument, use the `%s` format specifier. Here is an example:

```
String name = "Ringo";
System.out.printf("Your name is %s\n", name);
```

This code produces the following output:

```
Your name is Ringo
```

You can also use a field width when printing strings. For example, look at the following code:

```
String name1 = "George";
String name2 = "Franklin";
String name3 = "Jay";
String name4 = "Ozzy";
String name5 = "Carmine";
String name6 = "Dee";
System.out.printf("%10s%10s\n", name1, name2);
System.out.printf("%10s%10s\n", name3, name4);
System.out.printf("%10s%10s\n", name5, name6);
```

The `%10s` format specifier prints a string in a field that is ten spaces wide. This code displays the values of the variables in a table with three rows and two columns. Each column has a width of ten spaces. Here is the output of the code:

```
George Franklin
      Jay      Ozzy
Carmine      Dee
```

Notice that the strings are right-justified. You can use the minus flag (`-`) to left-justify a string within its field. The following code demonstrates:

```
String name1 = "George";
String name2 = "Franklin";
String name3 = "Jay";
String name4 = "Ozzy";
String name5 = "Carmine";
String name6 = "Dee";
System.out.printf("%-10s%-10s\n", name1, name2);
System.out.printf("%-10s%-10s\n", name3, name4);
System.out.printf("%-10s%-10s\n", name5, name6);
```

Here is the output of the code:

```
George  Franklin
Jay      Ozzy
Carmine  Dee
```

The following example shows how you can print arguments of different data types:

```
int hours = 40;
double pay = hours * 25;
String name = "Jay";
System.out.printf("Name: %s, Hours: %d, Pay: $%,.2f\n",
                  name, hours, pay);
```

In this example, we are displaying a `String`, an `int`, and a `double`. The code will produce the following output:



The format specifiers we have shown in this section are the basic ones. Java provides much more powerful format specifiers for more complex formatting needs. The API documentation gives an overview of them all.

The `String.format` Method

The `System.out.printf` method formats a string and displays it in the console window. Sometimes you need to format a string without displaying it in the console. For example, you might need to display formatted output in a graphical interface, such as a message dialog. When this is the case, you can use the `String.format` method.

The `String.format` method works exactly like the `System.out.printf` method, except that it does not display the formatted string on the screen. Instead, it returns a reference to the formatted string. You can assign the reference to a variable, and then use it later. Here is the method's general format:

```
String.format(FormatString, ArgumentList)
```

In the general format, *FormatString* is a string that contains text, special formatting specifiers, or both. *ArgumentList* is a list of zero or more additional arguments, which will be formatted according to the format specifiers listed in the format string. The syntax for writing the format specifiers is the same as with the `System.out.printf` method. The method creates a string in memory that is formatted as specified, and returns a reference to that string. For example, look at the program in [Code Listing 3-21](#). The program's output is shown in [Figure 3-22](#).




Figure 3-22 Output of Code Listing 3-21

(Oracle Corporate Counsel)

Code Listing 3-21 (`CurrencyFormat2.java`)

```

1 import javax.swing.JOptionPane;
2
3 /**
4     This program demonstrates how to use the String.format
5     method to format a number as currency.
6 */
7
8 public class CurrencyFormat2
9 {
10     public static void main(String[] args)
11     {
12         double monthlyPay = 5000.0;
13         double annualPay = monthlyPay * 12;
14         String output = String.format("Your annual pay is $%,.2f", annualPay);
15
16         JOptionPane.showMessageDialog(null, output);
17     }
18 }

```

Let's take a closer look at the program. Line 12 declares a double variable named `monthlyPay`, initialized with the value 5000.0, and line 13 declares a double variable named `annualPay`, initialized with the result of the calculation `monthlyPay * 12`. Line 14 declares a String variable named `output`, and initializes it with the string that is returned from the `String.format` method. In line 16, the `output` variable is passed as an argument to the `JOptionPane.showMessageDialog` method.

The program in [Code Listing 3-21](#) can be simplified. We can combine the steps of calling the `String.format` method, and passing the value that it returns to the `JOptionPane.showMessageDialog` method. This allows us to eliminate the declaration of the `output` variable. [Code Listing 3-22](#) shows how this is done. The program's output is the same as shown in [Figure 3-22](#).

Code Listing 3-22 (CurrencyFormat3.java)

```

1 import javax.swing.JOptionPane;
2
3 /**
4     This program demonstrates how to use the String.format
5     method to format a number as currency.
6 */
7
8 public class CurrencyFormat3
9 {
10     public static void main(String[] args)
11     {
12         double monthlyPay = 5000.0;
13         double annualPay = monthlyPay * 12;
14
15         JOptionPane.showMessageDialog(null,
16             String.format("Your annual pay is $%,.2f", annualPay));
17     }
18 }

```




1. 3.29 Assume the following variable declaration exists in a program:

```
double number = 1234567.456;
```

Write a statement that uses `System.out.printf` to display the value of the `number` variable formatted as:

```
1,234,567.46
```

2. 3.30 Assume the following variable declaration exists in a program:

```
double number = 123.456;
```

Write a statement that uses `System.out.printf` to display the value of the `number` variable rounded to one decimal place, in a field that is 10 spaces wide. (Do not use comma separators.)

3. 3.31 Assume the following variable declaration exists in a program:

```
double number = 123.456;
```

Write a statement that uses `System.out.printf` to display the value of the `number` variable padded with leading zeros, in a field that is eight spaces wide, rounded to one decimal place. (Do not use comma separators.)

4. 3.32 Assume the following variable declaration exists in a program:

```
int number = 123456;
```

Write a statement that uses `System.out.printf` to display the value of the `number` variable in a field that is 10 spaces wide, with comma separators.

5. 3.33 Assume the following variable declaration exists in a program:

```
double number = 123456.789;
```

Write a statement that uses `System.out.printf` to display the value of the `number` variable left-justified, with comma separators, in a field that is 20 spaces wide, rounded to two decimal places.

6. 3.34 Assume the following declaration exists in a program:

```
String name = "James";
```

Write a statement that uses `System.out.printf` to display the value of `name` in a field that is 20 spaces wide.

3.11 Common Errors to Avoid

The following list describes several errors that are commonly committed when learning this chapter's topics.

- **Using `=` instead of `==` to compare primitive values.** Remember, `=` is the assignment operator and `==` tests for equality.
- **Using `==` instead of the `equals` method to compare `String` objects.** You cannot use the `==` operator to compare the contents of a `String` object with another string. Instead you must use the `equals` or `compareTo` method.
- **Forgetting to enclose an `if` statement's `boolean` expression in parentheses.** Java requires that the `boolean` expression being tested by an `if` statement is enclosed in a set of parentheses. An error will result if you omit the parentheses or use any other grouping characters.
- **Writing a semicolon at the end of an `if` clause.** When you write a semicolon at the end of an `if` clause, Java assumes that the conditionally executed statement is a null or empty statement.
- **Forgetting to enclose multiple conditionally executed statements in braces.** Normally the `if` statement conditionally executes only one statement. To conditionally execute more than one statement, you must enclose them in braces.
- **Omitting the trailing `else` in an `if-else-if` statement.** This is not a syntax error, but can lead to logical errors. If you omit the trailing `else` from an `if-else-if` statement, no code will be executed if none of the statement's `boolean` expressions are `true`.
- **Not writing complete `boolean` expressions on both sides of a logical `&&` or `||` operator.** You must write a complete `boolean` expression on both sides of a logical `&&` or `||` operator. For example, the expression `x > 0 && < 10` is not valid because `< 10` is not a complete expression. The expression should be written as `x > 0 && x < 10`.
- **Trying to perform case-insensitive string comparisons with the `String` class's `equals` and `compareTo` methods.** To perform case-insensitive string comparisons, use the `String` class's `equalsIgnoreCase` and `compareToIgnoreCase` methods.
- **Using a `SwitchExpression` that is not an `int`, `short`, `byte`, `char`, or `String`.** The `switch` statement can only evaluate expressions that are of the `int`, `short`, `byte`, `char`, or `String` data types.
- **Using a `CaseExpression` that is not a literal or a `final` variable.** Because the compiler must determine the value of a `CaseExpression` at compile time, `CaseExpressions` must be either literal values or `final` variables.

- **Forgetting to write a colon at the end of a `case` statement.** A colon must appear after the *CaseExpression* in each `case` statement.
- **Forgetting to write a `break` statement in a `case` section.** This is not a syntax error, but it can lead to logical errors. The program does not branch out of a `switch` statement until it reaches a `break` statement or the end of the `switch` statement.
- **Forgetting to write a `default` section in a `switch` statement.** This is not a syntax error, but can lead to a logical error. If you omit the `default` section, no code will be executed if none of the *CaseExpressions* match the *SwitchExpression*.
- **Reversing the `?` and the `:` when using the conditional operator.** When using the conditional operator, the `?` character appears first in the conditional expression, then the `:` character.
- **When formatting a number with `System.out.printf` or `String.format`, writing the flags, field width, and precision in an incorrect order.**
- **When writing a format specifier for the `System.out.printf` or `String.format` methods, using the wrong type indicator (`%f` = floating-point, `%d` = integer, `%s` = string).**

Review Questions and Exercises

Multiple Choice and True/False

1. The `if` statement is an example of a .
 1. sequence structure
 2. decision structure
 3. pathway structure
 4. class structure
2. This type of expression has a value of either `true` or `false`.
 1. binary expression
 2. decision expression
 3. unconditional expression
 4. boolean expression
3. `>`, `<`, and `==` are .
 1. relational operators
 2. logical operators
 3. conditional operators
 4. ternary operators
4. `&&`, `||`, and `!` are .
 1. relational operators
 2. logical operators
 3. conditional operators
 4. ternary operators

5. This is an empty statement that does nothing.
 1. missing statement
 2. virtual statement
 3. `null` statement
 4. conditional statement
6. To create a block of statements, you enclose the statements in these.
 1. parentheses `()`
 2. square brackets `[]`
 3. angled brackets `<>`
 4. braces `{ }`
7. This is a `boolean` variable that signals when some condition exists in the program.
 1. flag
 2. signal
 3. sentinel
 4. siren
8. How does the character 'A' compare to the character 'B'?
 1. 'A' is greater than 'B'
 2. 'A' is less than 'B'
 3. 'A' is equal to 'B'
 4. You cannot compare characters
9. This is an `if` statement that appears inside another `if` statement.
 1. nested `if` statement
 2. tiered `if` statement
 3. dislodged `if` statement

4. structured `if` statement

10. An `else` clause always goes with .

1. the closest previous `if` clause that doesn't already have its own `else` clause
2. the closest `if` clause
3. the `if` clause that is randomly selected by the compiler
4. none of these

11. When determining whether a number is inside a range, it's best to use this operator.

1. `&&`
2. `!`
3. `||`
4. `? :`

12. This determines whether two different `String` objects contain the same string.

1. the `==` operator
2. the `=` operator
3. the `equals` method
4. the `stringCompare` method

13. The conditional operator takes this many operands.

1. one
2. two
3. three
4. four

14. This section of a `switch` statement is branched to if none of the `case` expressions match the `switch` expression.

1. `else`
2. `default`

3. `case`
 4. `otherwise`
15. You can use this method to display formatted output in a console window.
1. `Format.out.println`
 2. `Console.format`
 3. `System.out.printf`
 4. `System.out.formatted`
16. **True or False:** The `=` operator and the `==` operator perform the same operation.
17. **True or False:** A conditionally executed statement should be indented one level from the `if` clause.
18. **True or False:** All lines in a conditionally executed block should be indented one level.
19. **True or False:** When an `if` statement is nested in the `if` clause of another statement, the only time the inner `if` statement is executed is when the `boolean` expression of the outer `if` statement is `true`.
20. **True or False:** When an `if` statement is nested in the `else` clause of another statement, the only time the inner `if` statement is executed is when the `boolean` expression of the outer `if` statement is `true`.
21. **True or False:** The scope of a variable is limited to the block in which it is defined.

Find the Error

Find the errors in the following code:

1.

```
// Warning! This code contains ERRORS!
if (x == 1);
    y = 2;
else if (x == 2);
    y = 3;
else if (x == 3);
    y = 4;
```
2.

```
// Warning! This code contains an ERROR!
if (average = 100)
    System.out.println("Perfect Average!");
```


3. // Warning! This code contains ERRORS!

```
if (num2 == 0)
    System.out.println("Division by zero is not possible.");
    System.out.println("Please run the program again ");
    System.out.println("and enter a number besides zero.");
else
    Quotient = num1 / num2;
    System.out.print("The quotient of " + Num1);
    System.out.print(" divided by " + Num2 + " is ");
    System.out.println(Quotient);
```

4. // Warning! This code contains ERRORS!

```
switch (score)
{
    case (score > 90):
        grade = 'A';
        break;
    case(score > 80):
        grade = 'b';
        break;
    case(score > 70):
        grade = 'C';
        break;
    case (score > 60):
        grade = 'D';
        break;
    default:
        grade = 'F';
}
```

5. The following statement should determine whether `x` is not greater than 20. What is wrong with it?

```
if (!x > 20)
```

6. The following statement should determine whether `count` is within the range of 0 through 100. What is wrong with it?

```
if (count >= 0 || count <= 100)
```

7. The following statement should determine whether `count` is outside the range of 0 through 100. What is wrong with it?

```
if (count < 0 && count > 100)
```

8. The following statement should assign 0 to `z` if `a` is less than 10; otherwise, it should assign 7 to `z`. What is wrong with it?

```
z = (a < 10) : 0 ? 7;
```

9. Assume that `partNumber` references a `String` object. The following `if` statement should perform a case-insensitive comparison. What is wrong with it?

```
if (partNumber.equals("BQ789W4"))  
    available = true;
```

10. What is wrong with the following code?

```
double value = 12345.678;  
System.out.printf("%.2d", value);
```

Algorithm Workbench

1. Write an `if` statement that assigns 100 to `x` when `y` is equal to 0.
2. Write an `if-else` statement that assigns 0 to `x` when `y` is equal to 10. Otherwise, it should assign 1 to `x`.
3. Using the following chart, write an `if-else-if` statement that assigns .10, .15, or .20 to `commission`, depending on the value in `sales`.

Sales	Commission Rate
Up to \$10,000	10%
\$10,000 to \$15,000	15%
Over \$15,000	20%

4. Write an `if` statement that sets the variable `hours` to 10 when the `boolean` flag variable `minimum` is equal to `true`.
5. Write nested `if` statements that perform the following tests: If `amount1` is greater than 10 and `amount2` is less than 100, display the greater of the two.
6. Write an `if` statement that prints the message “The number is valid” if the variable `grade` is within the range 0 through 100.
7. Write an `if` statement that prints the message “The number is valid” if the variable `temperature` is within the range -50 through 150.
8. Write an `if` statement that prints the message “The number is not valid” if the variable `hours` is outside the range 0 through 80.

9. Write an `if-else` statement that displays the `String` objects `title1` and `title2` in alphabetical order.

10. Convert the following `if-else-if` statement into a `switch` statement:

```
if (choice == 1)
{
    System.out.println("You selected 1.");
}
else if (choice == 2 || choice == 3)
{
    System.out.println("You selected 2 or 3.");
}
else if (choice == 4)
{
    System.out.println("You selected 4.");
}
else
{
    System.out.println("Select again please.");
}
```

11. Match the conditional expression with the `if-else` statement that performs the same operation.

1. `q = x < y ? a + b : x * 2;`

2. `q = x < y ? x * 2 : a + b;`

3. `q = x < y ? 0 : 1;`

```
-----    if (x < y)
            q = 0;
else
            q = 1;
-----    if (x < y)
            q = a + b;
else
            q = x * 2;
-----    if (x < y)
            q = x * 2;
else
            q = a + b;
```

12. Assume the `double` variable `number` contains the value 12345.6789. Write a statement that uses `System.out.printf` to display the number as 12345.7.

13. Assume the `double` variable `number` contains the value 12345.6789. Write a statement that uses `System.out.printf` to display the number as 12,345.68.

14. Assume the `int` variable `number` contains the value 1234567. Write a statement that uses `System.out.printf` to display the number as 1,234,567.

Short Answer

1. Explain what is meant by the phrase “conditionally executed.”
2. Explain why a misplaced semicolon can cause an `if` statement to operate incorrectly.
3. Why is it good advice to indent all the statements inside a set of braces?
4. What happens when you compare two `String` objects with the `==` operator?
5. Explain the purpose of a flag variable. Of what data type should a flag variable be?
6. What risk does a programmer take when not placing a trailing `else` at the end of an `if-else-if` statement?
7. Briefly describe how the `&&` operator works.
8. Briefly describe how the `||` operator works.
9. Why are the relational operators called “relational”?
10. When does a constructor execute? What is its purpose?

Programming Challenges

MyProgrammingLab™ Visit www.myprogramminglab.com to complete many of these Programming Challenges online and get instant feedback.

1. Roman Numerals

Write a program that prompts the user to enter a number within the range of 1 through 10. The program should display the Roman numeral version of that number. If the number is outside the range of 1 through 10, the program should display an error message.

2. Magic Dates

The date June 10, 1960, is special because when we write it in the following format, the month times the day equals the year:

6/10/60

Write a program that asks the user to enter a month (in numeric form), a day, and a two-digit year. The program should then determine whether the month times the day is equal to the year. If so, it should display a message saying the date is magic. Otherwise, it should display a message saying the date is not magic.

3. Body Mass Index

Write a program that calculates and displays a person's body mass index (BMI). The BMI is often used to determine whether a person is overweight or underweight for his or her height. A person's BMI is calculated with the following formula:

$$B M I = W e i g h t \times 703 / H e i g h t ^2$$

where *weight* is measured in pounds and *height* is measured in inches. The program should display a message indicating whether the person has optimal weight, is underweight, or is overweight. A sedentary person's weight is considered optimal if his or her BMI is between 18.5 and 25. If the BMI is less than 18.5, the person is considered underweight. If the BMI value is greater than 25, the person is considered overweight.

4. Test Scores and Grade

Write a program that has variables to hold three test scores. The program should ask the user to enter three test scores and then assign the values entered to the variables. The program should display the average of the test scores and the letter grade that is assigned for the test score average. Use the grading scheme in the following table:

Test Score	Average Letter Grade

90–100	A

80–89	B

70–79	C

60–69	D

Below 60	F

5. Mass and Weight

Scientists measure an object's mass in kilograms and its weight in Newtons. If you know the amount of mass that an object has, you can calculate its weight, in Newtons, with the following formula:

$$\text{Weight} = \text{Mass} \times 9.8$$

Write a program that asks the user to enter an object's mass, and then calculate its weight. If the object weighs more than 1,000 Newtons, display a message indicating that it is too heavy. If the object weighs less than 10 Newtons, display a message indicating that the object is too light.

6. Time Calculator

Write a program that asks the user to enter a number of seconds.



- There are 60 seconds in a minute. If the number of seconds entered by the user is greater than or equal to 60, the program should display the number of minutes in that many seconds.
- There are 3,600 seconds in an hour. If the number of seconds entered by the user is greater than or equal to 3,600, the program should display the number of hours in that many seconds.
- There are 86,400 seconds in a day. If the number of seconds entered by the user is greater than or equal to 86,400, the program should display the number of days in that many seconds.

7. Sorted Names

Write a program that asks the user to enter three names, and then displays the names sorted in ascending order. For example, if the user entered “Charlie”, “Leslie”, and “Andy”, the program would display:

- Andy
- Charlie
- Leslie

8. Software Sales

A software company sells a package that retails for \$99. Quantity discounts are given according to the following table:

Quantity	Discount
10–19	20%
20–49	30%
50–99	40%
100 or more	50%

Write a program that asks the user to enter the number of packages purchased. The program should

then display the amount of the discount (if any) and the total amount of the purchase after the discount.

9. Shipping Charges

The Fast Freight Shipping Company charges the following rates:

Weight of Package	Rate per 500 Miles Shipped
2 pounds or less	\$1.10
Over 2 pounds but not more than 6 pounds	\$2.20
Over 6 pounds but not more than 10 pounds	\$3.70
Over 10 pounds	\$3.80

The shipping charges per 500 miles are not prorated. For example, if a 2-pound package is shipped 550 miles, the charges would be \$2.20. Write a program that asks the user to enter the weight of a package and then displays the shipping charges.

10. Fat Gram Calculator

Write a program that asks the user to enter the number of calories and fat grams in a food item. The program should display the percentage of the calories that come from fat. One gram of fat has 9 calories; therefore:

$$\text{Calories from fat} = \text{Fat grams} * 9$$

The percentage of calories from fat can be calculated as follows:

$$\text{Calories from fat} \div \text{Total calories}$$

If the calories from fat are less than 30 percent of the total calories of the food, it should also display a message indicating the food is low in fat.

 NOTE:

The number of calories from fat cannot be greater than the total number of calories in the food item. If the program determines that the number of calories from fat is greater than the number of calories in the food item, it should display an error message indicating that the input is invalid.

11. Running the Race

Write a program that asks for the names of three runners and the time, in minutes, it took each of them to finish a race. The program should display the names of the runners in the order that they finished.

12. The Speed of Sound

The following table shows the approximate speed of sound in air, water, and steel:

Medium	Speed
Air	1,100 feet per second
Water	4,900 feet per second
Steel	16,400 feet per second

Write a program that asks the user to enter “air”, “water”, or “steel”, and the distance that a sound wave will travel in the medium. The program should then display the amount of time it will take. You can calculate the amount of time it takes sound to travel in air with the following formula:

$$\text{Time} = \text{Distance} / 1,100$$

You can calculate the amount of time it takes sound to travel in water with the following formula:

$$\text{Time} = \text{Distance} / 4,900$$

You can calculate the amount of time it takes sound to travel in steel with the following formula:

$$\text{Time} = \text{Distance} / 16,400$$

13. Mobile Service Provider

A mobile phone service provider has three different subscription packages for its customers:

Package A: For \$39.99 per month, 450 minutes are provided. Additional minutes are \$0.45 per minute.

Package B: For \$59.99 per month, 900 minutes are provided. Additional minutes are \$0.40 per minute.

Package C: For \$69.99 per month, unlimited minutes are provided.

Write a program that calculates a customer's monthly bill. It should ask the user to enter the letter of the package the customer has purchased (A, B, or C) and the number of minutes that were used. The program should display the total charges.

14. Mobile Service Provider, Part 2

Modify the program you wrote for Programming Challenge 13 so it also calculates and displays the amount of money Package A customers would save if they purchased Package B or C, and the amount of money Package B customers would save if they purchased Package C. If there would be no savings, no message should be printed.

15. Bank Charges

A bank charges a base fee of \$10 per month, plus the following check fees for a commercial checking account:

- \$.10 each for less than 20 checks
- \$.08 each for 20–39 checks
- \$.06 each for 40–59 checks
- \$.04 each for 60 or more checks

Write a program that asks for the number of checks written for the month. The program should then calculate and display the bank's service fees for the month.

16. Book Club Points

Serendipity Booksellers has a book club that awards points to its customers based on the number of books purchased each month. The points are awarded as follows:

- If a customer purchases 0 books, he or she earns 0 points.
- If a customer purchases 1 book, he or she earns 5 points.
- If a customer purchases 2 books, he or she earns 15 points.
- If a customer purchases 3 books, he or she earns 30 points.
- If a customer purchases 4 or more books, he or she earns 60 points.

Write a program that asks the user to enter the number of books that he or she has purchased this month and then displays the number of points awarded.

17. Wi-Fi Diagnostic Tree

[Figure 3-23](#) shows a simplified flowchart for troubleshooting a bad Wi-Fi connection. Use the flowchart to create a program that leads a person through the steps of fixing a bad Wi-Fi connection. Here is an example of the program's output:




Figure 3-23 Troubleshooting a bad Wi-Fi connection

[Figure 3-23 Full Alternative Text](#)

```
Reboot the computer and try to connect.  
Did that fix the problem?  no [Enter]  
Reboot the router and try to connect.  
Did that fix the problem?  yes [Enter]
```

Notice that the program ends as soon as a solution is found to the problem. Here is another example of the program's output:

```
Reboot the computer and try to connect.  
Did that fix the problem?  no [Enter]  
Reboot the router and try to connect.  
Did that fix the problem?  no [Enter]  
Make sure the cables between the router & modem are plugged in firmly.  
Did that fix the problem?  no [Enter]  
Move the router to a new location.  
Did that fix the problem?  no [Enter]  
Get a new router.
```

18. Restaurant Selector

You have a group of friends coming to visit for your high school reunion, and you want to take them out to eat at a local restaurant. You aren't sure if any of them have dietary restrictions, but your restaurant choices are as follows:

- *Joe's Gourmet Burgers* – Vegetarian: No, Vegan: No, Gluten-Free: No
- *Main Street Pizza Company* – Vegetarian: Yes, Vegan: No, Gluten-Free: Yes
- *Corner Café* – Vegetarian: Yes, Vegan: Yes, Gluten-Free: Yes
- *Mama's Fine Italian* – Vegetarian: Yes, Vegan: No, Gluten-Free: No
- *The Chef's Kitchen* – Vegetarian: Yes, Vegan: Yes, Gluten-Free: Yes

Write a program that asks whether any members of your party are vegetarian, vegan, or gluten-free, and then displays only the restaurants that you may take the group to. Here is an example of the program's output:

```
Is anyone in your party a vegetarian? yes [Enter]
```

```
Is anyone in your party a vegan? no [Enter]  
Is anyone in your party gluten-free? yes [Enter]  
Here are your restaurant choices:  
    Main Street Pizza Company  
    Corner Cafe  
    The Chef's Kitchen
```

Here is another example of the program's output:

```
Is anyone in your party a vegetarian? yes [Enter]  
Is anyone in your party a vegan? yes [Enter]  
Is anyone in your party gluten-free? yes [Enter]  
Here are your restaurant choices:  
    Corner Cafe  
    The Chef's Kitchen
```

Chapter 4 Loops and Files

TOPICS

1. [4.1 The Increment and Decrement Operators](#)
2. [4.2 The `while` Loop](#)
3. [4.3 Using the `while` Loop for Input Validation](#)
4. [4.4 The `do-while` Loop](#)
5. [4.5 The `for` Loop](#)
6. [4.6 Running Totals and Sentinel Values](#)
7. [4.7 Nested Loops](#)
8. [4.8 The `break` and `continue` Statements \(Optional\)](#)
9. [4.9 Deciding Which Loop to Use](#)
10. [4.10 Introduction to File Input and Output](#)
11. [4.11 Generating Random Numbers with the `Random` Class](#)
12. [4.12 Common Errors to Avoid](#)

4.1 The Increment and Decrement Operators

CONCEPT: `++` and `--` are operators that add and subtract one from their operands.

To *increment* a value means to increase it by one, and to *decrement* a value means to decrease it by one. Both of the following statements increment the variable `number`:

```
number = number + 1;
number += 1;
```

And `number` is decremented in both of the following statements:

```
number = number - 1;
number -= 1;
```

Java provides a set of simple unary operators designed just for incrementing and decrementing variables. The increment operator is `++` and the decrement operator is `--`. The following statement uses the `++` operator to increment `number`:

```
number++;
```

And the following statement decrements `number`:

```
number--;
```



The expression `number++` is pronounced “number plus plus,” and `number--` is pronounced “number minus minus.”

The program in [Code Listing 4-1](#) demonstrates the `++` and `--` operators.

Code Listing 4-1 (IncrementDecrement.java)

```
1  /**
2   * This program demonstrates the ++ and -- operators.
3   */
4
5  public class IncrementDecrement
6  {
7      public static void main(String[] args)
8      {
9          int number = 4; // number starts out with 4
10
```

```

11 // Display the value in number.
12 System.out.println("number is " + number);
13 System.out.println("I will increment number.");
14
15 // Increment number.
16 number++;
17
18 // Display the value in number again.
19 System.out.println("Now, number is " + number);
20 System.out.println("I will decrement number.");
21
22 // Decrement number.
23 number--;
24
25 // Display the value in number once more.
26 System.out.println("Now, number is " + number);
27 }
28 }

```

Program Output

```

number is 4
I will increment number.
Now, number is 5
I will decrement number.
Now, number is 4

```

The statements in [Code Listing 4-1](#) show the increment and decrement operators used in *postfix mode*, which means the operator is placed after the variable. The operators also work in *prefix mode*, where the operator is placed before the variable name as follows:

```

++number;
--number;

```

In both postfix and prefix mode, these operators add one to or subtract one from their operand. [Code Listing 4-2](#) demonstrates this.

Code Listing 4-2 (Prefix.java)

```

1  /**
2   * This program demonstrates the ++ and -- operators
3   * in prefix mode.
4   */
5
6  public class Prefix
7  {
8      public static void main(String[] args)
9      {
10         int number = 4; // number starts out with 4
11

```

```

12      // Display the value in number.
13      System.out.println("number is " + number);
14      System.out.println("I will increment number.");
15
16      // Increment number.
17      ++number;
18
19      // Display the value in number again.
20      System.out.println("Now, number is " + number);
21      System.out.println("I will decrement number.");
22
23      // Decrement number.
24      --number;
25
26      // Display the value in number once again.
27      System.out.println("Now, number is " + number);
28  }
29  }

```

Program Output

```

number is 4
I will increment number.
Now, number is 5
I will decrement number.
Now, number is 4

```

The Difference between Postfix and Prefix Modes

In [Code Listings 4-1](#) and [4-2](#), the statements `number++` and `++number` increment the variable `number`, while the statements `number--` and `--number` decrement the variable `number`. In these simple statements, it doesn't matter whether the operator is used in postfix or prefix mode. The difference is important, however, when these operators are used in statements that do more than just increment or decrement. For example, look at the following code:

```

number = 4;
System.out.println(number++);

```

The statement that calls the `println` method does two things: (1) calls `println` to display the value of `number`, and (2) increments `number`. But which happens first? The `println` method will display a different value if `number` is incremented first than if `number` is incremented last. The answer depends upon the mode of the increment operator.

Postfix mode causes the increment to happen after the value of the variable is used in the expression. In the previously shown statement, the `println` method will display 4 and then `number` will be incremented to 5. Prefix mode, however, causes the increment to happen first. Here is an example:

```
number = 4;  
System.out.println(++number);
```

In these statements, `number` is incremented to 5, then `println` will display the value in `number` (which is 5). For another example, look at the following code:

```
int x = 1, y;  
y = x++;           // Postfix increment
```

The first statement declares the variable `x` (initialized with the value 1) and the variable `y`. The second statement does the following:

- It assigns the value of `x` to the variable `y`.
- The variable `x` is incremented.

The value that will be stored in `y` depends on when the increment takes place. Because the `++` operator is used in postfix mode, it acts after the assignment takes place. So, this code will store 1 in `y`. After the code has executed, `x` will contain 2. Let's look at the same code, but with the `++` operator used in prefix mode as follows:

```
int x = 1, y;  
y = ++x;           // Prefix increment
```

The first statement declares the variable `x` (initialized with the value 1) and the variable `y`. In the second statement, the `++` operator is used in prefix mode, so it acts on the variable before the assignment takes place. So, this code will store 2 in `y`. After the code has executed, `x` will also contain 2.

Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 4.1 What will the following program segments display?

1. `x = 2;`
 `y = x++;`
 `System.out.println(y);`
2. `x = 2;`
 `System.out.println(x++);`
3. `x = 2;`
 `System.out.println(--x);`
4. `x = 8;`
 `y = x--;`

```
System.out.println(y);
```

4.2 The while Loop

CONCEPT: A loop is part of a program that repeats.

In [Chapter 3](#), you were introduced to the concept of control structures, which direct the flow of a program. A *loop* is a control structure that causes a statement or group of statements to repeat. Java has three looping control structures: the `while` loop, the `do-while` loop, and the `for` loop. The difference among each of these is how they control the repetition. In this section we will focus on the `while` loop.

The `while` loop has two important parts: (1) a `boolean` expression that is tested for a `true` or `false` value, and (2) a statement or block of statements that is repeated as long as the expression is `true`. [Figure 4-1](#) shows the logic of a `while` loop.




Figure 4-1 Logic of a `while` loop



The while Loop

Here is the general format of the `while` loop:

```
while (BooleanExpression)
    Statement;
```

In the general format, *BooleanExpression* is any valid `boolean` expression, and *Statement* is any valid Java statement. The first line shown in the format is sometimes called the *loop header*. It consists of the key word `while` followed by the *BooleanExpression* enclosed in parentheses.

Here's how the loop works: The *BooleanExpression* is tested, and if it is `true`, the *Statement* is executed. Then, the *BooleanExpression* is tested again. If it is `true`, the *Statement* is executed. This cycle repeats until the *BooleanExpression* is `false`.

The statement that is repeated is known as the *body* of the loop. It is also considered a conditionally executed statement, because it is only executed under the condition that the *BooleanExpression* is `true`.

Notice there is no semicolon at the end of the loop header. Like the `if` statement, the `while` loop is not complete without the conditionally executed statement that follows it.

If you wish the `while` loop to repeat a block of statements, the format is as follows:

```
while (BooleanExpression)
{
    Statement;
    Statement;
    // Place as many statements here
    // as necessary.
}
```

The `while` loop works like an `if` statement that executes over and over. As long as the expression in the parentheses is `true`, the conditionally executed statement or block will repeat. The program in [Code Listing 4-3](#) uses the `while` loop to print “Hello” five times.

Code Listing 4-3 (WhileLoop.java)

```
1  /**
2   * This program demonstrates the while loop.
3   */
4
5  public class WhileLoop
6  {
7      public static void main(String[] args)
8      {
9          int number = 1;
10
11         while (number <= 5)
12         {
13             System.out.println("Hello");
14             number++;
15         }
16
17         System.out.println("That's all!");
18     }
19 }
```

Program Output

```
Hello
Hello
Hello
Hello
Hello
Hello
That's all!
```

Let's take a closer look at this program. An integer variable, `number`, is declared and initialized with the value 1. The `while` loop begins with the following statement:

```
while (number <= 5)
```

This statement tests the variable `number` to determine whether it is less than or equal to 5. If it is, then the statements in the body of the loop are executed as follows:

```
System.out.println("Hello");
number++;
```

The first statement in the body of the loop prints the word “Hello”. The second statement uses the increment operator to add one to `number`. This is the last statement in the body of the loop, so after it executes, the loop starts over. It tests the `boolean` expression again, and if it is `true`, the statements in the body of the loop are executed. This cycle repeats until the `boolean` expression `number <= 5` is `false`, as illustrated in [Figure 4-2](#).




Figure 4-2 The `while` loop

[Figure 4-2 Full Alternative Text](#)

Each repetition of a loop is known as an *iteration*. This loop will perform five iterations because the variable `number` is initialized with the value 1, and it is incremented each time the body of the loop is executed. When the expression `number <= 5` is tested and found to be `false`, the loop will terminate and the program will resume execution at the statement that immediately follows the loop. [Figure 4-3](#) shows the logic of this loop.




Figure 4-3 Logic of the example `while` loop

[Figure 4-3 Full Alternative Text](#)

In this example, the `number` variable is referred to as the *loop control variable* because it controls the number of times that the loop iterates.

The `while` Loop Is a Pretest Loop

The `while` loop is known as a *pretest* loop, which means it tests its expression before each iteration. Notice the variable declaration of `number` in [Code Listing 4-3](#):

```
int number = 1;
```

The `number` variable is initialized with the value 1. If `number` had been initialized with a value that is greater than 5, as shown in the following program segment, the loop would never execute:

```
int number = 6;
while (number <= 5)
{
    System.out.println("Hello");
    number++;
}
```

An important characteristic of the `while` loop is that the loop will never iterate if the `boolean` expression is `false` to start with. If you want to be sure that a `while` loop executes the first time, you must initialize the relevant data in such a way that the `boolean` expression starts out as `true`.

Infinite Loops

In all but rare cases, loops must contain a way to terminate within themselves. This means that something inside the loop must eventually make the `boolean` expression `false`. The loop in [Code Listing 4-3](#) stops when the expression `number <= 5` is `false`.

If a loop does not have a way of stopping, it is called an *infinite loop* as it continues to repeat until the program is interrupted. Here is an example of an infinite loop:

```
int number = 1;
while (number <= 5)
{
    System.out.println("Hello");
}
```

This is an infinite loop because it does not contain a statement that changes the value of the `number` variable. Each time the `boolean` expression is tested, `number` will contain the value 1.

It's also possible to create an infinite loop by accidentally placing a semicolon after the first line of the `while` loop. Here is an example:

```
int number = 1;
while (number <= 5); // This semicolon is an ERROR!
{
    System.out.println("Hello");
    number++;
}
```

The semicolon at the end of the first line is assumed to be a `null` statement and disconnects the `while` statement from the block that comes after it. To the compiler, this loop looks like the following:

```
while (number <= 5);
```

This `while` loop will forever execute the null statement, which does nothing. The program will appear to have “gone into space” because there is nothing to display screen output or show activity.

Don't Forget the Braces with a Block of Statements

If you are using a block of statements, don't forget to enclose all of the statements in a set of braces. If the braces are accidentally left out, the `while` statement conditionally executes only the very next statement. For example, look at the following code:

```
int number = 1;
// This loop is missing its braces!
while (number <= 5)
```

```
System.out.println("Hello");  
number++;
```

In this code the `number++` statement is not in the body of the loop. Because the braces are missing, the `while` statement executes only the statement that immediately follows it. This loop will execute infinitely because there is no code in its body that changes the `number` variable.

Programming Style and the `while` Loop

It's possible to create loops that look like the following:

```
while (number <= 5) { System.out.println("Hello"); number++; }
```

Avoid this style of programming. The programming style you should use with the `while` loop is similar to that of the `if` statement as follows:

- If there is only one statement repeated by the loop, it should appear on the line after the `while` statement and be indented one additional level. The statement can optionally appear inside a set of braces.
- If the loop repeats a block, each line inside the braces should be indented.

This programming style should visually set the body of the loop apart from the surrounding code. In general, you'll find a similar style being used with the other types of loops presented in this chapter.

In the Spotlight:



Designing a Program with a `while` Loop

A project currently underway at Chemical Labs, Inc., requires that a substance be continually heated in a vat. A technician must check the substance's temperature every 15 minutes. If the substance's temperature does not exceed 102.5 degrees Celsius, then the technician does nothing. However, if the temperature is greater than 102.5 degrees Celsius, the technician must turn down the vat's thermostat, wait 5 minutes, and check the temperature again. The technician repeats these steps until the temperature does not exceed 102.5 degrees Celsius. The director of engineering has asked you to write a program that guides the technician through this process.

Here is the algorithm:

1. *Prompt the user to enter the substance's temperature.*

2. *Repeat the following steps as long as the temperature is greater than 102.5 degrees Celsius:*
 1. *Tell the technician to turn down the thermostat, wait 5 minutes, and check the temperature again.*
 2. *Prompt the user to enter the substance's temperature.*
3. *After the loop finishes, tell the technician that the temperature is acceptable and to check it again in 15 minutes.*

After reviewing this algorithm, you realize that Steps 2(a) and 2(b) should not be performed if the test condition (temperature is greater than 102.5) is `false` to begin with. The `while` loop will work well in this situation, because it will not execute even once if its condition is `false`. [Code Listing 4-4](#) shows the program.

Code Listing 4-4 (CheckTemperature.java)

```
1  import java.util.Scanner;
2
3  /**
4   * This program assists a technician in the process
5   * of checking a substance's temperature.
6   */
7  public class CheckTemperature
8  {
9      public static void main(String[] args)
10     {
11         final double MAX_TEMP = 102.5;    // Maximum temperature
12         double temperature;               // To hold the temperature
13
14         // Create a Scanner object for keyboard input.
15         Scanner keyboard = new Scanner(System.in);
16
17         // Get the current temperature.
18         System.out.print("Enter the substance's Celsius temperature: ");
19         temperature = keyboard.nextDouble();
20
21         // As long as necessary, instruct the technician
22         // to adjust the temperature.
23         while (temperature > MAX_TEMP)
24         {
25             System.out.println("The temperature is too high. Turn the");
26             System.out.println("thermostat down and wait 5 minutes.");
27             System.out.println("Then, take the Celsius temperature again");
28             System.out.print("and enter it here: ");
29             temperature = keyboard.nextDouble();
30         }
31
32         // Remind the technician to check the temperature
33         // again in 15 minutes.
34         System.out.println("The temperature is acceptable.");
35         System.out.println("Check it again in 15 minutes.");
```

```
36     }  
37 }
```

Program Output with Example Input Shown in Bold

Enter the substance's Celsius temperature: **104.7** [Enter]
The temperature is too high. Turn the thermostat down and wait 5 minutes.
Then, take the Celsius temperature again and enter it here: **103.2** [Enter]
The temperature is too high. Turn the thermostat down and wait 5 minutes.
Then, take the Celsius temperature again and enter it here: **102.1** [Enter]
The temperature is acceptable.
Check it again in 15 minutes.



Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 4.2 How many times will "Hello World" be printed in the following program segment?

```
int count = 10;  
while (count < 1)  
{  
    System.out.println("Hello World");  
    count++;  
}
```

2. 4.3 How many times will "I love Java programming!" be printed in the following program segment?

```
int count = 0;  
while (count < 10)  
    System.out.println("I love Java programming!");
```

4.3 Using the `while` Loop for Input Validation

CONCEPT: The `while` loop can be used to create input routines that repeat until acceptable data is entered.

Perhaps the most famous saying of the computer industry is “garbage in, garbage out.” The integrity of a program’s output is only as good as its input, so you should try to make sure garbage does not go into your programs. *Input validation* is the process of inspecting data given to a program by the user and determining whether it is valid. A good program should give clear instructions about the kind of input that is acceptable, and not assume the user has followed those instructions.

The `while` loop is especially useful for validating input. If an invalid value is entered, a loop can require that the user reenter it as many times as necessary. For example, the loop in the following code asks for a number in the range of 1 through 100:

```
int number;
// Create a Scanner object for keyboard input.
Scanner keyboard = new Scanner(System.in);
// Get a number from the user.
System.out.print("Enter a number in the range of 1 through 100: ");
number = keyboard.nextInt();
// Validate the input.
while (number < 1 || number > 100)
{
    System.out.print("Invalid input. Enter a number in the range " +
                    "of 1 through 100: ");
    number = keyboard.nextInt();
}
```

This code first allows the user to enter a number. This takes place just before the loop. If the input is valid, the loop will not execute. If the input is invalid, however, the loop will display an error message and require the user to enter another number. The loop will continue to execute until the user enters a valid number. The general logic of performing input validation is shown in [Figure 4-4](#).




Figure 4-4 Input validation logic

[Figure 4-4 Full Alternative Text](#)


The read operation that takes place just before the loop is called a *priming read*. It provides the first value for the loop to test. Subsequent values are obtained by the loop.

The program in [Code Listing 4-5](#) calculates the number of soccer teams a youth league may create, based on a given number of players and a maximum number of players per team. The program uses `while` loops (in lines 28 through 36 and lines 44 through 49) to validate the user's input. [Figure 4-5](#) shows an example of interaction with the program.

This input dialog box appears first. The user enters 4 (an invalid value) and clicks the OK button.




This input dialog box appears next. The user enters 12 and clicks the OK button.



This input dialog box appears next. The user enters -142 (an invalid value) and clicks the OK button.



This input dialog box appears next. The user enters 142 and clicks the OK button.



This message dialog box appears next.




Figure 4-5 Interaction with the SoccerTeams program

(Oracle Corporate Counsel)

[Figure 4-5 Full Alternative Text](#)

Code Listing 4-5 (SoccerTeams.java)


```

1 import javax.swing.JOptionPane;
2
3 /**
4  This program calculates the number of soccer teams
5  that a youth league may create from the number of
6  available players. Input validation is demonstrated
7  with while loops.
8  */
9
10 public class SoccerTeams
11 {
12     public static void main(String[] args)
13     {
14         final int MIN_PLAYERS = 9;        // Minimum players per team
15         final int MAX_PLAYERS = 15;       // Maximum players per team
16         int players;                      // Number of available players
17         int teamSize;                     // Number of players per team
18         int teams;                       // Number of teams
19         int leftOver;                    // Number of leftover players
20         String input;                     // To hold the user input
21
22         // Get the number of players per team.
23         input = JOptionPane.showInputDialog("Enter the number of " +
24                                             "players per team.");
25         teamSize = Integer.parseInt(input);
26
27         // Validate the number entered.
28         while (teamSize < MIN_PLAYERS || teamSize > MAX_PLAYERS)
29         {
30             input = JOptionPane.showInputDialog("The number must " +
31                                                 "be at least " + MIN_PLAYERS +
32                                                 " and no more than " +
33                                                 MAX_PLAYERS + ".\n Enter " +
34                                                 "the number of players.");
35             teamSize = Integer.parseInt(input);
36         }
37
38         // Get the number of available players.
39         input = JOptionPane.showInputDialog("Enter the available " +
40                                             "number of players.");
41         players = Integer.parseInt(input);
42
43         // Validate the number entered.
44         while (players < 0)
45         {
46             input = JOptionPane.showInputDialog("Enter 0 or " +
47                                                 "greater.");
48             players = Integer.parseInt(input);
49         }
50
51         // Calculate the number of teams.
52         teams = players / teamSize;
53
54         // Calculate the number of leftover players.
55         leftOver = players % teamSize;
56
57         // Display the results.
58         JOptionPane.showMessageDialog(null, "There will be " +

```

```

59
60
61
62         System.exit(0);
63     }
64 }

teams + " teams with " +
leftOver +
" players left over.");

```



MyProgrammingLab™ www.myprogramminglab.com

1. 4.4 Write an input validation loop that asks the user to enter a number in the range of 10 through 24.
2. 4.5 Write an input validation loop that asks the user to enter ‘Y’, ‘y’, ‘N’, or ‘n’.
3. 4.6 Write an input validation loop that asks the user to enter “Yes” or “No”.

4.4 The do-while Loop

CONCEPT: The `do-while` loop is a posttest loop, which means its `boolean` expression is tested after each iteration.

The `do-while` loop looks something like an inverted `while` loop. Here is the `do-while` loop's format when the body of the loop contains only a single statement:

```
do
    Statement;
while (BooleanExpression);
```

Here is the format of the `do-while` loop when the body of the loop contains multiple statements:

```
do
{
    Statement;
    Statement;
    // Place as many statements here as necessary.
} while (BooleanExpression);
```

NOTE:

The `do-while` loop must be terminated with a semicolon.

The `do-while` loop is a *posttest* loop. This means it does not test its `boolean` expression until it has completed an iteration. As a result, the `do-while` loop always performs at least one iteration, even if the `boolean` expression is `false` to begin with. This differs from the behavior of a `while` loop, which you will recall is a pretest loop. For example, in the following `while` loop the `println` statement will not execute at all:

```
int x = 1;
while (x < 0)
    System.out.println(x);
```

But the `println` statement in the following `do-while` loop will execute once because the `do-while` loop does not evaluate the expression `x < 0` until the end of the iteration:

```
int x = 1;
do
    System.out.println(x);
while (x < 0);
```

[Figure 4-6](#) illustrates the logic of the `do-while` loop.




Figure 4-6 Logic of the do-while loop

You should use the `do-while` loop when you want to make sure the loop executes at least once. For example, the program in [Code Listing 4-6](#) averages a series of three test scores for a student. After the average is displayed, it asks the user whether he or she wants to average another set of test scores. The program repeats as long as the user enters `y` for yes.

Code Listing 4-6 (TestAverage1.java)

```

1  import java.util.Scanner;
2
3  /**
4   * This program demonstrates a user controlled loop.
5   */
6
7  public class TestAverage1
8  {
9      public static void main(String[] args)
10     {
11         int score1, score2, score3;           // Three test scores
12         double average;                       // Average test score
13         char repeat;                          // To hold 'y' or 'n'
14         String input;                         // To hold input
15
16         System.out.println("This program calculates the " +
17                             "average of three test scores.");
18
19         // Create a Scanner object for keyboard input.
20         Scanner keyboard = new Scanner(System.in);
21
22         // Get as many sets of test scores as the user wants.
23         do
24         {
25             // Get the first test score in this set.
26             System.out.print("Enter score #1: ");

```

```

27         score1 = keyboard.nextInt();
28
29         // Get the second test score in this set.
30         System.out.print("Enter score #2: ");
31         score2 = keyboard.nextInt();
32
33         // Get the third test score in this set.
34         System.out.print("Enter score #3: ");
35         score3 = keyboard.nextInt();
36
37         // Consume the remaining newline.
38         keyboard.nextLine();
39
40         // Calculate and print the average test score.
41         average = (score1 + score2 + score3) / 3.0;
42         System.out.println("The average is " + average);
43         System.out.println(); // Prints a blank line
44
45         // Does the user want to average another set?
46         System.out.println("Would you like to average " +
47                             "another set of test scores?");
48         System.out.print("Enter Y for yes or N for no: ");
49         input = keyboard.nextLine(); // Read a line.
50         repeat = input.charAt(0);    // Get the first char.
51
52     } while (repeat == 'Y' || repeat == 'y');
53 }
54 }

```

Program Output with Example Input Shown in Bold

This program calculates the average of three test scores.

Enter score #1: **89 [Enter]**

Enter score #2: **90 [Enter]**

Enter score #3: **97 [Enter]**

The average is 92.0

Would you like to average another set of test scores?

Enter Y for yes or N for no: **y [Enter]**

Enter score #1: **78 [Enter]**

Enter score #2: **65 [Enter]**

Enter score #3: **88 [Enter]**

The average is 77.0

Would you like to average another set of test scores?

Enter Y for yes or N for no: **n [Enter]**

When this program was written, the programmer had no way of knowing the number of times the loop would iterate. This is because the loop asks the user whether he or she wants to repeat the process. This type of loop is known as a *user controlled loop*, because it allows the user to decide the number of iterations.

4.5 The `for` Loop

CONCEPT: The `for` loop is ideal for performing a known number of iterations.

In general, there are two categories of loops: conditional loops and count-controlled loops. A *conditional loop* executes as long as a particular condition exists. For example, an input validation loop executes as long as the input value is invalid. When you write a conditional loop, you have no way of knowing the number of times it will iterate.

Sometimes you do know the exact number of iterations that a loop must perform. A loop that repeats a specific number of times is known as a *count-controlled loop*. For example, if a loop asks the user to enter the sales amounts for each month in the year, it will iterate 12 times. In essence, the loop counts to 12 and asks the user to enter a sales amount each time it makes a count.

A count-controlled loop must possess three elements:

1. It must initialize a control variable to a starting value.
2. It must test the control variable by comparing it to a maximum value. When the control variable reaches its maximum value, the loop terminates.
3. It must update the control variable during each iteration. This is usually done by incrementing the variable.

In Java, the `for` loop is ideal for writing count-controlled loops. It is specifically designed to initialize, test, and update a loop control variable. Here is the format of the `for` loop when used to repeat a single statement:

```
for (Initialization; Test; Update)  
    Statement;
```

The format of the `for` loop when used to repeat a block is as follows:

```
for (Initialization; Test; Update)  
{  
    Statement;  
    Statement;  
    // Place as many statements here as necessary.  
}
```

The first line of the `for` loop is known as the *loop header*. After the key word `for`, there are three expressions inside the parentheses, separated by semicolons. (Notice there is not a semicolon after the third expression.) The first expression is the *initialization expression*. It is normally used to initialize a control variable to its starting value. This is the first action performed by the loop, and it is done only once. The second expression is the *test expression*. This is a `boolean` expression that

controls the execution of the loop. As long as this expression is `true`, the body of the `for` loop will repeat. The `for` loop is a pretest loop, so it evaluates the test expression before each iteration. The third expression is the *update expression*. It executes at the end of each iteration. Typically, this is a statement that increments the loop's control variable.

Here is an example of a simple `for` loop that prints "Hello" five times:

```
for (count = 1; count <= 5; count++)  
    System.out.println("Hello");
```

In this loop, the initialization expression is `count = 1`, the test expression is `count <= 5`, and the update expression is `count++`. The body of the loop has one statement, which is the `println` statement. [Figure 4-7](#) illustrates the sequence of events that takes place during the loop's execution. Notice that Steps 2 through 4 are repeated as long as the test expression is `true`.




Figure 4-7 Sequence of events in the `for` loop

[Figure 4-7 Full Alternative Text](#)

[Figure 4-8](#) shows the loop's logic in the form of a flowchart.




Figure 4-8 Logic of the `for` loop

[Figure 4-8 Full Alternative Text](#)

Notice how the control variable, `count`, is used to control the number of times that the loop iterates. During the execution of the loop, this variable takes on the values 1 through 5, and when the test expression `count <= 5` is *false*, the loop terminates. Because this variable keeps a count of the number of iterations, it is often called a *counter variable*.

Also notice that in this example the `count` variable is used only in the loop header, to control the number of loop iterations. It is not used for any other purpose. It is also possible to use the control variable within the body of the loop. For example, look at the following code:

```
for (number = 1; number <= 10; number++)
    System.out.print(number + " ");
```

The control variable in this loop is `number`. In addition to controlling the number of iterations, it is also used in the body of the loop. This loop will produce the following output:

```
1 2 3 4 5 6 7 8 9 10
```

As you can see, the loop displays the contents of the `number` variable during each iteration. The program in [Code Listing 4-7](#) shows another example of a `for` loop that uses its control variable within the body of the loop. This program displays a table showing the numbers 1 through 10 and their squares.

Code Listing 4-7 (`Squares.java`)

```
1  /**
2      This program demonstrates the for loop.
```

```

3  */
4
5  public class Squares
6  {
7      public static void main(String[] args)
8      {
9          int number; // Loop control variable
10
11         System.out.println("Number Number Squared");
12         System.out.println("-----");
13
14         for (number = 1; number <= 10; number++)
15         {
16             System.out.println(number + "\t\t" +
17                                 number * number);
18         }
19     }
20 }

```

Program Output

Number Number Squared

1	1
2	4
3	9
4	16
5	25
6	36
7	49
8	64
9	81

[Figure 4-9](#) illustrates the sequence of events performed by this `for` loop.




Figure 4-9 Sequence of events with the `for` loop in [Code Listing 4-7](#)

[Figure 4-9 Full Alternative Text](#)

[Figure 4-10](#) shows the logic of the loop.




Figure 4-10 Logic of the `for` loop in [Code Listing 4-7](#)

[Figure 4-10 Full Alternative Text](#)

The `for` Loop Is a Pretest Loop

Because the `for` loop tests its `boolean` expression before it performs an iteration, it is a pretest loop. It is possible to write a `for` loop in such a way that it will never iterate. Here is an example:

```
for (count = 11; count <= 10; count++)  
    System.out.println("Hello");
```

Because the variable `count` is initialized to a value that makes the `boolean` expression `false` from the beginning, this loop terminates as soon as it begins.

Avoid Modifying the Control Variable in the Body of the `for` Loop

Be careful not to place a statement that modifies the control variable in the body of the `for` loop. All modifications of the control variable should take place in the update expression, which is automatically executed at the end of each iteration. If a statement in the body of the loop also modifies the control variable, the loop probably will not terminate when you expect it to. The following loop, for example, increments `x` twice for each iteration:

```
for (x = 1; x <= 10; x++)  
{  
    System.out.println(x);  
    x++; // Wrong!  
}
```

Other Forms of the Update Expression

You are not limited to using increment statements in the update expression. Here is a loop that displays all the even numbers from 2 through 100 by adding 2 to its counter:

```
for (number = 2; number <= 100; number += 2)  
    System.out.println(number);
```

And here is a loop that counts backward from 10 to 0:

```
for (number = 10; number >= 0; number--)  
    System.out.println(number);
```

Declaring a Variable in the `for` Loop's Initialization Expression

Not only may the control variable be initialized in the initialization expression, but also it may be declared there. The following code shows an example. The following is a modified version of the loop in [Code Listing 4-7](#):

```
for (int number = 1; number <= 10; number++)  
{  
    System.out.println(number + "\t\t" +  
                        number * number);  
}
```

In this loop, the `number` variable is both declared and initialized in the initialization expression. If the control variable is used only in the loop, it makes sense to declare it in the loop header. This makes the variable's purpose clearer.

When a variable is declared in the initialization expression of a `for` loop, the scope of the variable is limited to the loop. This means you cannot access the variable in statements outside the loop. For example, the following program segment will not compile because the last `println` statement cannot access the variable `count`:

```
for (int count = 1; count <= 10; count++)  
    System.out.println(count);  
System.out.println("count is now " + count); // ERROR!
```

Creating a User Controlled `for` Loop

Sometimes you want the user to determine the maximum value of the control variable in a `for` loop, and therefore determine the number of times the loop iterates. For example, look at the program in [Code Listing 4-8](#). It is a modification of [Code Listing 4-7](#). Instead of displaying the numbers 1 through 10 and their squares, this program allows the user to enter the maximum value to display.

Code Listing 4-8 (`UserSquares.java`)

```
1  import java.util.Scanner;  
2  
3  /**  
4      This program demonstrates a user controlled for loop.  
5  */  
6
```

```

7 public class UserSquares
8 {
9     public static void main(String[] args)
10    {
11        int number;        // Loop control variable
12        int maxValue;      // Maximum value to display
13
14        System.out.println("I will display a table of " +
15                            "numbers and their squares.");
16
17        // Create a Scanner object for keyboard input.
18        Scanner keyboard = new Scanner(System.in);
19
20        // Get the maximum value to display.
21        System.out.print("How high should I go? ");
22        maxValue = keyboard.nextInt();
23
24        // Display the table.
25        System.out.println("Number Number Squared");
26        System.out.println("-----");
27        for (number = 1; number <= maxValue; number++)
28        {
29            System.out.println(number + "\t\t" +
30                                number * number);
31        }
32    }
33 }

```

Program Output with Example Input Shown in Bold

```

I will display a table of numbers and their squares.
How high should I go? 7 [Enter]

```

Number Number Squared

```

1      1
2      4
3      9
4     16
5     25

```

6 36

7 49

In lines 21 and 22, which are before the loop, this program asks the user to enter the highest value to display. This value is stored in the `maxValue` variable as follows:

```
System.out.print("How high should I go? ");
maxValue = keyboard.nextInt();
```

In line 27, the for loop’s test expression uses the value in the `maxValue` variable as the upper limit for the control variable as follows:

```
for (number = 1; number <= maxValue; number++)
```

In this loop, the `number` variable takes on the values 1 through `maxValue`, and then the loop terminates.

Using Multiple Statements in the Initialization and Update Expressions

It is possible to execute more than one statement in the initialization expression and the update expression. When using multiple statements in either of these expressions, simply separate the statements with commas. For example, look at the loop in the following code, which has two statements in the initialization expression:

```
int x, y;
for (x = 1, y = 1; x <= 5; x++)
{
    System.out.println(x + " plus " + y +
                        " equals " + (x + y));
}
```

This loop’s initialization expression is as follows:

```
x = 1, y = 1
```

This initializes two variables, `x` and `y`. The output produced by this loop is as follows:

```
1 plus 1 equals 2
2 plus 1 equals 3
3 plus 1 equals 4
4 plus 1 equals 5
5 plus 1 equals 6
```

We can further modify the loop to execute two statements in the update expression. Here is an example:

```
int x, y;
for (x = 1, y = 1; x <= 5; x++, y++)
{
    System.out.println(x + " plus " + y +
                        " equals " + (x + y));
}
```

The loop's update expression is as follows:

`x++, y++`

This update expression increments both the `x` and `y` variables. The output produced by this loop is as follows:

```
1 plus 1 equals 2
2 plus 2 equals 4
3 plus 3 equals 6
4 plus 4 equals 8
5 plus 5 equals 10
```

Connecting multiple statements with commas works well in the initialization and update expressions, but don't try to connect multiple `boolean` expressions this way in the test expression. If you wish to combine multiple `boolean` expressions in the test expression, you must use the `&&` or `||` operators.

In the Spotlight:



Designing a Count-Controlled `for` Loop

Your friend Amanda just inherited a European sports car from her uncle. Amanda lives in the United States, and she is afraid she will get a speeding ticket because the car's speedometer indicates kilometers per hour (KPH). She has asked you to write a program that displays a table of speeds in kilometers per hour with their values converted to miles per hour (MPH). The formula for converting KPH to MPH is

$$M P H = K P H * 0.6214$$

In the formula, *MPH* is the speed in miles per hour and *KPH* is the speed in kilometers per hour.

The table that your program displays should show speeds from 60 kilometers per hour through 130 kilometers per hour, in increments of 10, along with their values converted to miles per hour. The table should look something like this:

KPH MPH

60 37.3

70 43.5

80 49.7

etc.

130 80.8

After thinking about this table of values, you decide that you will write a `for` loop that uses a counter variable to hold the KPH speeds. The counter’s starting value will be 60, its ending value will be 130, and you will add 10 to the counter variable after each iteration. Inside the loop you will use the counter variable to calculate a speed in MPH. [Code Listing 4-9](#) shows the code.

Code Listing 4-9 (SpeedConverter.java)

```
1  /**
2   This program displays a table of speeds in
3   kph converted to mph.
4   */
5
6  public class SpeedConverter
7  {
8      public static void main(String[] args)
9      {
10         // Constants
11         final int STARTING_KPH = 60;           // Starting speed
12         final int MAX_KPH = 130;              // Maximum speed
13         final int INCREMENT = 10;             // Speed increment
14
15         // Variables
16         int kph;           // To hold the speed in kph
17         double mph;        // To hold the speed in mph
18
19         // Display the table headings.
20         System.out.println("KPH\t\tMPH");
21         System.out.println("-----");
22
23         // Display the speeds.
```

```

24     for (kph = STARTING_KPH; kph <= MAX_KPH; kph += INCREMENT)
25     {
26         // Calculate the mph.
27         mph = kph * 0.6214;
28
29         // Display the speeds in kph and mph.
30         System.out.printf("%d\t\t%.1f\n", kph, mph);
31     }
32 }
33 }

```

Program Output

KPH MPH

60 37.3

70 43.5

80 49.7

90 55.9

100 62.1

110 68.4

120 74.6

130 80.8



Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 4.7 Name the three expressions that appear inside the parentheses in the `for` loop's header.

2. 4.8 You want to write a `for` loop that displays “I love to program” 50 times. Assume that you will use a control variable named `count`.
 1. What initialization expression will you use?
 2. What test expression will you use?
 3. What update expression will you use?
 4. Write the loop.
3. 4.9 What will the following program segments display?
 1.

```
for (int count = 0; count < 6; count++)
    System.out.println(count + count);
```
 2.

```
for (int value = -5; value < 5; value++)
    System.out.println(value);
```
 3.

```
int x;
for (x = 5; x <= 14; x += 3)
    System.out.println(x);
System.out.println(x);
```
4. 4.10 Write a `for` loop that displays your name 10 times.
5. 4.11 Write a `for` loop that displays all of the odd numbers, 1 through 49.
6. 4.12 Write a `for` loop that displays every fifth number, zero through 100.

4.6 Running Totals and Sentinel Values

CONCEPT: A running total is a sum of numbers that accumulates with each iteration of a loop. The variable used to keep the running total is called an accumulator. A sentinel is a value that signals when the end of a list of values has been reached.

Many programming tasks require you to calculate the total of a series of numbers. For example, suppose you are writing a program that calculates a business's total sales for a week. The program would read the sales for each day as input and calculate the total of those numbers.

Programs that calculate the total of a series of numbers typically use two elements:

- A loop that reads each number in the series.
- A variable that accumulates the total of the numbers as they are read.

The variable used to accumulate the total of the numbers is called an *accumulator*. It is often said that the loop keeps a *running total* because it accumulates the total as it reads each number in the series. [Figure 4-11](#) shows the general logic of a loop that calculates a running total.




Figure 4-11 Logic for calculating a running total

When the loop finishes, the accumulator will contain the total of the numbers that were read by the loop. Notice that the first step in the flowchart is to set the accumulator variable to 0. This is a critical step. Each time the loop reads a number, it adds it to the accumulator. If the accumulator starts with any value other than 0, it will not contain the correct total when the loop finishes.

Let's look at a program that calculates a running total. [Code Listing 4-10](#) calculates a company's total sales over a period of time by taking daily sales figures as input and calculating a running total of them as they are gathered. [Figure 4-12](#) shows an example of interaction with the program.

Code Listing 4-10 (TotalSales.java)

```
1  import javax.swing.JOptionPane;
2
3  /**
4   * This program calculates a running total.
5   */
6
7  public class TotalSales
8  {
9      public static void main(String[] args)
10     {
11         int days;           // The number of days
12         double sales;       // A day's sales figure
13         double totalSales;  // Accumulator
14         String input;       // To hold the user's input
15
16         // Get the number of days.
17         input = JOptionPane.showInputDialog("For how many days " +
18                                           "do you have sales figures?");
19         days = Integer.parseInt(input);
20
21         // Set the accumulator to 0.
22         totalSales = 0.0;
23
24         // Get the sales figures and calculate a running total.
25         for (int count = 1; count <= days; count++)
26         {
27             input = JOptionPane.showInputDialog("Enter the sales " +
28                                               "for day " + count + ": ");
29             sales = Double.parseDouble(input);
30             totalSales += sales; // Add sales to totalSales.
31         }
32
33         // Display the total sales.
34         JOptionPane.showMessageDialog(null,
35             String.format("The total sales are $%,.2f", totalSales));
36
37         System.exit(0);
38     }
39 }
```




Figure 4-12 Interaction with the TotalSales program

(Oracle Corporate Counsel)

[Figure 4-12 Full Alternative Text](#)

Let's take a closer look at this program. In lines 17 and 18 the user is asked to enter the number of days for which he or she has sales figures. The number is read from an input dialog box and assigned to the `days` variable. Then, in line 22 the `totalSales` variable is assigned 0.0. In general programming terms, the `totalSales` variable is referred to as an accumulator. An *accumulator* is a variable that is initialized with a starting value, which is usually zero, and then accumulates a sum of numbers by having the numbers added to it. As you will see, it is critical that the accumulator is set to zero before values are added to it.

Next, the `for` loop in lines 25 through 31 executes. During each iteration of the loop, the user enters the amount of sales for a specific day, which are assigned to the `sales` variable. This is done in lines 27 through 29. Then, in line 30 the contents of `sales` is added to the existing value in the `totalSales` variable. (Note that line 30 does not assign `sales` to `totalSales`, but adds `sales` to `totalSales`. Put another way, this line increases `totalSales` by the amount in `sales`.)

Because `totalSales` was initially assigned 0.0, after the first iteration of the loop, `totalSales` will be set to the same value as `sales`. After each subsequent iteration, `totalSales` will be increased by the amount in `sales`. After the loop has finished, `totalSales` will contain the total of all the daily sales figures entered. Now it should be clear why we assigned 0.0 to `totalSales` before the loop executed. If `totalSales` started at any other value, the total would be incorrect.

Using a Sentinel Value

The program in [Code Listing 4-10](#) requires the user to know in advance the number of days for which

he or she has sales figures. Sometimes the user has a very long list of input values, and doesn't know the exact number of items. In other cases, the user might be entering values from several lists and it is impractical to require that every item in every list is counted.

A technique that can be used in these situations is to ask the user to enter a sentinel value at the end of the list. A *sentinel value* is a special value that cannot be mistaken as a member of the list, and signals that there are no more values to be entered. When the user enters the sentinel value, the loop terminates.

The program in [Code Listing 4-11](#) shows an example. It calculates the total points earned by a soccer team over a series of games. It allows the user to enter the series of game points, and then -1 to signal the end of the list.

Code Listing 4-11 (SoccerPoints.java)

```
1  import java.util.Scanner;
2
3  /**
4   * This program calculates the total number of points a
5   * soccer team has earned over a series of games. The user
6   * enters a series of point values, then -1 when finished.
7   */
8
9  public class SoccerPoints
10 {
11     public static void main(String[] args)
12     {
13         int points;                // Game points
14         int totalPoints = 0;        // Accumulator initialized to 0
15
16         // Create a Scanner object for keyboard input.
17         Scanner keyboard = new Scanner(System.in);
18
19         // Display general instructions.
20         System.out.println("Enter the number of points your team");
21         System.out.println("has earned for each game this season.");
22         System.out.println("Enter -1 when finished.");
23         System.out.println();
24
25         // Get the first number of points.
26         System.out.print("Enter game points or -1 to end: ");
27         points = keyboard.nextInt();
28
29         // Accumulate the points until -1 is entered.
30         while (points != -1)
31         {
32             // Add points to totalPoints.
33             totalPoints += points;
34
35             // Get the next number of points.
36             System.out.print("Enter game points or -1 to end: ");
37             points = keyboard.nextInt();
```

```

38     }
39
40     // Display the total number of points.
41     System.out.println("The total points are " +
42         totalPoints);
43 }
44 }

```

Program Output with Example Input Shown in Bold

Enter the number of points your team
has earned for each game this season.
Enter -1 when finished.

```

Enter game points or -1 to end: 7 [Enter]
Enter game points or -1 to end: 9 [Enter]
Enter game points or -1 to end: 4 [Enter]
Enter game points or -1 to end: 6 [Enter]
Enter game points or -1 to end: 8 [Enter]
Enter game points or -1 to end: -1 [Enter]
The total points are 34

```

The value -1 was chosen for the sentinel because it is not possible for a team to score negative points. Notice that this program performs a priming read to get the first value. This makes it possible for the loop to terminate immediately if the user enters -1 as the first value. Also note that the sentinel value is not included in the running total.

Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 4.13 Write a `for` loop that repeats seven times, asking the user to enter a number. The loop should also calculate the sum of the numbers entered.
2. 4.14 In the following program segment, which variable is the loop control variable (also known as the counter variable) and which is the accumulator?

```

int a, x = 0, y = 0;
while (x < 10)
{
    a = x * 2;
    y += a;
    x++;
}
System.out.println("The sum is " + y);

```


3. 4.15 Why should you be careful when choosing a sentinel value?

4.7 Nested Loops

CONCEPT: A loop that is inside another loop is called a nested loop.

Nested loops are necessary when a task performs a repetitive operation and that task itself must be repeated. A clock is a good example of something that works like a nested loop. The program in [Code Listing 4-12](#) uses nested loops to simulate a clock.

Code Listing 4-12 (Clock.java)

```
1  /**
2      This program uses nested loops to simulate a clock.
3  */
4
5  public class Clock
6  {
7      public static void main(String[] args)
8      {
9          // Simulate the clock.
10         for (int hours = 1; hours <= 12; hours++)
11         {
12             for (int minutes = 0; minutes <= 59; minutes++)
13             {
14                 for (int seconds = 0; seconds <= 59; seconds++)
15                 {
16                     System.out.printf("%02d:%02d:%02d\n", hours, minutes, seconds);
17                 }
18             }
19         }
20     }
21 }
```

Program Output

```
01:00:00
01:00:01
01:00:02
01:00:03
(The loop continues to count ... )
12:59:57
12:59:58
12:59:59
```

The innermost loop (which begins at line 14) will iterate 60 times for each single iteration of the middle loop. The middle loop (which begins at line 12) will iterate 60 times for each single iteration of the outermost loop. When the outermost loop (which begins at line 10) has iterated 12 times, the

middle loop will have iterated 720 times and the innermost loop will have iterated 43,200 times.

The simulated clock example brings up a few points about nested loops:

- An inner loop goes through all of its iterations for each iteration of an outer loop.
- Inner loops complete their iterations before outer loops do.
- To get the total number of iterations of a nested loop, multiply the number of iterations of all the loops.

The program in [Code Listing 4-13](#) shows another example. It is a program that a teacher might use to get the average of each student's test scores. In line 22 the user enters the number of students, and in line 26 the user enters the number of test scores per student. The `for` loop that begins in line 29 iterates once for each student. The nested inner `for` loop, in lines 36 through 41, iterates once for each test score.

Code Listing 4-13 (TestAverage2.java)

```
1  import java.util.Scanner;
2
3  /**
4   * This program demonstrates a nested loop.
5   */
6
7  public class TestAverage2
8  {
9      public static void main(String [] args)
10     {
11         int numStudents,      // Number of students
12             numTests,         // Number of tests per student
13             score,             // Test score
14             total;             // Accumulator for test scores
15         double average;       // Average test score
16
17         // Create a Scanner object for keyboard input.
18         Scanner keyboard = new Scanner(System.in);
19
20         // Get the number of students.
21         System.out.print("How many students do you have? ");
22         numStudents = keyboard.nextInt();
23
24         // Get the number of test scores per student.
25         System.out.print("How many test scores per student? ");
26         numTests = keyboard.nextInt();
27
28         // Process all the students.
29         for (int student = 1; student <= numStudents; student++)
30         {
31             total = 0; // Set the accumulator to zero.
32
```

```

33 // Get the test scores for a student.
34 System.out.println("Student number " + student);
35 System.out.println("-----");
36 for (int test = 1; test <= numTests; test++)
37 {
38     System.out.print("Enter score " + test + ": ");
39     score = keyboard.nextInt();
40     total += score; // Add score to total.
41 }
42
43 // Calculate and display the average.
44 average = total / numTests;
45 System.out.printf("The average for student %d is %.1f.\n\n",
46                 student, average);
47 }
48 }
49 }

```

Program Output with Example Input Shown in Bold

How many students do you have? **3** [Enter]
 How many test scores per student? **3** [Enter]

Student number 1

 Enter score 1: **100** [Enter]
 Enter score 2: **95** [Enter]
 Enter score 3: **90** [Enter]
 The average for student number 1 is 95.0.

Student number 2

 Enter score 1: **80** [Enter]
 Enter score 2: **81** [Enter]
 Enter score 3: **82** [Enter]
 The average for student number 2 is 81.0.

Student number 3

 Enter score 1: **75** [Enter]
 Enter score 2: **85** [Enter]
 Enter score 3: **80** [Enter]
 The average for student number 3 is 80.0.

In the Spotlight:



Using Nested Loops to Print Patterns

One interesting way to learn about nested loops is to use them to display patterns on the screen. Let's look at a simple example. Suppose we want to print asterisks on the screen in the following rectangular pattern:

```
*****
*****
*****
*****
*****
*****
*****
*****
```

If you think of this pattern as having rows and columns, you can see that it has eight rows, and each row has six columns. The following code can be used to display one row of asterisks:

```
final int COLS = 6;
for (int col = 0; col < COLS; col++)
{
    System.out.print("*");
}
```

If we run this code in a program, it will produce the following output:

```
*****
```

To complete the entire pattern, we need to execute this loop eight times. We can place the loop inside another loop that iterates eight times, as shown here:

```
1      final int COLS = 6;
2      final int ROWS = 8;
3      for (int row = 0; row < ROWS; row++)
4      {
5          for (int col = 0; col < COLS; col++)
6          {
7              System.out.print("*");
8          }
9          System.out.println();
10     }
```

The outer loop iterates eight times. Each time it iterates, the inner loop iterates six times. (Notice that in line 9, after each row has been printed, we call the `System.out.println()` method. We have to do that to advance the screen cursor to the next line at the end of each row. Without that statement, all the asterisks will be printed in one long row on the screen.)

We could easily write a program that prompts the user for the number of rows and columns, as shown in [Code Listing 4-14](#).

Code Listing 4-14 (RectangularPattern.java)

```
1  import java.util.Scanner;
2
3  /**
4   * This program displays a rectangular pattern
5   * of asterisks.
6   */
7
8  public class RectangularPattern
9  {
10     public static void main(String[] args)
11     {
12         int rows, cols;
13
14         // Create a Scanner object for keyboard input.
15         Scanner keyboard = new Scanner(System.in);
16
17         // Get the number of rows and columns.
18         System.out.print("How many rows? ");
19         rows = keyboard.nextInt();
20         System.out.print("How many columns? ");
21         cols = keyboard.nextInt();
22
23         for (int r = 0; r < rows; r++)
24         {
25             for (int c = 0; c < cols; c++)
26             {
27                 System.out.print("*");
28             }
29             System.out.println();
30         }
31     }
32 }
```

Program Output with Example Input Shown in Bold

```
How many rows? 5 [Enter]
How many columns? 10 [Enter]
*****
*****
*****
*****
*****
```

Let's look at another example. Suppose you want to print asterisks in a pattern that looks like the following triangle:

*

```

*
**
***
****
*****
*****
*****
*****
*****

```

Once again, think of the pattern as being arranged in rows and columns. The pattern has a total of eight rows. In the first row, there is one column. In the second row, there are two columns. In the third row, there are three columns. This continues to the eighth row, which has eight columns. [Code Listing 4-15](#) shows the program that produces this pattern.

Code Listing 4-15 (TrianglePattern.java)

```

1  import java.util.Scanner;
2
3  /**
4   * This program displays a triangle pattern.
5   */
6
7  public class TrianglePattern
8  {
9      public static void main(String[] args)
10     {
11         final int BASE_SIZE = 8;
12
13         for (int r = 0; r < BASE_SIZE; r++)
14         {
15             for (int c = 0; c < (r + 1); c++)
16             {
17                 System.out.print("*");
18             }
19             System.out.println();
20         }
21     }
22 }

```

Program Output

```

*
**
***
****
*****
*****
*****
*****
*****

```

The outer loop (which begins in line 13) will iterate eight times. As the loop iterates, the variable `r`

will be assigned the values 0 through 7.

For each iteration of the outer loop, the inner loop will iterate $r + 1$ times. So,

- During the outer loop's first iteration, the variable r is assigned 0. The inner loop iterates one time, printing one asterisk.
- During the outer loop's second iteration, the variable r is assigned 1. The inner loop iterates two times, printing two asterisks.
- During the outer loop's third iteration, the variable r is assigned 2. The inner loop iterates three times, printing three asterisks.
- And so forth.

Let's look at another example. Suppose you want to display the following stair-step pattern:

```
#  
 #  
  #  
   #  
    #  
     #
```

The pattern has six rows. In general, we can describe each row as having some number of spaces followed by a # character. Here's a row-by-row description:

First row: 0 spaces followed by a # character.

Second row: 1 space followed by a # character.

Third row: 2 spaces followed by a # character.

Fourth row: 3 spaces followed by a # character.

Fifth row: 4 spaces followed by a # character.

Sixth row: 5 spaces followed by a # character.

To display this pattern, we can write code containing a pair of nested loops that work in the following manner:

- The outer loop will iterate six times. Each iteration will perform the following:
- The inner loop will display the correct number of spaces, side by side.
- Then, a # character will be displayed.

[Code Listing 4-16](#) shows the Java code.

Code Listing 4-16 (StairStepPattern.java)

```

1  import java.util.Scanner;
2
3  /**
4   This program displays a stairstep pattern.
5   */
6
7  public class StairStepPattern
8  {
9      public static void main(String[] args)
10     {
11         final int NUM_STEPS = 6;
12
13         for (int r = 0; r < NUM_STEPS; r++)
14         {
15             for (int c = 0; c < r; c++)
16             {
17                 System.out.print(" ");
18             }
19             System.out.println("#");
20         }
21     }
22 }
```

Program Output

```

#
 #
  #
   #
    #
     #
```

The outer loop (which begins in line 13) will iterate six times. As the loop iterates, the variable `r` will be assigned the values 0 through 5.

For each iteration of the outer loop, the inner loop will iterate `r` times. So,

- During the outer loop's first iteration, the variable `r` is assigned 0. The inner loop will not execute at this time.

- During the outer loop's second iteration, the variable r is assigned 1. The inner loop iterates one time, printing one space.
- During the outer loop's third iteration, the variable r is assigned 2. The inner loop iterates two times, printing two spaces.
- And so forth.

4.8 The break and continue Statements (Optional)

CONCEPT: The `break` statement causes a loop to terminate early. The `continue` statement causes a loop to stop its current iteration and begin the next one.

The `break` statement, which was used with the `switch` statement in [Chapter 3](#), can also be placed inside a loop. When it is encountered, the loop stops and the program jumps to the statement immediately following the loop. Although it is perfectly acceptable to use the `break` statement in a `switch` statement, it is considered taboo to use it in a loop. This is because it bypasses the normal condition that is required to terminate the loop, and it makes code difficult to understand and debug. For this reason, you should avoid using the `break` statement in a loop when possible.

The `continue` statement causes the current iteration of a loop to end immediately. When `continue` is encountered, all the statements in the body of the loop that appear after it are ignored, and the loop prepares for the next iteration. In a `while` loop, this means the program jumps to the `boolean` expression at the top of the loop. As usual, if the expression is still `true`, the next iteration begins. In a `do-while` loop, the program jumps to the `boolean` expression at the bottom of the loop, which determines whether the next iteration will begin. In a `for` loop, `continue` causes the update expression to be executed, and then the test expression is evaluated.

The `continue` statement should also be avoided. Like the `break` statement, it bypasses the loop's logic and makes the code difficult to understand and debug.

4.9 Deciding Which Loop to Use

CONCEPT: Although most repetitive algorithms can be written with any of the three types of loops, each works best in different situations.

Each of Java's three loops is ideal to use in different situations. The following is a short summary of when each loop should be used:

- The `while` loop. The `while` loop is a pretest loop. It is ideal in situations where you do not want the loop to iterate if the condition is `false` from the beginning. It is also ideal if you want to use a sentinel value to terminate the loop.
- The `do-while` loop. The `do-while` loop is a posttest loop. It is ideal in situations where you always want the loop to iterate at least once.
- The `for` loop. The `for` loop is a pretest loop that has built-in expressions for initializing, testing, and updating. These expressions make it very convenient to use a loop control variable as a counter. The `for` loop is ideal in situations where the exact number of iterations is known.

4.10 Introduction to File Input and Output

CONCEPT: The Java API provides several classes that you can use for writing data to a file and reading data from a file. To write data to a file, you can use the `PrintWriter` class and, optionally, the `FileWriter` class. To read data from a file, you can use the `Scanner` class and the `File` class.

The programs you have written so far require you to reenter data each time the program runs. This is because the data stored in variables and objects in RAM disappears once the program stops running. To retain data between the times it runs, a program must have a way of saving the data.

Data may be saved in a *file*, which is usually stored on a computer's disk. Once the data is saved in a file, it will remain there after the program stops running. The data can then be retrieved and used at a later time. In general, there are three steps that are taken when a file is used by a program:

1. The file must be *opened*. When the file is opened, a connection is created between the file and the program.
2. Data is then written to the file or read from the file.
3. When the program is finished using the file, the file must be *closed*.

In this section we will discuss how to write Java programs that write data to files and read data from files. The terms *input file* and *output file* are commonly used. An *input file* is a file that a program reads data from. It is called an input file because the data stored in it serves as input to the program. An *output file* is a file that a program writes data to. It is called an output file because the program stores output in the file.

In general, there are two types of files: text and binary. A *text file* contains data that has been encoded as text, using a scheme such as Unicode. Even if the file contains numbers, those numbers are stored in the file as a series of characters. As a result, the file may be opened and viewed in a text editor such as Notepad. A *binary file* contains data that has not been converted to text. As a consequence, you cannot view the contents of a binary file with a text editor. In this chapter, we will discuss how to work with text files. Binary files are discussed in [Chapter 11](#).

The Java API provides a number of classes that you will use to work with files. To use these classes, you will place the following `import` statement near the top of your program:

```
import java.io.*;
```

Using the `PrintWriter` Class to Write Data to a File

To write data to a file you will create an instance of the `PrintWriter` class. The `PrintWriter` class allows you to open a file for writing. It also allows you to write data to the file using the same `print` and `println` methods that you have been using to display data on the screen. You pass the name of the file that you wish to open, as a string, to the `PrintWriter` class's constructor. For example, the following statement creates a `PrintWriter` object and passes the file name *StudentData.txt* to the constructor.

```
PrintWriter outputFile = new PrintWriter("StudentData.txt");
```

This statement will create an empty file named *StudentData.txt* and establish a connection between it and the `PrintWriter` object that is referenced by `outputFile`. The file will be created in the current directory or folder.

You may also pass a reference to a `String` object as an argument to the `PrintWriter` constructor. For example, in the following code the user specifies the name of the file.

```
String filename;  
filename = JOptionPane.showInputDialog("Enter the filename.");  
PrintWriter outputFile = new PrintWriter(filename);
```



If the file that you are opening with the `PrintWriter` object already exists, it will be erased and an empty file by the same name will be created.

Once you have created an instance of the `PrintWriter` class and opened a file, you can write data to the file using the `print` and `println` methods. You already know how to use `print` and `println` with `System.out` to display data on the screen. They are used the same way with a `PrintWriter` object to write data to a file. For example, assuming that `outputFile` references a `PrintWriter` object, the following statement writes the string "Jim" to the file:

```
outputFile.println("Jim");
```

When the program is finished writing data to the file, it must close the file. To close the file use the `PrintWriter` class's `close` method. Here is an example of the method's use:

```
outputFile.close();
```

Your application should always close files when finished with them. This is because the system creates one or more buffers when a file is opened. A *buffer* is a small “holding section” of memory. When a program writes data to a file, that data is first written to the buffer. When the buffer is filled, all the information stored there is written to the file. This technique increases the system's performance because writing data to memory is faster than writing it to a disk. The `close` method writes any unsaved data remaining in the file buffer.

Once a file is closed, the connection between it and the `PrintWriter` object is removed. In order to

perform further operations on the file, it must be opened again.

More about the `PrintWriter` Class's `println` Method

The `PrintWriter` class's `println` method writes a line of data to a file. For example, assume an application creates a file and writes three students' first names and their test scores to the file with the following code:

```
PrintWriter outputFile = new PrintWriter("StudentData.txt");
outputFile.println("Jim");
outputFile.println(95);
outputFile.println("Karen");
outputFile.println(98);
outputFile.println("Bob");
outputFile.println(82);
outputFile.close();
```

The `println` method writes data to the file and then writes a newline character immediately after the data. You can visualize the data written to the file in the following manner:

`Jim<newline>95<newline>Karen<newline>98<newline>Bob<newline>82<newline>`

The newline characters are represented here as `<newline>`. You do not actually see the newline characters, but when the file is opened in a text editor such as Notepad, its contents will appear as shown in [Figure 4-13](#). As you can see from the figure, each newline character causes the data that follows it to be displayed on a new line.




Figure 4-13 File contents displayed in Notepad

(Microsoft Corporation)

In addition to separating the contents of a file into lines, the newline character also serves as a

delimiter. A *delimiter* is an item that separates other items. When you write data to a file using the `println` method, newline characters will separate the individual items of data. Later you will see that the individual items of data in a file must be separated in order for them to be read from the file.

The `PrintWriter` Class's `print` Method

The `print` method is used to write an item of data to a file without writing the newline character. For example, look at the following code:

```
String name = "Jeffrey Smith";
String phone = "555-7864";
int idNumber = 47895;
PrintWriter outputFile = new PrintWriter("PersonalData.txt");
outputFile.print(name + " ");
outputFile.print(phone + " ");
outputFile.println(idNumber);
outputFile.close();
```

This code uses the `print` method to write the contents of the `name` object to the file, followed by a space (" "). Then it uses the `print` method to write the contents of the `phone` object to the file, followed by a space. Then it uses the `println` method to write the contents of the `idNumber` variable, followed by a newline character. [Figure 4-14](#) shows the contents of the file displayed in Notepad.




Figure 4-14 Contents of file displayed in Notepad

(Microsoft Corporation)

Adding a `throws` Clause to the Method Header

When an unexpected event occurs in a Java program, it is said that the program throws an exception. For now, you can think of an *exception* as a signal indicating that the program cannot continue until the unexpected event has been dealt with. For example, suppose you create a `PrintWriter` object

and pass the name of a file to its constructor. The `PrintWriter` object attempts to create the file, but unexpectedly, the disk is full and the file cannot be created. Obviously the program cannot continue until this situation has been dealt with, so an exception is thrown, which causes the program to suspend normal execution.

When an exception is thrown, the method that is executing must either deal with the exception or throw it again. If the `main` method throws an exception, the program halts and an error message is displayed. Because `PrintWriter` objects are capable of throwing exceptions, we must either write code that deals with the possible exceptions, or allow our methods to rethrow the exceptions when they occur. In [Chapter 11](#) you will learn all about exceptions and how to respond to them, but for now, we will simply allow our methods to rethrow any exceptions that might occur.

To allow a method to rethrow an exception that has not been dealt with, you simply write a `throws` clause in the method header. The `throws` clause must indicate the type of exception that might be thrown. The following is an example:

```
public static void main(String[] args) throws IOException
```

This header indicates that the `main` method is capable of throwing an exception of the `IOException` type. This is the type of exception that `PrintWriter` objects are capable of throwing. So, any method that uses `PrintWriter` objects and does not respond to their exceptions must have this `throws` clause listed in its header.

In addition, any method that calls a method that uses a `PrintWriter` object should have a `throws IOException` clause in its header. For example, suppose the `main` method does not perform any file operations, but calls a method named `buildFile` that opens a file and writes data to it. Both the `buildFile` and `main` methods should have a `throws IOException` clause in their headers. Otherwise a compiler error will occur.

An Example Program

Let's look at an example program that writes data to a file. The program in [Code Listing 4-17](#) writes the names of your friends to a file.

Code Listing 4-17 (`FileWriteDemo.java`)

```
1  import java.util.Scanner;           // Needed for Scanner class
2  import java.io.*;                  // Needed for File I/O classes
3
4  /**
5   * This program writes data to a file.
6   */
7
8  public class FileWriteDemo
9  {
```

```

10 public static void main(String[] args) throws IOException
11 {
12     String filename;           // File name
13     String friendName;         // Friend's name
14     int numFriends;            // Number of friends
15
16     // Create a Scanner object for keyboard input.
17     Scanner keyboard = new Scanner(System.in);
18
19     // Get the number of friends.
20     System.out.print("How many friends do you have? ");
21     numFriends = keyboard.nextInt();
22
23     // Consume the remaining newline character.
24     keyboard.nextLine();
25
26     // Get the filename.
27     System.out.print("Enter the filename: ");
28     filename = keyboard.nextLine();
29
30     // Open the file.
31     PrintWriter outputFile = new PrintWriter(filename);
32
33     // Get data and write it to the file.
34     for (int i = 1; i <= numFriends; i++)
35     {
36         // Get the name of a friend.
37         System.out.print("Enter the name of friend " +
38             "number " + i + ": ");
39         friendName = keyboard.nextLine();
40
41         // Write the name to the file.
42         outputFile.println(friendName);
43     }
44
45     // Close the file.
46     outputFile.close();
47     System.out.println("Data written to the file.");
48 }
49 }

```

Program Output with Example Input Shown in Bold

```

How many friends do you have? 5 [Enter]
Enter the filename: MyFriends.txt [Enter]
Enter the name of friend number 1: Joe [Enter]
Enter the name of friend number 2: Rose [Enter]
Enter the name of friend number 3: Greg [Enter]
Enter the name of friend number 4: Kirk [Enter]
Enter the name of friend number 5: Renee [Enter]
Data written to the file.

```

The `import` statement in line 2 is necessary because this program uses the `PrintWriter` class. In addition, the `main` method header, in line 10, has a `throws IOException` clause because objects of the `PrintWriter` class can potentially throw an `IOException`.

This program asks the user to enter the number of friends he or she has (in lines 20 and 21), then a name for the file that will be created (in lines 27 and 28). The `filename` variable references the name of the file, and is used in the following statement, in line 31:

```
PrintWriter outputFile = new PrintWriter(filename);
```

This statement opens the file and creates a `PrintWriter` object that can be used to write data to the file. The `for` loop in lines 34 through 43 performs an iteration for each friend that the user has, each time asking for the name of a friend. The user's input is referenced by the `friendName` variable. Once the name is entered, it is written to the file with the following statement, which appears in line 42:

```
outputFile.println(friendName);
```

After the loop finishes, the file is closed in line 46. After the program is executed with the input shown in the example run, the file *MyFriends.txt* will be created. If we open the file in Notepad, we will see its contents as shown in [Figure 4-15](#).




Figure 4-15 Contents of the file displayed in Notepad

(Microsoft Corporation)

Review

Before moving on, let's review the basic steps necessary when writing a program that writes data to a file:

1. You need the `import java.io.*;` statement in the top section of your program.

2. Because we have not yet learned how to respond to exceptions, any method that uses a `PrintWriter` object must have a `throws IOException` clause in its header.
3. You create a `PrintWriter` object and pass the name of the file as a string to the constructor.
4. You use the `PrintWriter` class's `print` and `println` methods to write data to the file.
5. When finished writing to the file, you use the `PrintWriter` class's `close` method to close the file.

Appending Data to a File

When you pass the name of a file to the `PrintWriter` constructor, and the file already exists, it will be erased and a new empty file with the same name will be created. Sometimes, however, you want to preserve an existing file and append new data to its current contents. Appending to a file means writing new data to the end of the data that already exists in the file.

To append data to an existing file, you first create an instance of the `FileWriter` class. You pass two arguments to the `FileWriter` constructor: a string containing the name of the file, and the `boolean` value `true`. Here is an example:

```
FileWriter fwriter = new FileWriter("MyFriends.txt", true);
```

This statement creates a `FileWriter` object and opens the file *MyFriends.txt* for writing. Any data written to the file will be appended to the file's existing contents. (If the file does not exist, it will be created.)

You still need to create a `PrintWriter` object so you can use the `print` and `println` methods to write data to the file. When you create the `PrintWriter` object, you pass a reference to the `FileWriter` object as an argument to the `PrintWriter` constructor. For example, look at the following code:

```
FileWriter fwriter = new FileWriter("MyFriends.txt", true);  
PrintWriter outputFile = new PrintWriter(fwriter);
```

This creates a `PrintWriter` object that can be used to write data to the file *MyFriends.txt*. Any data that is written to the file will be appended to the file's existing contents. For example, assume the file *MyFriends.txt* exists and contains the following data:

```
Joe  
Rose  
Greg  
Kirk  
Renee
```

The following code opens the file and appends additional data to its existing contents:

```
FileWriter fwriter = new FileWriter("MyFriends.txt", true);
PrintWriter outputFile = new PrintWriter(fwriter);
outputFile.println("Bill");
outputFile.println("Steven");
outputFile.println("Sharon");
outputFile.close();
```

After this code executes, the *MyFriends.txt* file will contain the following data:

```
Joe
Rose
Greg
Kirk
Renee
Bill
Steven
Sharon
```

NOTE:

The `FileWriter` class also throws an `IOException` if the file cannot be opened for any reason.

Specifying the File Location

When you open a file you may specify its path along with its filename. On a Windows computer, paths contain backslash characters. Remember that when a single backslash character appears in a string literal, it marks the beginning of an escape sequence such as `"\n"`. Two backslash characters in a string literal represent a single backslash. So, when you provide a path in a string literal, and the path contains backslash characters, you must use two backslash characters in the place of each single backslash character.

For example, the path `"E:\\Names.txt"` specifies that *Names.txt* is in the root folder of drive E:, and the path `"C:\\MyData\\Data.txt"` specifies that *Data.txt* is in the *MyData* folder on drive C:. In the following statement, the file *Pricelist.txt* is created in the root folder of drive H:

```
PrintWriter outputFile = new PrintWriter("H:\\PriceList.txt");
```

You only need to use double backslashes if the file's path is in a string literal. If your program asks the user to enter a path into a `String` object, which is then passed to the `PrintWriter` or `FileWriter` constructor, the user does not have to enter double backslashes.

TIP:

Java allows you to substitute forward slashes for backslashes in a Windows path. For example, the

path `"C:\\MyData\\Data.txt"` could be written as `"C:/MyData/Data.txt"`. This eliminates the need to use double backslashes.

On a UNIX or Linux computer, you can provide a path without any modifications. Here is an example:

```
PrintWriter outputFile = new PrintWriter("/home/rharrison/names.txt");
```

Reading Data from a File

In [Chapter 2](#) you learned how to use the `Scanner` class to read input from the keyboard. To read keyboard input, recall that we create a `Scanner` object, passing `System.in` to the `Scanner` class constructor. Here is an example:

```
Scanner keyboard = new Scanner(System.in);
```

Recall that the `System.in` object represents the keyboard. Passing `System.in` as an argument to the `Scanner` constructor specifies that the keyboard is the `Scanner` object's source of input.

You can also use the `Scanner` class to read input from a file. Instead of passing `System.in` to the `Scanner` class constructor, you pass a reference to a `File` object. Here is an example:

```
File myFile = new File("Customers.txt");
Scanner inputFile = new Scanner(myFile);
```

The first statement creates an instance of the `File` class. The `File` class is in the Java API, and is used to represent a file. Notice that we have passed the string `"Customers.txt"` to the constructor. This creates a `File` object that represents the file *Customers.txt*.

In the second statement we pass a reference to this `File` object as an argument to the `Scanner` class constructor. This creates a `Scanner` object that uses the file *Customers.txt* as its source of input. You can then use the same `Scanner` class methods that you learned about in [Chapter 2](#) to read items from the file. (See Table 2-17 for a list of commonly used methods.)

When you are finished reading from the file, you use the `Scanner` class's `close` method to close the file. For example, assuming the variable `inputFile` references a `Scanner` object, the following statement closes the file that is the object's source of input:

```
inputFile.close();
```

Reading Lines from a File with the `nextLine` Method

The `Scanner` class's `nextLine` method reads a line of input, and returns the line as a `String`. The

program in [Code Listing 4-18](#) demonstrates how the `nextLine` method can be used to read a line from a file. This program asks the user to enter a filename. It then displays the first line in the file on the screen.

Code Listing 4-18 (ReadFirstLine.java)

```
1  import java.util.Scanner; // Needed for Scanner class
2  import java.io.*;         // Needed for File and IOException
3
4  /**
5   * This program reads the first line from a file.
6   */
7
8  public class ReadFirstLine
9  {
10     public static void main(String[] args) throws IOException
11     {
12         // Create a Scanner object for keyboard input.
13         Scanner keyboard = new Scanner(System.in);
14
15         // Get the file name.
16         System.out.print("Enter the name of a file: ");
17         String filename = keyboard.nextLine();
18
19         // Open the file.
20         File file = new File(filename);
21         Scanner inputFile = new Scanner(file);
22
23         // Read the first line from the file.
24         String line = inputFile.nextLine();
25
26         // Display the line.
27         System.out.println("The first line in the file is:");
28         System.out.println(line);
29
30         // Close the file.
31         inputFile.close();
32     }
33 }
```

Program Output with Example Input Shown in Bold

```
Enter the name of a file: MyFriends.txt [Enter]
The first line in the file is:
Joe
```

This program gets the name of a file from the user in line 17. A `File` object is created in line 20 to represent the file, and a `Scanner` object is created in line 21 to read data from the file. Line 24 reads

a line from the file. After this statement executes, the `line` variable references a `String` object holding the line that was read from the file. The line is displayed on the screen in line 28, and the file is closed in line 31.

It's worth pointing out that this program creates two separate `Scanner` objects. The `Scanner` object that is created in line 13 reads data from the keyboard, and the `Scanner` object that is created in line 21 reads data from a file.

When a file is opened for reading, a special value known as a *read position* is internally maintained for that file. A file's read position marks the location of the next item that will be read from the file. When a file is opened, its read position is set to the first item in the file. When the item is read, the read position is advanced to the next item in the file. As subsequent items are read, the internal read position advances through the file. For example, consider the file *Quotation.txt*, shown in [Figure 4-16](#). As you can see from the figure, the file has three lines.




Figure 4-16 File with three lines

(Microsoft Corporation)

You can visualize that the data is stored in the file in the following manner:

```
Imagination is more<newline>important than knowledge.<newline>
Albert Einstein<newline>
```

Suppose a program opens the file with the following code:

```
File file = new File("Quotation.txt");
Scanner inputFile = new Scanner(file);
```

When this code opens the file, its read position is at the beginning of the first line, as illustrated in [Figure 4-17](#).

Read position → Imagination is more
important than knowledge.
Albert Einstein

Figure 4-17 Initial read position

Now, suppose the program uses the following statement to read a line from the file:

```
String str = inputFile.nextLine();
```

This statement will read a line from the file, beginning at the current read position. After the statement executes, the object referenced by `str` will contain the string “Imagination is more”. The file’s read position will be advanced to the next line, as illustrated in [Figure 4-18](#).

Read position → Imagination is more
important than knowledge.
Albert Einstein

Figure 4-18 Read position after first line is read

If the `nextLine` method is called again, the second line will be read from the file and the file’s read position will be advanced to the third line. After all the lines have been read, the read position will be at the end of the file.



NOTE:

The string that is returned from the `nextLine` method will not contain the newline character.

Adding a throws Clause to the Method Header

When you pass a `File` object reference to the `Scanner` class constructor, the constructor will throw an exception of the `IOException` type if the specified file is not found. So, you will need to write a `throws IOException` clause in the header of any method that passes a `File` object reference to the `Scanner` class constructor.

Detecting the End of a File

Quite often a program must read the contents of a file without knowing the number of items that are stored in the file. For example, the *MyFriends.txt* file that was created by the program in [Code](#)

[Listing 4-17](#) can have any number of names stored in it. This is because the program asks the user for the number of friends that he or she has. If the user enters 5 for the number of friends, the program creates a file with five names in it. If the user enters 100, the program creates a file with 100 names in it.

The `Scanner` class has a method named `hasNext` that can be used to determine whether the file has more data that can be read. You call the `hasNext` method before you call any other methods to read from the file. If there is more data that can be read from the file, the `hasNext` method returns `true`. If the end of the file has been reached and there is no more data to read, the `hasNext` method returns `false`.

[Code Listing 4-19](#) shows an example. The program reads the file containing the names of your friends, which was created by the program in [Code Listing 4-17](#).

Code Listing 4-19 (FileReadDemo.java)

```
1  import java.util.Scanner;    // Needed for the Scanner class
2  import java.io.*;           // Needed for the File and IOException
3
4  /**
5   * This program reads data from a file.
6   */
7
8  public class FileReadDemo
9  {
10     public static void main(String[] args) throws IOException
11     {
12         // Create a Scanner object for keyboard input.
13         Scanner keyboard = new Scanner(System.in);
14
15         // Get the filename.
16         System.out.print("Enter the filename: ");
17         String filename = keyboard.nextLine();
18
19         // Open the file.
20         File file = new File(filename);
21         Scanner inputFile = new Scanner(file);
22
23         // Read lines from the file until no more are left.
24         while (inputFile.hasNext())
25         {
26             // Read the next name.
27             String friendName = inputFile.nextLine();
28
29             // Display the last name read.
30             System.out.println(friendName);
31         }
32
33         // Close the file.
34         inputFile.close();
35     }
36 }
```

Program Output with Example Input Shown in Bold

Enter the filename: **MyFriends.txt** [Enter]

Joe
Rose
Greg
Kirk
Renee

The file is opened and a `Scanner` object to read it is created in line 21. The loop in lines 24 through 31 reads all of the lines from the file and displays them. In line 24 the loop calls the `Scanner` object's `hasNext` method. If the method returns `true`, then the file has more data to read. In that case, the next line is read from the file in line 27, and is displayed in line 30. The loop repeats until the `hasNext` method returns `false` in line 24. [Figure 4-19](#) shows the logic of reading a file until the end is reached.




Figure 4-19 Logic of reading a file until the end is reached

[Figure 4-19 Full Alternative Text](#)

Reading Primitive Values from a File

Recall from [Chapter 2](#) that the `Scanner` class provides methods for reading primitive values. These methods are named `nextByte`, `nextDouble`, `nextFloat`, `nextInt`, `nextLine`, `nextLong`, and `nextShort`. Table 2-17 gives more information on each of these methods, which can be used to read primitive values from a file.

The program in [Code Listing 4-20](#) demonstrates how the `nextDouble` method can be used to read floating-point values from a file. The program reads the contents of a file named *Numbers.txt*. The contents of the *Numbers.txt* file are shown in [Figure 4-20](#). As you can see, the file contains a series of floating-point numbers. The program reads all of the numbers from the file and calculates their total.




Figure 4-20 Contents of *Numbers.txt*

(Microsoft Corporation)

Code Listing 4-20 (`FileSum.java`)

```
1  import java.util.Scanner;
2  import java.io.*;
3
4  /**
5   * This program reads a series of numbers from a file and
6   * accumulates their sum.
7   */
8
9  public class FileSum
10 {
11     public static void main(String[] args) throws IOException
12     {
13         double sum = 0.0; // Accumulator, initialized to 0
14
15         // Open the file for reading.
16         File file = new File("Numbers.txt");
17         Scanner inputFile = new Scanner(file);
```

```

18
19 // Read all of the values from the file
20 // and calculate their total.
21 while (inputFile.hasNext())
22 {
23     // Read a value from the file.
24     double number = inputFile.nextDouble();
25
26     // Add the number to sum.
27     sum = sum + number;
28 }
29
30 // Close the file.
31 inputFile.close();
32
33 // Display the sum of the numbers.
34 System.out.println("The sum of the numbers in " +
35                    "Numbers.txt is " + sum);
36 }
37 }

```

Program Output

The sum of the numbers in Numbers.txt is 41.4

Review

Let's quickly review the steps necessary when writing a program that reads data from a file:

1. You will need the `import java.util.Scanner;` statement in the top section of your program, so you can use the `Scanner` class. You will also need the `import java.io.*;` statement in the top section of your program. This is required by the `File` class.
2. Because we have not yet learned how to respond to exceptions, any method that uses a `Scanner` object to open a file must have a `throws IOException` clause in its header.
3. You create a `File` object and pass the name of the file as a string to the constructor.
4. You create a `Scanner` object and pass a reference to the `File` object as an argument to the constructor.
5. You use the `Scanner` class's `nextLine` method to read a line from the file. The method returns the line of data as a string. To read primitive values, use methods such as `nextInt`, `nextDouble`, and so forth.
6. Call the `Scanner` class's `hasNext` method to determine whether there is more data to read from the file. If the method returns `true`, then there is more data to read. If the method returns `false`,

you have reached the end of the file.

7. When finished writing to the file, you use the `Scanner` class's `close` method to close the file.

Checking for a File's Existence

It's usually a good idea to make sure that a file exists before you try to open it for input. If you attempt to open a file for input, and the file does not exist, the program will throw an exception and halt. For example, the program you saw in [Code Listing 4-20](#) will throw an exception at line 17 if the file *Numbers.txt* does not exist. Here is an example of the error message that will be displayed when this happens:

```
Exception in thread "main" java.io.FileNotFoundException: Numbers.txt (The
system cannot find the file specified)
    at java.io.FileInputStream.open(Native Method)
    at java.io.FileInputStream.<init>(FileInputStream.java:106)
    at java.util.Scanner.<init>(Scanner.java:636)
    at FileSum.main(FileSum.java:17)
```

Rather than allowing the exception to be thrown and permitting this cryptic error message to be displayed, your program can check for the file's existence before it attempts to open the file. If the file does not exist, the program can display a more user-friendly error message and gracefully shut down.

After you create a `File` object representing the file that you want to open, you can use the `File` class's `exists` method to determine whether the file exists. The method returns `true` if the file exists, or `false` if the file does not exist. [Code Listing 4-21](#) shows how to use the method. This is a modification of the *FileSum* program in [Code Listing 4-20](#). This version of the program checks for the existence of the file *Numbers.txt* before it attempts to open it.

Code Listing 4-21 (*FileSum2.java*)

```
1  import java.util.Scanner;
2  import java.io.*;
3
4  /**
5   * This version of the program confirms that the
6   * Numbers.txt file exists before opening it.
7   */
8
9  public class FileSum2
10 {
11     public static void main(String[] args) throws IOException
12     {
13         double sum = 0.0; // Accumulator, initialized to 0
14
15         // Make sure the file exists.
16         File file = new File("Numbers.txt");
17         if (!file.exists())
```

```

18     {
19         System.out.println("The file Numbers.txt is not found.");
20         System.exit(0);
21     }
22
23     // Open the file for reading.
24     Scanner inputFile = new Scanner(file);
25
26     // Read all of the values from the file
27     // and calculate their total.
28     while (inputFile.hasNext())
29     {
30         // Read a value from the file.
31         double number = inputFile.nextDouble();
32
33         // Add the number to sum.
34         sum = sum + number;
35     }
36
37     // Close the file.
38     inputFile.close();
39
40     // Display the sum of the numbers.
41     System.out.println("The sum of the numbers in " +
42                       "Numbers.txt is " + sum);
43 }
44 }

```

Program Output (Assuming Numbers.txt Does Not Exist)

The file Numbers.txt is not found.

In line 16 the program creates a `File` object to represent the *Numbers.txt* file. In line 17, the `if` statement calls the `file.exists()` method. Notice the use of the `!` operator. If the method returns `false`, indicating that the file does not exist, the code in lines 19 and 20 executes. Line 19 displays an error message, and line 20 calls the `System.exit(0)` method to shut the program down.

The previous example shows you how to make sure that a file exists before trying to open it for input. But, when you are opening a file for output, sometimes you want to make sure the file does *not* exist. When you use a `PrintWriter` object to open a file, the file will be erased if it already exists. If you do not want to erase the existing file, you have to check for its existence before creating the `PrintWriter` object. [Code Listing 4-22](#) shows you how to use the `File` class's `exists` method in this type of situation. This is a modification of the program you saw in [Code Listing 4-17](#).

Code Listing 4-22 (FileWriteDemo2.java)

```

1  import java.util.Scanner;           // Needed for Scanner class

```

```

2 import java.io.*; // Needed for File and IOException
3
4 /**
5  This program writes data to a file. It makes sure the
6  specified file does not exist before opening it.
7  */
8
9 public class FileWriteDemo2
10 {
11     public static void main(String[] args) throws IOException
12     {
13         String filename; // Filename
14         String friendName; // Friend's name
15         int numFriends; // Number of friends
16
17         // Create a Scanner object for keyboard input.
18         Scanner keyboard = new Scanner(System.in);
19
20         // Get the number of friends.
21         System.out.print("How many friends do you have? ");
22         numFriends = keyboard.nextInt();
23
24         // Consume the remaining newline character.
25         keyboard.nextLine();
26
27         // Get the filename.
28         System.out.print("Enter the filename: ");
29         filename = keyboard.nextLine();
30
31         // Make sure the file does not exist.
32         File file = new File(filename);
33         if (file.exists())
34         {
35             System.out.println("The file " + filename +
36                               " already exists.");
37             System.exit(0);
38         }
39
40         // Open the file.
41         PrintWriter outputFile = new PrintWriter(file);
42
43         // Get data and write it to the file.
44         for (int i = 1; i <= numFriends; i++)
45         {
46             // Get the name of a friend.
47             System.out.print("Enter the name of friend " +
48                             "number " + i + ": ");
49             friendName = keyboard.nextLine();
50
51             // Write the name to the file.
52             outputFile.println(friendName);
53         }
54
55         // Close the file.
56         outputFile.close();
57         System.out.println("Data written to the file.");
58     }

```


Program Output with Example Input Shown in Bold

```
How many friends do you have? 2 [Enter]  
Enter the filename: MyFriends.txt [Enter]  
The file MyFriends.txt already exists.
```

Line 32 creates a `File` object representing the file. The `if` statement in line 33 calls the `file.exists()` method. If the method returns `true`, then the file exists. In this case the code in lines 35 through 37 executes. This code displays an error message and shuts the program down. If the file does not exist, the rest of the program executes.

Notice that in line 41 we pass a reference to the `File` object to the `PrintWriter` constructor. In previous programs that created an instance `PrintWriter`, we passed a filename to the constructor. If you have a reference to a `File` object that represents the file you wish to open, as we do in this program, you have the option of passing it to the `PrintWriter` constructor.



Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 4.16 What is the difference between an input file and an output file?
2. 4.17 What `import` statement will you need in a program that performs file operations?
3. 4.18 What class do you use to write data to a file?
4. 4.19 Write code that does the following: opens a file named *MyName.txt*, writes your first name to the file, and then closes the file.
5. 4.20 What classes do you use to read data from a file?
6. 4.21 Write code that does the following: opens a file named *MyName.txt*, reads the first line from the file and displays it, and then closes the file.
7. 4.22 You are opening an existing file for output. How do you open the file without erasing it, and at the same time make sure that new data that is written to the file is appended to the end of the file's existing data?
8. 4.23 What clause must you write in the header of a method that performs a file operation?

4.11 Generating Random Numbers with the Random Class

CONCEPT: Random numbers are used in a variety of applications. Java provides the `Random` class that you can use to generate random numbers.

Random numbers are useful for lots of different programming tasks. The following are just a few examples.

- Random numbers are commonly used in games. For example, computer games that let the player roll dice use random numbers to represent the values of the dice. Programs that show cards being drawn from a shuffled deck use random numbers to represent the face values of the cards.
- Random numbers are useful in simulation programs. In some simulations, the computer must randomly decide how a person, animal, insect, or other living being will behave. Formulas can be constructed in which a random number is used to determine various actions and events that take place in the program.
- Random numbers are useful in statistical programs that must randomly select data for analysis.
- Random numbers are commonly used in computer security to encrypt sensitive data.

The Java API provides a class named `Random` that you can use to generate random numbers. The class is part of the `java.util` package, so any program that uses it will need an `import` statement such as:

```
import java.util.Random;
```

You create an object from the `Random` class with a statement such as this:

```
Random randomNumbers = new Random();
```

This statement does the following:

- It declares a variable named `randomNumbers`. The data type is the `Random` class.
- The expression `new Random()` creates an instance of the `Random` class.
- The equal sign assigns the address of the `Random` class to the `randomNumbers` variable.

After this statement executes, the `randomNumbers` variable will reference a `Random` object. Once you have created a `Random` object, you can call its `nextInt` method to get a random integer number. The following code shows an example:

```
// Declare an int variable.
```

```
int number;
// Create a Random object.
Random randomNumbers = new Random();
// Get a random integer and assign it to number.
number = randomNumbers.nextInt();
```

After this code executes, the `number` variable will contain a random integer. If you call the `nextInt` method with no arguments, as shown in this example, the returned integer is somewhere between $-2,147,483,648$ and $+2,147,483,647$. Alternatively, you can pass an argument that specifies an upper limit to the generated number’s range. In the following statement, the value assigned to `number` is somewhere between 0 and 99:

```
number = randomNumbers.nextInt(100);
```

You can add or subtract a value to shift the numeric range upward or downward. In the following statement, we call the `nextInt` method to get a random number in the range of 0 through 9, and then we add 1 to it. So, the number assigned to `number` will be somewhere in the range of 1 through 10:

```
number = randomNumbers.nextInt(10) + 1;
```

The following statement shows another example. It assigns a random integer to `number` between -50 and $+49$:

```
number = randomNumbers.nextInt(100) - 50
```

The `Random` class has other methods for generating random numbers, and [Table 4-1](#) summarizes several of them.

Table 4-1 Some of the Random class’s methods

Method	Description
<code>nextDouble()</code>	Returns the next random number as a <code>double</code> . The number will be within the range of <code>0.0</code> through <code>1.0</code> .
<code>nextFloat()</code>	Returns the next random number as a <code>float</code> . The number will be within the range of <code>0.0</code> through <code>1.0</code> .
<code>nextInt()</code>	Returns the next random number as an <code>int</code> . The number will be within the range of an <code>int</code> , which is $-2,147,483,648$ to $+2,147,483,648$.

`nextInt(int n)` This method accepts an integer argument, *n*. It returns a random number as an `int`. The number will be within the range of 0 through *n*.

`nextLong()` Returns the next random number as a `long`. The number will be within the range of a `long`, which is $-9,223,372,036,854,775,808$ to $+9,223,372,036,854,775,808$.

The program in [Code Listing 4-23](#) demonstrates using the `Random` class.

Code Listing 4-23 (MathTutor.java)

```
1  import java.util.Scanner;        // Needed for the Scanner class
2  import java.util.Random;        // Needed for the Random class
3
4  /**
5   * This program demonstrates the Random class.
6   */
7
8  public class MathTutor
9  {
10     public static void main(String[] args)
11     {
12         int number1;              // A number
13         int number2;              // Another number
14         int sum;                  // The sum of the numbers
15         int userAnswer;           // The user's answer
16
17         // Create a Scanner object for keyboard input.
18         Scanner keyboard = new Scanner(System.in);
19
20         // Create a Random class object.
21         Random randomNumbers = new Random();
22
23         // Get two random numbers.
24         number1 = randomNumbers.nextInt(100);
25         number2 = randomNumbers.nextInt(100);
26
27         // Display an addition problem.
28         System.out.println("What is the answer to the " +
29                             "following problem?");
30         System.out.print(number1 + " + " +
31                             number2 + " = ? ");
32
33         // Calculate the answer.
34         sum = number1 + number2;
35
36         // Get the user's answer.
37         userAnswer = keyboard.nextInt();
38
```

```

39 // Display the user's results.
40 if (userAnswer == sum)
41     System.out.println("Correct!");
42 else
43 {
44     System.out.println("Sorry, wrong answer. " +
45                         "The correct answer is " +
46                         sum);
47 }
48 }
49 }

```

Program Output with Example Input Shown in Bold

```

What is the answer to the following problem?
52 + 19 = ? 71 [Enter]
Correct!

```

Program Output with Example Input Shown in Bold

```

What is the answer to the following problem?
27 + 73 = ? 101 [Enter]
Sorry, wrong answer. The correct answer is 100

```

In the Spotlight:



Using Random Numbers

Dr. Kimura teaches an introductory statistics class, and has asked you to write a program that he can use in class to simulate the rolling of dice. The program should randomly generate two numbers in the range of 1 through 6 and display them. In your interview with Dr. Kimura, you learn that he would like to use the program to simulate several rolls of the dice, one after the other. Here is the pseudocode for the program:

While the user wants to roll the dice:

- *Display a random number in the range of 1 through 6*

- *Display another random number in the range of 1 through 6*
- *Ask the user if he or she wants to roll the dice again*

You will write a `while` loop that simulates one roll of the dice, and then asks the user whether another roll should be performed. As long as the user answers “y” for yes, the loop will repeat. [Code Listing 4-24](#) shows the program.

Code Listing 4-24 (RollDice.java)

```

1  import java.util.Scanner;
2  import java.util.Random;
3
4  /**
5   * This program simulates the rolling of dice.
6   */
7
8  public class RollDice
9  {
10     public static void main(String[] args)
11     {
12         String again = "y";    // To control the loop
13         int die1;              // To hold the value of die #1
14         int die2;              // to hold the value of die #2
15
16         // Create a Scanner object to read keyboard input.
17         Scanner keyboard = new Scanner(System.in);
18
19         // Create a Random object to generate random numbers.
20         Random rand = new Random();
21
22         // Simulate rolling the dice.
23         while (again.equalsIgnoreCase("y"))
24         {
25             System.out.println("Rolling the dice . . . ");
26             die1 = rand.nextInt(6) + 1;
27             die2 = rand.nextInt(6) + 1;
28             System.out.println("Their values are:");
29             System.out.println(die1 + " " + die2);
30
31             System.out.print("Roll them again (y = yes)? ");
32             again = keyboard.nextLine();
33         }
34     }
35 }

```

Program Output with Example Input Shown in Bold

```
Rolling the dice . . .
Their values are:
4 3
Roll them again (y = yes)?    y [Enter]
Rolling the dice . . .
Their values are:
2 6
Roll them again (y = yes)?    y [Enter]
Rolling the dice . . .
Their values are:
1 5
Roll them again (y = yes)?    n [Enter]
```



In the Spotlight:

Using Random Numbers to Represent Other Values

Dr. Kimura was so happy with the dice rolling simulator that you wrote for him, he has asked you to write one more program. He would like a program that he can use to simulate ten coin tosses, one after the other. Each time the program simulates a coin toss, it should randomly display either “Heads” or “Tails”.

You decide that you can simulate the tossing of a coin by randomly generating a number in the range of 0 through 1. You will write an `if` statement that displays “Tails” if the random number is 0, or “Heads” otherwise. Here is the pseudocode:

Repeat 10 times:

- *If a random number in the range of 0 through 1 equals 0, then:*
 - *Display “Tails”*
- *Else:*
 - *Display “Heads”*

Because the program should simulate 10 tosses of a coin, you decide to use a `for` loop. The program is shown in [Code Listing 4-25](#).

Code Listing 4-25 (CoinToss.java)

```
1 import java.util.Random;
2
```



```

3  /**
4   This program simulates 10 tosses of a coin.
5  */
6
7  public class CoinToss
8  {
9      public static void main(String[] args)
10     {
11         // Create a Random object to generate random numbers.
12         Random rand = new Random();
13
14         // Simulate the coin tosses.
15         for (int count = 0; count < 10; count++)
16         {
17             if (rand.nextInt(2) == 0)
18                 System.out.println("Tails");
19             else
20                 System.out.println("Heads");
21         }
22     }
23 }

```

Program Output

```

Tails
Tails
Heads
Tails
Heads
Heads
Heads
Heads
Tails
Heads
Tails

```



MyProgrammingLab™ www.myprogramminglab.com

1. 4.24 Assume `x` is an `int` variable, and `rand` references a `Random` object. What does the following statement do?

```
x = rand.nextInt();
```

2. 4.25 Assume `x` is an `int` variable, and `rand` references a `Random` object. What does the following statement do?

```
x = rand.nextInt(100);
```

3. 4.26 Assume `x` is an `int` variable, and `rand` references a `Random` object. What does the following statement do?

```
x = rand.nextInt(9) + 1;
```

4. 4.27 Assume `x` is a `double` variable, and `rand` references a `Random` object. What does the following statement do?

```
x = rand.nextDouble();
```

4.12 Common Errors to Avoid

The following list describes several errors that are commonly committed when learning this chapter's topics.

- **Using the increment or decrement operator in the wrong mode.** When the increment or decrement operator is placed in front of (to the left of) its operand, it is used in prefix mode. When either of these operators is placed behind (to the right of) its operand, it is used in postfix mode.
- **Forgetting to enclose the `boolean` expression in a `while` loop or a `do-while` loop inside parentheses.**
- **Placing a semicolon at the end of a `while` or `for` loop's header.** When you write a semicolon at the end of a `while` or `for` loop's header, Java assumes that the conditionally executed statement is a `null` or empty statement. This usually results in an infinite loop.
- **Forgetting to write the semicolon at the end of the `do-while` loop.** The `do-while` loop must be terminated with a semicolon.
- **Forgetting to enclose multiple statements in the body of a loop in braces.** Normally a loop conditionally executes only one statement. To conditionally execute more than one statement, you must place the statements in braces.
- **Using commas instead of semicolons to separate the initialization, test, and update expressions in a `for` loop.**
- **Forgetting to write code in the body of a `while` or `do-while` loop that modifies the loop control variable.** If a `while` or `do-while` loop's `boolean` expression never becomes `false`, the loop will repeat indefinitely. You must have code in the body of the loop that modifies the loop control variable so that the `boolean` expression will at some point become `false`.
- **Using a sentinel value that can also be a valid data value.** Remember, a sentinel is a special value that cannot be mistaken as a member of a list of data items and signals that there are no more data items from the list to be processed. If you choose as a sentinel a value that might also appear in the list, the loop will prematurely terminate if it encounters the value in the list.
- **Forgetting to initialize an accumulator to zero.** In order for an accumulator to keep a correct running total, it must be initialized to zero before any values are added to it.

Review Questions and Exercises

Multiple Choice and True/False

1. What will the `println` statement in the following program segment display?

```
int x = 5;  
System.out.println(x++);
```

1. 5
2. 6
3. 0
4. None of these

2. What will the `println` statement in the following program segment display?

```
int x = 5;  
System.out.println(++x);
```

1. 5
2. 6
3. 0
4. None of these

3. In the expression `number++`, the `++` operator is in what mode?

1. prefix
2. pretest
3. postfix
4. posttest

4. What is each repetition of a loop known as?

1. cycle
2. revolution

3. orbit
4. iteration
5. This is a variable that controls the number of iterations performed by a loop.
 1. loop control variable
 2. accumulator
 3. iteration register variable
 4. repetition meter
6. The `while` loop is this type of loop.
 1. pretest
 2. posttest
 3. prefix
 4. postfix
7. The `do-while` loop is this type of loop.
 1. pretest
 2. posttest
 3. prefix
 4. postfix
8. The `for` loop is this type of loop.
 1. pretest
 2. posttest
 3. prefix
 4. postfix
9. This type of loop has no way of ending and repeats until the program is interrupted.
 1. indeterminate

2. interminable
3. infinite
4. timeless

10. This type of loop always executes at least once.

1. `while`
2. `do-while`
3. `for`
4. any of these

11. This expression is executed by the `for` loop only once, regardless of the number of iterations.

1. initialization expression
2. test expression
3. update expression
4. pre-increment expression

12. This is a variable that keeps a running total.

1. sentinel
2. sum
3. total
4. accumulator

13. This is a special value that signals when there are no more items from a list of items to be processed. This value cannot be mistaken as an item from the list.

1. sentinel
2. flag
3. signal
4. accumulator

14. To open a file for writing, you use the following class.

1. `PrintWriter`
2. `FileOpen`
3. `OutputFile`
4. `FileReader`

15. To open a file for reading, you use the following classes.

1. `File` and `Writer`
2. `File` and `Output`
3. `File` and `Input`
4. `File` and `Scanner`

16. When a program is finished using a file, it should do this.

1. erase the file
2. close the file
3. throw an exception
4. reset the read position

17. This class allows you to use the `print` and `println` methods to write data to a file.

1. `File`
2. `FileReader`
3. `OutputFile`
4. `PrintWriter`

18. This class allows you to read a line from a file.

1. `FileWriter`
2. `Scanner`
3. `InputFile`
4. `FileReader`

19. True or False: The `while` loop is a pretest loop.
20. True or False: The `do-while` loop is a pretest loop.
21. True or False: The `for` loop is a posttest loop.
22. True or False: It is not necessary to initialize accumulator variables.
23. True or False: One limitation of the `for` loop is that only one variable may be initialized in the initialization expression.
24. True or False: A variable may be defined in the initialization expression of the `for` loop.
25. True or False: In a nested loop, the inner loop goes through all of its iterations for every iteration of the outer loop.
26. True or False: To calculate the total number of iterations of a nested loop, add the number of iterations of all the loops.

Find the Error

Find the errors in the following code:

1.

```
// This code contains ERRORS!
// It adds two numbers entered by the user.
int num1, num2;
String input;
char again;
Scanner keyboard = new Scanner(System.in);
while (again == 'y' || again == 'Y')
    System.out.print("Enter a number: ");
    num1 = keyboard.nextInt();
    System.out.print("Enter another number: ");
    num2 = keyboard.nextInt();
    System.out.println("Their sum is "+ (num1 + num2));
    System.out.println("Do you want to do this again? ");
    keyboard.nextLine(); // Consume remaining newline
    input = keyboard.nextLine();
    again = input.charAt(0);
```
2.

```
// This code contains ERRORS!
int count = 1, total;
while (count <= 100)
    total += count;
System.out.print("The sum of the numbers 1 - 100 is ");
System.out.println(total);
```
3.

```
// This code contains ERRORS!
int choice, num1, num2;
Scanner keyboard = new Scanner(System.in);
```



```

do
{
    System.out.print("Enter a number: ");
    num1 = keyboard.nextInt();
    System.out.print("Enter another number: ");
    num2 = keyboard.nextInt();
    System.out.println("Their sum is " + (num1 + num2));
    System.out.println("Do you want to do this again? ");
    System.out.print("1 = yes, 0 = no ");
    choice = keyboard.nextInt();
} while (choice = 1)

```

```

4. // This code contains ERRORS!
// Print the numbers 1 through 10.
for (int count = 1, count <= 10, count++;)
{
    System.out.println(count);
    count++;
}

```

Algorithm Workbench

1. Write a `while` loop that lets the user enter a number. The number should be multiplied by 10, and the result stored in the variable `product`. The loop should iterate as long as `product` contains a value less than 100.
2. Write a `do-while` loop that asks the user to enter two numbers. The numbers should be added and the sum displayed. The loop should ask the user whether he or she wishes to perform the operation again. If so, the loop should repeat; otherwise it should terminate.

3. Write a `for` loop that displays the following set of numbers:

0, 10, 20, 30, 40, 50 . . . 1000

4. Write a loop that asks the user to enter a number. The loop should iterate 10 times and keep a running total of the numbers entered.

5. Write a `for` loop that calculates the total of the following series of numbers:

$1 \cdot 30 + 2 \cdot 29 + 3 \cdot 28 + \cdots + 30 \cdot 1$

6. Write a nested loop that displays 10 rows of '#' characters. There should be 15 '#' characters in each row.

7. Convert the `while` loop in the following code to a `do-while` loop:

```

Scanner keyboard = new Scanner(System.in);
int x = 1;
while (x > 0)
{

```

```
System.out.print("Enter a number: ");
x = keyboard.nextInt();
}
```

8. Convert the `do-while` loop in the following code to a `while` loop:

```
Scanner keyboard = new Scanner(System.in);
String input;
char sure;
do
{
System.out.print("Are you sure you want to quit? ");
input = keyboard.next();
sure = input.charAt(0);
} while (sure != 'Y' && sure != 'N');
```

9. Convert the following `while` loop to a `for` loop:

```
int count = 0;
while (count < 50)
{
    System.out.println("count is " + count);
    count++;
}
```

10. Convert the following `for` loop to a `while` loop:

```
for (int x = 50; x > 0; x--)
{
    System.out.println(x + " seconds to go.");
}
```

11. Write an input validation loop that asks the user to enter a number in the range of 1 through 4.
12. Write an input validation loop that asks the user to enter the word “yes” or “no”.
13. Write nested loops to draw this pattern:

```
*****
*****
*****
****
***
**
*
```

14. Write nested loops to draw this pattern:

```
##
# #
#  #
#   #
#    #
#     #
```

15. Complete the following program so it displays a random integer in the range of 1 through 10.

```
// Write the necessary import statement(s) here.
public class ReviewQuestion15
{
    public static void main(String[] args)
    {
        // Write the necessary code here.
    }
}
```

16. Complete the following program so it performs the following actions 10 times:

- Generates a random number that is either 0 or 1.
- Displays either the word “Yes” or the word “No” depending on the random number that was generated.

```
// Write the necessary import statement(s) here.
public class ReviewQuestion16
{
    public static void main(String[] args)
    {
        // Write the necessary code here.
    }
}
```

17. Write code that does the following: opens a file named *NumberList.txt*, uses a loop to write the numbers 1 through 100 to the file, and then closes the file.

18. Write code that does the following: opens the *NumberList.txt* file that was created by the code in [Question 17](#), reads all of the numbers from the file and displays them, and then closes the file.

19. Modify the code you wrote in [Question 18](#) so it adds all of the numbers read from the file and displays their total.

20. Write code that opens a file named *NumberList.txt* for writing, but does not erase the file’s contents if it already exists.

Short Answer

1. Briefly describe the difference between the prefix and postfix modes used by the increment and decrement operators.
2. Why should you indent the statements in the body of a loop?
3. Describe the difference between pretest loops and posttest loops.
4. Why are the statements in the body of a loop called conditionally executed statements?

5. Describe the difference between the `while` loop and the `do-while` loop.
6. Which loop should you use in situations where you want the loop to repeat until the `boolean` expression is `false`, and the loop should not execute if the test expression is `false` to begin with?
7. Which loop should you use in situations where you want the loop to repeat until the `boolean` expression is `false`, but the loop should execute at least once?
8. Which loop should you use when you know the number of required iterations?
9. Why is it critical that accumulator variables are properly initialized?
10. What is an infinite loop? Write the code for an infinite loop.
11. Describe a programming problem that would require the use of an accumulator.
12. What does it mean to let the user control a loop?
13. What is the advantage of using a sentinel?
14. Why must the value chosen for use as a sentinel be carefully selected?
15. Describe a programming problem requiring the use of nested loops.
16. How does a file buffer increase a program's performance?
17. Why should a program close a file when it's finished using it?
18. What is a file's read position? Where is the read position when a file is first opened for reading?
19. When writing data to a file, what is the difference between the `print` and the `println` methods?
20. What does the `Scanner` class's `hasNext` method return when the end of the file has been reached?
21. What is a potential error that can occur when a file is opened for reading?
22. What does it mean to append data to a file?
23. How do you open a file so that new data will be written to the end of the file's existing data?

Programming Challenges

MyProgrammingLab™ Visit www.myprogramminglab.com to complete many of these Programming Challenges online and get instant feedback.

1. Sum of Numbers

Write a program that asks the user for a positive nonzero integer value. The program should use a loop to get the sum of all the integers from 1 up to the number entered. For example, if the user enters 50, the loop will find the sum of 1, 2, 3, 4, . . . 50.

2. Distance Traveled

The distance a vehicle travels can be calculated as follows:

$$\text{Distance} = \text{Speed} * \text{Time}$$

For example, if a train travels 40 miles-per-hour for three hours, the distance traveled is 120 miles. Write a program that asks for the speed of a vehicle (in miles-per-hour) and the number of hours it has traveled. It should use a loop to display the distance a vehicle has traveled for each hour of a time period specified by the user. For example, if a vehicle is traveling at 40 mph for a three-hour time period, it should display a report similar to the one that follows:

Hour	Distance Traveled

1	40

2	80

3	120

Input Validation: Do not accept a negative number for speed and do not accept any value less than 1 for time traveled.

3. Distance File

Modify the program you wrote for Programming Challenge 2 (Distance Traveled) so it writes the report to a file instead of the screen. Open the file in Notepad or another text editor to confirm the output.

4. Pennies for Pay

Write a program that calculates the amount a person would earn over a period of time if his or her salary is one penny the first day, two pennies the second day, and continues to double each day. The program should display a table showing the salary for each day, and then show the total pay at the end of the period. The output should be displayed in a dollar amount, not the number of pennies.



The Pennies for Pay Problem

Input Validation: Do not accept a number less than 1 for the number of days worked.

5. Letter Counter

Write a program that asks the user to enter a string, and then asks the user to enter a character. The program should count and display the number of times that the specified character appears in the string.

6. File Letter Counter

Write a program that asks the user to enter the name of a file, and then asks the user to enter a character. The program should count and display the number of times that the specified character appears in the file. Use Notepad or another text editor to create a simple file that can be used to test the program.

7. Hotel Occupancy

A hotel's occupancy rate is calculated as follows:

$$\text{Occupancy rate} = \text{Number of rooms occupied} \div \text{Total number of rooms}$$

Write a program that calculates the occupancy rate for each floor of a hotel. The program should start

by asking for the number of floors in the hotel. A loop should then iterate once for each floor. During each iteration, the loop should ask the user for the number of rooms on the floor and the number of them that are occupied. After all the iterations, the program should display the number of rooms the hotel has, the number of them that are occupied, the number that are vacant, and the occupancy rate for the hotel.

Input Validation: Do not accept a value less than 1 for the number of floors. Do not accept a number less than 10 for the number of rooms on a floor.

8. Average Rainfall

Write a program that uses nested loops to collect data and calculate the average rainfall over a period of years. First the program should ask for the number of years. The outer loop will iterate once for each year. The inner loop will iterate 12 times, once for each month. Each iteration of the inner loop will ask the user for the inches of rainfall for that month. After all iterations, the program should display the number of months, the total inches of rainfall, and the average rainfall per month for the entire period.

Input Validation: Do not accept a number less than 1 for the number of years. Do not accept negative numbers for the monthly rainfall.

9. Population

Write a program that will predict the size of a population of organisms. The program should ask for the starting number of organisms, their average daily population increase (as a percentage), and the number of days they will multiply. For example, a population might begin with two organisms, have an average daily increase of 50 percent, and will be allowed to multiply for seven days. The program should use a loop to display the size of the population for each day.

Input Validation: Do not accept a number less than 2 for the starting size of the population. Do not accept a negative number for average daily population increase. Do not accept a number less than 1 for the number of days they will multiply.

10. Largest and Smallest

Write a program with a loop that lets the user enter a series of integers. The user should enter -99 to signal the end of the series. After all the numbers have been entered, the program should display the largest and smallest numbers entered.

11. Celsius to Fahrenheit Table

Write a program that displays a table of the Celsius temperatures 0 through 20 and their Fahrenheit equivalents. The formula for converting a temperature from Celsius to Fahrenheit is

$$F = 9/5 \cdot C + 32$$

where F is the Fahrenheit temperature and C is the Celsius temperature. Your program must use a loop to display the table.

12. Bar Chart

Write a program that asks the user to enter today's sales for five stores. The program should display a bar chart comparing each store's sales. Create each bar in the bar chart by displaying a row of asterisks. Each asterisk should represent \$100 of sales. Here is an example of the program's output:

```
Enter today's sales for store 1: 1000 [Enter]
Enter today's sales for store 2: 1200 [Enter]
Enter today's sales for store 3: 1800 [Enter]
Enter today's sales for store 4: 800 [Enter]
Enter today's sales for store 5: 1900 [Enter]
```

```
SALES BAR CHART
Store 1: *****
Store 2: *****
Store 3: *****
Store 4: *****
Store 5: *****
```

13. File Head Display

Write a program that asks the user for the name of a file. The program should display only the first five lines of the file's contents. If the file contains fewer than five lines, it should display the file's entire contents.

14. Line Numbers

Write a program that asks the user for the name of a file. The program should display the contents of the file with each line preceded with a line number followed by a colon. The line numbering should start at 1.

15. Uppercase File Converter

Write a program that asks the user for the names of two files. The first file should be opened for

reading and the second file should be opened for writing. The program should read the contents of the first file, change all characters to uppercase, and store the results in the second file. The second file will be a copy of the first file, except that all the characters will be uppercase. Use Notepad or another text editor to create a simple file that can be used to test the program.

16. Budget Analysis

Write a program that asks the user to enter the amount that he or she has budgeted for a month. A loop should then prompt the user to enter each of his or her expenses for the month, and keep a running total. When the loop finishes, the program should display the amount that the user is over or under budget.

17. Random Number Guessing Game

Write a program that generates a random number and asks the user to guess what the number is. If the user's guess is higher than the random number, the program should display "Too high, try again." If the user's guess is lower than the random number, the program should display "Too low, try again." The program should use a loop that repeats until the user correctly guesses the random number.

18. Random Number Guessing Game Enhancement

Enhance the program that you wrote for Programming Challenge 17 so it keeps a count of the number of guesses that the user makes. When the user correctly guesses the random number, the program should display the number of guesses.

19. ESP Game

Write a program that tests your ESP (extrasensory perception). The program should randomly select the name of a color from the following list of words:

Red, Green, Blue, Orange, Yellow

To select a word, the program can generate a random number. For example, if the number is 0, the selected word is *Red*, if the number is 1, the selected word is *Green*, and so forth.

Next, the program should ask the user to enter the color that the computer has selected. After the user has entered his or her guess, the program should display the name of the randomly selected color. The program should repeat this 10 times and then display the number of times the user correctly guessed

the selected color.

20. Square Display

Write a program that asks the user for a positive integer no greater than 15. The program should then display a square on the screen using the character 'X'. The number entered by the user will be the length of each side of the square. For example, if the user enters 5, the program should display the following:

```
XXXXX
XXXXX
XXXXX
XXXXX
XXXXX
```

If the user enters 8, the program should display the following:

```
XXXXXXXX
XXXXXXXX
XXXXXXXX
XXXXXXXX
XXXXXXXX
XXXXXXXX
XXXXXXXX
XXXXXXXX
XXXXXXXX
```

21. Dice Game

Write a program that plays a simple dice game between the computer and the user. When the program runs, a loop should repeat 10 times. Each iteration of the loop should do the following:

- Generate a random integer in the range of 1 through 6. This is the value of the computer's die.
- Generate another random integer in the range of 1 through 6. This is the value of the user's die.
- The die with the highest value wins. (In case of a tie, there is no winner for that particular roll of the dice.)

As the loop iterates, the program should keep count of the number of times the computer wins, and the number of times that the user wins. After the loop performs all of its iterations, the program should display who was the grand winner, the computer or the user.

22. Slot Machine Simulation

A slot machine is a gambling device that the user inserts money into and then pulls a lever (or presses a button). The slot machine then displays a set of random images. If two or more of the images match, the user wins an amount of money that the slot machine dispenses back to the user.

Create a program that simulates a slot machine. When the program runs, it should do the following:

- Asks the user to enter the amount of money he or she wants to enter into the slot machine.
- Instead of displaying images, the program will randomly select a word from the following list:

Cherries, Oranges, Plums, Bells, Melons, Bars

To select a word, the program can generate a random number in the range of 0 through 5. If the number is 0, the selected word is *Cherries*; if the number is 1, the selected word is *Oranges*; and so forth. The program should randomly select a word from this list three times and display all three of the words.

- If none of the randomly selected words match, the program will inform the user that he or she has won \$0. If two of the words match, the program will inform the user that he or she has won two times the amount entered. If three of the words match, the program will inform the user that he or she has won three times the amount entered.
- The program will ask whether the user wants to play again. If so, these steps are repeated. If not, the program displays the total amount of money entered into the slot machine and the total amount won.

23. Personal Web Page Generator

Write a program that asks the user for his or her name, and then asks the user to enter a sentence that describes him or herself. Here is an example of the program's screen:

Enter your name: **Julie Taylor [Enter]**

Describe yourself: **I am a computer science major, a member of the Jazz club, and**

Once the user has entered the requested input, the program should create an HTML file, containing the input, for a simple Web page. Here is an example of the HTML content, using the sample input previously shown:

```
<html>
<head>
</head>
<body>
  <center>
    <h1>Julie Taylor</h1>
  </center>
  <hr />
```

I am a computer science major, a member of the Jazz club,
and I hope to work as a mobile app developer after I graduate.

<hr />

</body>

</html>

Chapter 5 Methods

TOPICS

1. [5.1 Introduction to Methods](#)
2. [5.2 Passing Arguments to a Method](#)
3. [5.3 More about Local Variables](#)
4. [5.4 Returning a Value from a Method](#)
5. [5.5 Problem Solving with Methods](#)
6. [5.6 Common Errors to Avoid](#)

5.1 Introduction to Methods

CONCEPT: Methods can be used to break a complex program into small, manageable pieces. A void method simply executes a group of statements and then terminates. A value-returning method returns a value to the statement that called it.

In a general sense, a method is a collection of statements that performs a specific task. So far you have experienced methods in two ways: (1) You have created a method named `main` in every program you've written, and (2) you have executed predefined methods from the Java API, such as `System.out.println`, `Integer.parseInt`, and `Math.pow`. In this chapter, you will learn how to create your own methods, other than `main`, that can be executed just as you execute the API methods.

Methods are commonly used to break a problem into small, manageable pieces. Instead of writing one long method that contains all of the statements necessary to solve a problem, several small methods that each solve a specific part of the problem can be written. These small methods can then be executed in the desired order to solve the problem. This approach is sometimes called *divide and conquer* because a large problem is divided into several smaller problems that are easily solved. [Figure 5-1](#) illustrates this idea by comparing two programs: one that uses a long, complex method containing all of the statements necessary to solve a problem, and another that divides a problem into smaller problems, each of which is handled by a separate method.

In this program the problem has been divided into smaller problems, each of which is handled by a separate method.




Figure 5-1 Using methods to divide and conquer a problem

Figure 5-1 Full Alternative Text

Another reason to write methods is that they simplify programs. If a specific task is performed in several places in a program, a method can be written once to perform that task, and then be executed any time it is needed. This benefit of using methods is known as *code reuse* because you are writing the code to perform a task once and then reusing it each time you need to perform the task.

First, we will look at the general ways in which methods operate. At the end of the chapter we will discuss in greater detail how methods can be used in problem solving.

void Methods and Value-Returning Methods

In this chapter, you will learn about two general categories of methods: `void` methods and value-returning methods. A *void method* is one that simply performs a task and then terminates. `System.out.println` is an example of a `void` method. For example, look at the following code:

```
1    int number = 7;
2    System.out.println(number);
3    number = 0;
```

The statement in line 1 declares the `number` variable and initializes it with the value 7. The statement in line 2 calls the `System.out.println` method, passing `number` as an argument. The method does its job, which is to display a value on the screen, and then terminates. The code then resumes at line 3.

A *value-returning method* not only performs a task but also sends a value back to the code that called it. The `Random` class's `nextInt` method is an example of a value-returning method. For example, look at the following code:

```
1    int number;
2    Random rand = new Random();
3    number = rand.nextInt();
```

The statement in line 1 declares the `number` variable. Line 2 creates a `Random` object and assigns its address to a variable named `rand`. Line 3 is an assignment statement, which assigns a value to the `number` variable. Notice that on the right side of the `=` operator is a call to the `rand.nextInt` method. The method executes, and then returns a value. The value that is returned from the method is assigned to the `number` variable.

Defining a void Method

To create a method you must write its *definition*, which consists of two general parts: a header and a body. You learned about both of these in [Chapter 2](#), but let's briefly review. The *method header*, which appears at the beginning of a method definition, lists several important things about the method, including the method's name. The *method body* is a collection of statements that are performed when the method is executed. These statements are enclosed inside a set of curly braces. [Figure 5-2](#) points out the header and body of a `main` method.




Figure 5-2 The header and body of a main method

[Figure 5-2 Full Alternative Text](#)

As you already know, every complete Java program must have a `main` method. Java programs can have other methods as well. Here is an example of a simple method that displays a message on the screen:

```
public static void displayMessage()  
{  
    System.out.println("Hello from the displayMessage method.");  
}
```

This method has a header and a body. [Figure 5-3](#) shows the different parts of the method header.




Figure 5-3 Parts of the method header

[Figure 5-3 Full Alternative Text](#)

Let's take a closer look at the parts identified in the figure as follows:

- **Method modifiers**—The key words `public` and `static` are modifiers. You don't need to be too concerned with these modifiers now, but if your curiosity is getting the best of you, here's a brief explanation: The word `public` means that the method is publicly available to code outside the class. The word `static` means that the method belongs to the class, not a specific object. You will learn more about these modifiers in later chapters. For this chapter, every method that we write will begin with `public static`.
- **Return type**—Recall our previous discussion of `void` and value-returning methods. When the key word `void` appears here, it means that the method is a `void` method, and does not return a value. As you will see later, a value-returning method lists a data type here.
- **Method name**—You should give each method a descriptive name. In general, the same rules that apply to variable names also apply to method names. This method is named `displayMessage`,

so we can easily guess what the method does: It displays a message.

- Parentheses—In the header, the method name is always followed by a set of parentheses. As you will learn later in this chapter, methods can be capable of receiving arguments. When this is the case, a list of one or more variable declarations will appear inside the parentheses. The method in this example does not receive any arguments, so the parentheses are empty.



The method header is never terminated with a semicolon.

Calling a Method

A method executes when it is called. The `main` method is automatically called when a program starts, but other methods are executed by method call statements. When a method is called, the JVM branches to that method and executes the statements in its body. Here is an example of a method call statement that calls the `displayMessage` method we previously examined:

```
displayMessage();
```

The statement is simply the name of the method followed by a set of parentheses. Because it is a complete statement, it is terminated with a semicolon.



Notice that the method modifiers and the `void` return type are not written in the method call statement. They are written only in the method header.

The program in [Code Listing 5-1](#) demonstrates.

Code Listing 5-1 (SimpleMethod.java)

```
1  /**
2   * This program defines and calls a simple method.
3   */
4
5  public class SimpleMethod
6  {
7      public static void main(String[] args)
8      {
9          System.out.println("Hello from the main method.");
10     }
```

```

10         displayMessage();
11         System.out.println("Back in the main method.");
12     }
13
14     /**
15      * The displayMessage method displays a greeting.
16      */
17
18     public static void displayMessage()
19     {
20         System.out.println("Hello from the displayMessage method.");
21     }
22 }

```

Program Output

```

Hello from the main method.
Hello from the displayMessage method.
Back in the main method.

```

Notice how the program flows. It starts, of course, in the `main` method. When the call to the `displayMessage` method in line 10 is encountered, the JVM branches to that method and performs the statement in its body (at line 20). Once the `displayMessage` method has finished executing, the JVM branches back to the `main` method and resumes at line 11 with the statement that follows the method call. This is illustrated in [Figure 5-4](#).




Figure 5-4 Branching in the SimpleMethod.java program

[Figure 5-4 Full Alternative Text](#)

Method call statements may be used in control structures like loops, `if` statements, and `switch` statements. The program in [Code Listing 5-2](#) places the `displayMessage` method call inside a loop.

Code Listing 5-2 (LoopCall.java)

```
1  /**
2   This program defines and calls a simple method.
3  */
4
5  public class LoopCall
6  {
7      public static void main(String[] args)
8      {
9          System.out.println("Hello from the main method.");
10         for (int i = 0; i < 5; i++)
11             displayMessage();
12         System.out.println("Back in the main method.");
13     }
14
15     /**
16      The displayMessage method displays a greeting.
17     */
18
19     public static void displayMessage()
20     {
21         System.out.println("Hello from the displayMessage method.");
22     }
23 }
```

Program Output

```
Hello from the main method.
Hello from the displayMessage method.
Hello from the displayMessage method.
Hello from the displayMessage method.
Hello from the displayMessage method.
Hello from the displayMessage method.
Hello from the displayMessage method.
Back in the main method.
```

The program in [Code Listing 5-3](#) shows another example. It asks the user to enter his or her annual salary and credit rating. The program then determines whether the user qualifies for a credit card. One of two void methods, `qualify` or `noQualify`, is called to display a message. [Figures 5-5](#) and [5-6](#) show example interactions with the program.

Code Listing 5-3 (CreditCard.java)

```
1  import javax.swing.JOptionPane;
2
3  /**
4   This program uses two void methods.
5  */
6
```

```

7 public class CreditCard
8 {
9     public static void main(String[] args)
10    {
11        double salary;        // Annual salary
12        int creditRating;    // Credit rating
13        String input;        // To hold the user's input
14
15        // Get the user's annual salary.
16        input = JOptionPane.showInputDialog("What is " +
17                                           "your annual salary?");
18        salary = Double.parseDouble(input);
19
20        // Get the user's credit rating (1 through 10).
21        input = JOptionPane.showInputDialog("On a scale of " +
22                                           "1 through 10, what is your credit rating?\n" +
23                                           "(10 = excellent, 1 = very bad)");
24        creditRating = Integer.parseInt(input);
25
26        // Determine whether the user qualifies.
27        if (salary >= 20000 && creditRating >= 7)
28            qualify();
29        else
30            noQualify();
31
32        System.exit(0);
33    }
34
35    /**
36     The qualify method informs the user that he
37     or she qualifies for the credit card.
38    */
39
40    public static void qualify()
41    {
42        JOptionPane.showMessageDialog(null, "Congratulations! " +
43                                       "You qualify for the credit card!");
44    }
45
46    /**
47     The noQualify method informs the user that he
48     or she does not qualify for the credit card.
49    */
50
51    public static void noQualify()
52    {
53        JOptionPane.showMessageDialog(null, "I'm sorry. You " +
54                                       "do not qualify for the credit card.");
55    }
56 }

```



Figure 5-5 Interaction with the `CreditCard.java` program

(Oracle Corporate Counsel)

[Figure 5-5 Full Alternative Text](#)



Figure 5-6 Interaction with the `CreditCard.java` program

(Oracle Corporate Counsel)

[Figure 5-6 Full Alternative Text](#)

Layered Method Calls

Methods can also be called in layered fashion. In other words, method A can call method B, which can then call method C. When method C finishes, the JVM returns to method B. When method B finishes, the JVM returns to method A. The program in [Code Listing 5-4](#) demonstrates this with three methods: `main`, `deep`, and `deeper`. The `main` method calls the `deep` method, which then calls the `deeper` method.

Code Listing 5-4 (DeepAndDeeper.java)

```
1  /**
2   This program demonstrates layered method calls.
3  */
4
5  public class DeepAndDeeper
6  {
7      public static void main(String[] args)
8      {
9          System.out.println("I am starting in main.");
10         deep();
11         System.out.println("Now I am back in main.");
12     }
13
14     /**
15      The deep method displays a message and then calls
16      the deeper method.
17     */
18
19     public static void deep()
20     {
21         System.out.println("I am now in deep.");
22         deeper();
23         System.out.println("Now I am back in deep.");
24     }
25
26     /**
27      The deeper method simply displays a message.
28     */
29
30     public static void deeper()
31     {
32         System.out.println("I am now in deeper.");
33     }
34 }
```

Program Output

```
I am starting in main.
I am now in deep.
I am now in deeper.
Now I am back in deep.
Now I am back in main.
```

Using Documentation Comments with Methods

You should always document a method by writing comments that appear just before the method's definition. The comments should provide a brief explanation of the method's purpose. Notice that the

programs we've looked at in this chapter use documentation comments. Recall from [Chapter 2](#) that documentation comments begin with `/**` and end with `*/`. These types of comments can be read and processed by a program named `javadoc`, which produces attractive HTML documentation. As we progress through this chapter, you will learn more about documentation comments and how they can be used with methods.

Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 5.1 What is the difference between a `void` method and a value-returning method?

2. 5.2 Is the following line of code a method header or a method call?

```
calcTotal();
```

3. 5.3 Is the following line of code a method header or a method call?

```
public static void calcTotal()
```

4. 5.4 What message will the following program display if the user enters 5? What if the user enters 10? What if the user enters 100?

```
import javax.swing.JOptionPane;
public class Checkpoint
{
    public static void main(String[] args)
    {
        String input;
        int number;

        input = JOptionPane.showInputDialog("Enter a number.");
        number = Integer.parseInt(input);

        if (number < 10)
            method1();
        else
            method2();

        System.exit(0);
    }

    public static void method1()
    {
        JOptionPane.showMessageDialog(null, "Able was I.");
    }
    public static void method2()
    {
        JOptionPane.showMessageDialog(null, "I saw Elba.");
    }
}
```


}

5. 5.5 Write a `void` method that displays your full name. The method should be named `myName`.

5.2 Passing Arguments to a Method

CONCEPT: A method may be written so it accepts arguments. Data can then be passed into the method when it is called.



Passing Arguments to a Method

Values that are sent into a method are called *arguments*. You're already familiar with how to use arguments in a method call. For example, look at the following statement:

```
System.out.println("Hello");
```

This statement calls the `System.out.println` method and passes "Hello" as an argument. Here is another example:

```
number = Integer.parseInt(str);
```

This statement calls the `Integer.parseInt` method and passes the contents of the `str` variable as an argument. By using parameter variables, you can design your own methods that accept data this way. A *parameter variable*, sometimes simply referred to as a *parameter*, is a special variable that holds a value being passed into a method. Here is the definition of a method that uses a parameter:

```
public static void displayValue(int num)
{
    System.out.println("The value is " + num);
}
```

Notice the integer variable declaration that appears inside the parentheses (`int num`). This is the declaration of a parameter variable, which enables the `displayValue` method to accept an integer value as an argument. Here is an example of a call to the `displayValue` method, passing 5 as an argument:

```
displayValue(5);
```

This statement executes the `displayValue` method. The argument that is listed inside the parentheses is copied into the method's parameter variable, `num`. This is illustrated in [Figure 5-7](#).




Figure 5-7 Passing 5 to the `displayValue` method

[Figure 5-7 Full Alternative Text](#)

Inside the `displayValue` method, the variable `num` will contain the value of whatever argument was passed into it. If we pass 5 as the argument, the method will display as follows:

```
The value is 5
```

You may also pass the contents of variables and the values of expressions as arguments. For example, the following statements call the `displayValue` method with various arguments passed:

```
displayValue(x);
displayValue(x * 4);
displayValue(Integer.parseInt("700"));
```

The first statement is simple. It passes the value in the variable `x` as the argument to the `displayValue` method. The second statement is also simple, but it does a little more work: It passes the result of the expression `x * 4` as the argument to the `displayValue` method. The third statement does even more work. It passes the value returned from the `Integer.parseInt` method as the argument to the `displayValue` method. (The `Integer.parseInt` method is called first, and its return value is passed to the `displayValue` method.) The program in [Code Listing 5-5](#) demonstrates these method calls.

Code Listing 5-5 (`PassArg.java`)

```
1  /**
2   * This program demonstrates a method with a parameter.
3   */
4
5  public class PassArg
6  {
7      public static void main(String[] args)
8      {
```

```

9         int x = 10;
10
11         System.out.println("I am passing values to displayValue.");
12         displayValue(5); // Pass 5
13         displayValue(x); // Pass 10
14         displayValue(x * 4); // Pass 40
15         displayValue(Integer.parseInt("700")); // Pass 700
16         System.out.println("That's all!");
17     }
18
19     /**
20      * The displayValue method displays the value
21      * of its integer parameter.
22      */
23
24     public static void displayValue(int num)
25     {
26         System.out.println("The value is " + num);
27     }
28 }

```

Program Output

```

I am passing values to displayValue.
The value is 5
The value is 10
The value is 40
The value is 700
That's all!

```

WARNING!

When passing a variable as an argument, simply write the variable name inside the parentheses of the method call. Do not write the data type of the argument variable in the method call. For example, the following statement will cause an error:

```
displayValue(int x); // Error!
```

The method call should appear as follows:

```
displayValue(x); // Correct
```

NOTE:

In this text, the values that are passed into a method are called arguments, and the variables that receive those values are called parameters. There are several variations of these terms in use. In

some circles these terms are switched in meaning. Also, some call the arguments *actual parameters* and call the parameters *formal parameters*. Others use the terms *actual argument* and *formal argument*. Regardless of which set of terms you use, it is important to be consistent.

Argument and Parameter Data Type Compatibility

When you pass an argument to a method, be sure that the argument's data type is compatible with the parameter variable's data type. Java will automatically perform a widening conversion if the argument's data type is ranked lower than the parameter variable's data type. For example, the `displayValue` method has an `int` parameter variable. Both of the following code segments will work because the `short` and `byte` arguments are automatically converted to an `int`:

```
short s = 1;
displayValue(s);    // Converts short to int

byte b = 2;
displayValue(b);    // Converts byte to int
```

However, Java will not automatically convert an argument to a lower-ranking data type. This means that a `long`, `float`, or `double` value cannot be passed to a method that has an `int` parameter variable. For example, the following code will cause a compiler error:

```
double d = 1.0;
displayValue(d);    // Error! Can't convert double to int.
```



You can use a cast operator to convert a value manually to a lower-ranking data type. For example, the following code will compile:

```
double d = 1.0;
displayValue((int)d);    // This will work.
```

Parameter Variable Scope

Recall from [Chapter 2](#) that a variable's scope is the part of the program where the variable may be accessed by its name. A variable is visible only to statements inside the variable's scope. A parameter variable's scope is the method in which the parameter is declared. No statement outside the method can access the parameter variable by its name.

Passing Multiple Arguments

Often it is useful to pass more than one argument to a method. Here is a method that accepts two arguments:

```
public static void showSum(double num1, double num2)
{
    double sum; // To hold the sum

    sum = num1 + num2;
    System.out.println("The sum is " + sum);
}
```

Notice that two parameter variables, `num1` and `num2`, are declared inside the parentheses in the method header. This is often referred to as a *parameter list*. Also notice that a comma separates the declarations. Here is an example of a statement that calls the method:

```
showSum(5, 10);
```

This statement passes the arguments 5 and 10 into the method. The arguments are passed into the parameter variables in the order that they appear in the method call. In other words, the first argument is passed into the first parameter variable, the second argument is passed into the second parameter variable, and so forth. So, this statement causes 5 to be passed into the `num1` parameter and 10 to be passed into the `num2` parameter. This is illustrated in [Figure 5-8](#).




Figure 5-8 Multiple arguments passed into multiple parameters

Suppose we were to reverse the order in which the arguments are listed in the method call, as shown here:

```
showSum(10, 5);
```

This would cause 10 to be passed into the `num1` parameter and 5 to be passed into the `num2` parameter. The following code segment shows one more example. This time we are passing variables as arguments.

```
double value1 = 2.5;
double value2 = 3.5;
showSum(value1, value2);
```

When the `showSum` methods executes as a result of this code, the `num1` parameter will contain 2.5 and the `num2` parameter will contain 3.5.



Each parameter variable in a parameter list must have a data type listed before its name. For example, a compiler error would occur if the parameter list for the `showSum` method were defined as shown in the following header:

```
public static void showSum(double num1, num2) // Error!
```

A data type for both the `num1` and `num2` parameter variables must be listed, as shown here:

```
public static void showSum(double num1, double num2)
```

See the program *TwoArgs.java* in this chapter's source code folder for a complete program that demonstrates the `showSum` method. You can download the book's source code from www.pearsonhighered.com/gaddis.

Arguments Are Passed by Value

In Java, all arguments of the primitive data types are *passed by value*, which means that only a copy of an argument's value is passed into a parameter variable. A method's parameter variables are separate and distinct from the arguments that are listed inside the parentheses of a method call. If a parameter variable is changed inside a method, it has no effect on the original argument. For example, look at the program in [Code Listing 5-6](#).

Code Listing 5-6 (PassByValue.java)

```

2   This program demonstrates that only a copy of an argument
3   is passed into a method.
4   */
5
6   public class PassByValue
7   {
8       public static void main(String[] args)
9       {
10          int number = 99; // number starts with 99
11
12          // Display the value in number.
13          System.out.println("number is " + number);
14
15          // Call changeMe, passing the value in number
16          // as an argument.
17          changeMe(number);
18
19          // Display the value in number again.
20          System.out.println("number is " + number);
21      }
22
23      /**
24       The changeMe method accepts an argument and then
25       changes the value of the parameter.
26      */
27
28      public static void changeMe(int myValue)
29      {
30          System.out.println("I am changing the value.");
31
32          // Change the myValue parameter variable to 0.
33          myValue = 0;
34
35          // Display the value in myValue.
36          System.out.println("Now the value is " + myValue);
37      }
38  }

```

Program Output

```

number is 99
I am changing the value.
Now the value is 0
number is 99

```

Even though the parameter variable `myValue` is changed in the `changeMe` method, the argument `number` is not modified. The `myValue` variable contains only a copy of the `number` variable.

Passing Object References to a Method

So far you've seen examples of methods that accept primitive values as arguments. You can also

write methods that accept references to objects as arguments. For example, look at the following method:

```
public static void showLength(String str)
{
    System.out.println(str + " is " + str.length() +
        " characters long.");
}
```

This method accepts a `String` object reference as its argument, and displays a message showing the number of characters in the object. The following code shows an example of how to call the method:

```
String name = "Warren";
showLength(name);
```

When this code executes, the `showLength` method will display the following:

```
Warren is 6 characters long.
```

When an object, such as a `String`, is passed as an argument, it is actually a reference to the object that is passed. In this example code, the `name` variable is a `String` reference variable. It is passed as an argument to the `showLength` method. The `showLength` method has a parameter variable, `str`, which is also a `String` reference variable, that receives the argument.

Recall that a reference variable holds the memory address of an object. When the `showLength` method is called, the address that is stored in `name` is passed into the `str` parameter variable. This is illustrated in [Figure 5-9](#). This means that when the `showLength` method is executing, both `name` and `str` reference the same object. This is illustrated in [Figure 5-10](#).




Figure 5-9 Passing a reference as an

argument

[Figure 5-9 Full Alternative Text](#)




Figure 5-10 Both name and str reference the same object

This might lead you to the conclusion that a method can change the contents of any `String` object that has been passed to it as an argument. After all, the parameter variable references the same object as the argument. However, `String` objects in Java are *immutable*, which means that they cannot be changed. For example, look at the program in [Code Listing 5-7](#). It passes a `String` object to a method, which appears to change the object. In reality, the object is not changed.

Code Listing 5-7 (PassString.java)

```
1  /**
2   * This program demonstrates that String arguments
3   * cannot be changed.
4   */
5
6  public class PassString
7  {
8      public static void main(String[] args)
9      {
10         // Create a String object containing "Shakespeare".
11         // The name variable references the object.
12         String name = "Shakespeare";
13
14         // Display the String referenced by the name variable.
15         System.out.println("In main, the name is " +
16                             name);
17
18         // Call the changeName method, passing the
19         // contents of the name variable as an argument.
```

```

20     changeName(name);
21
22     // Display the String referenced by the name variable.
23     System.out.println("Back in main, the name is " +
24         name);
25 }
26
27 /**
28     The changeName method accepts a String as its argument
29     and assigns the str parameter to a new String.
30 */
31
32 public static void changeName(String str)
33 {
34     // Create a String object containing "Dickens".
35     // Assign its address to the str parameter variable.
36     str = "Dickens";
37
38     // Display the String referenced by str.
39     System.out.println("In changeName, the name " +
40         "is now " + str);
41 }
42 }

```

Program Output

```

In main, the name is Shakespeare
In changeName, the name is now Dickens
Back in main, the name is Shakespeare

```

Let's take a closer look at this program. After line 12 executes, the `name` variable references a `String` object containing "Shakespeare". In line 20 the `changeName` method is called and the `name` variable is passed as an argument. This passes the address of the `String` object into the `str` parameter variable. At this point, both `name` and `str` reference the same object, as shown in [Figure 5-11](#).




Figure 5-11 Before line 36 executes, both `name` and `str` reference the same object

In the `changeName` method, line 36 executes as follows:

```
str = "Dickens";
```

At first, you might think that this statement changes the `String` object's contents to `"Dickens"`. What actually happens is that a new `String` object containing `"Dickens"` is created and its address is stored in the `str` variable. After this statement executes, the `name` variable and the `str` parameter variable reference different objects. This is shown in [Figure 5-12](#).

After Line 36 executes:




Figure 5-12 After line 36 executes, `name` and `str` reference different objects

In [Chapter 9](#) we will discuss the immutability of `String` objects in greater detail. Until then, just remember the following point: `String` objects cannot be changed. Any time you use the `=` operator to assign a string literal to a `String` reference variable, a new `String` object is created in memory.

Using the `@param` Tag in Documentation Comments

When writing the documentation comments for a method, you can provide a description of each parameter by using a `@param` tag. When the `javadoc` utility sees a `@param` tag inside a method's documentation comments, it knows that the documentation for a parameter variable appears next. The file `TwoArgs2.java`, in this chapter's source code (available at www.pearsonhighered.com/gaddis),

has the following method, which uses `@param` tags in its documentation comments:

```
/**
 * The showSum method displays the sum of two numbers.
 * @param num1 The first number.
 * @param num2 The second number.
 */

public static void showSum(double num1, double num2)
{
    double sum;    // To hold the sum

    sum = num1 + num2;
    System.out.println("The sum is " + sum);
}
```

The general format of a `@param` tag comment is as follows:

```
@param  parameterName Description
```

In the general format, `parameterName` is the name of the parameter and `Description` is a description of the parameter. Remember the following points about `@param` tag comments:

- All `@param` tags in a method's documentation comment must appear after the general description of the method.
- The description can span several lines. It ends at the end of the documentation comment (the `*/` symbol), or at the beginning of another tag.

When a method's documentation comments contain one or more `@param` tags, the `javadoc` utility will create a **Parameters** section in the method's documentation. This is where the descriptions of the method's parameters will be listed. [Figure 5-13](#) shows the documentation generated by `javadoc` for the `showSum` method in the *TwoArgs2.java* file.

showSum

```
public static void showSum(double num1,
                          double num2)
```

The showSum method displays the sum of two numbers.

Parameters:

num1 - The first number.

num2 - The second number.

Figure 5-13 Documentation for the showSum method in *TwoArgs2.java*



MyProgrammingLab™ www.myprogramminglab.com

1. 5.6 What is the difference between an argument and a parameter?
2. 5.7 Look at the following method header:

```
public static void myMethod(int num)
```

Which of the following calls to the method will cause a compiler error?

1. `myMethod(7);`
2. `myMethod(6.2);`
3. `long x = 99;`
`myMethod(x);`
4. `short s = 2;`
`myMethod(s);`
3. 5.8 Suppose a method named `showValues` accepts two `int` arguments. Which of the following method headers is written correctly?
 1. `public static void showValues()`
 2. `public static void showValues(int num1, num2)`
 3. `public static void showValues(num1, num2)`
 4. `public static void showValues(int num1, int num2)`
4. 5.9 In Java, method arguments are passed by value. What does this mean?
5. 5.10 What will the following program display?

```
public class Checkpoint
{
    public static void main(String[] args)
    {
        int num1 = 99;
        double num2 = 1.5;
```

```
        System.out.println(num1 + " " + num2);
        myMethod(num1, num2);
        System.out.println(num1 + " " + num2);
    }

    public static void myMethod(int i, double d)
    {
        System.out.println(i + " " + d);
        i = 0;
        d = 0.0;
        System.out.println(i + " " + d);
    }
}
```

5.3 More about Local Variables

CONCEPT: A local variable is declared inside a method and is not accessible to statements outside the method. Different methods can have local variables with the same names because the methods cannot see each other's local variables.

In [Chapter 2](#) we introduced the concept of local variables, which are variables that are declared inside a method. They are called *local* because they are local to the method in which they are declared. Statements outside a method cannot access that method's local variables.

Because a method's local variables are hidden from other methods, the other methods may have their own local variables with the same name. For example, look at the program in [Code Listing 5-8](#). In addition to the `main` method, this program has two other methods: `texas` and `california`. These two methods each have a local variable named `birds`.

Code Listing 5-8 (LocalVars.java)

```
1  /**
2      This program demonstrates that two methods may have
3      local variables with the same name.
4  */
5
6  public class LocalVars
7  {
8      public static void main(String[] args)
9      {
10         texas();
11         california();
12     }
13
14     /**
15         The texas method has a local variable named birds.
16     */
17
18     public static void texas()
19     {
20         int birds = 5000;
21
22         System.out.println("In texas there are " +
23                             birds + " birds.");
24     }
25
26     /**
27         The california method also has a local variable named birds.
28     */
29     public static void california()
30     {
31         int birds = 3500;
```



```

32
33         System.out.println("In california there are " +
34                               birds + " birds.");
35     }
36 }

```

Program Output

```

In texas there are 5000 birds.
In california there are 3500 birds.

```

Although there are two variables named `birds`, the program can see only one of them at a time because they are in different methods. When the `texas` method is executing, the `birds` variable declared inside `texas` is visible. When the `california` method is executing, the `birds` variable declared inside `california` is visible.

Local Variable Lifetime

A method's local variables exist only while the method is executing. This is known as the *lifetime* of a local variable. When the method begins, its local variables and its parameter variables are created in memory, and when the method ends, the local variables and parameter variables are destroyed. This means that any value stored in a local variable is lost between calls to the method in which the variable is declared.

Initializing Local Variables with Parameter Values

It is possible to use a parameter variable to initialize a local variable. Sometimes this simplifies the code in a method. For example, recall the following `showSum` method we discussed earlier:

```

public static void showSum(double num1, double num2)
{
    double sum;    // To hold the sum

    sum = num1 + num2;
    System.out.println("The sum is " + sum);
}

```

In the body of the method, the `sum` variable is declared and then a separate assignment statement assigns `num1 + num2` to `sum`. We can combine these statements into one, as shown in the following modified version of the method.

```

public static void showSum(double num1, double num2)
{

```

```
double sum = num1 + num2;  
System.out.println("The sum is " + sum);  
}
```

Because the scope of a parameter variable is the entire method in which it is declared, we can use parameter variables to initialize local variables.

WARNING!

Local variables are not automatically initialized with a default value. They must be given a value before they can be used. If you attempt to use a local variable before it has been given a value, a compiler error will result. For example, look at the following method:

```
public static void myMethod()  
{  
    int x;  
    System.out.println(x); //Error! x has no value.  
}
```

This code will cause a compiler error because the variable `x` has not been given a value, and it is being used as an argument to the `System.out.println` method.

5.4 Returning a Value from a Method

CONCEPT: A method may send a value back to the statement that called the method.

You've seen that data may be passed into a method by way of parameter variables. Data may also be returned from a method, back to the statement that called it. Methods that return a value are appropriately known as *value-returning methods*.



Returning a Value from a Method

You are already experienced at using value-returning methods. For instance, you have used the wrapper class parse methods, such as `Integer.parseInt`. Here is an example:

```
int num;
num = Integer.parseInt("700");
```

The second line in this code calls the `Integer.parseInt` method, passing "700" as the argument. The method returns the integer value 700, which is assigned to the `num` variable by the `=` operator. You have also seen the `Math.pow` method, which returns a value. Here is an example:

```
double x;
x = Math.pow(4.0, 2.0);
```

The second line in this code calls the `Math.pow` method, passing 4.0 and 2.0 as arguments. The method calculates the value of 4.0 raised to the power of 2.0 and returns that value. The value, which is 16.0, is assigned to the `x` variable by the `=` operator.

In this section, we will discuss how you can write your own value-returning methods.

Defining a Value-Returning Method

When you are writing a value-returning method, you must decide what type of value the method will return. This is because you must specify the data type of the return value in the method header. Recall that a `void` method, which does not return a value, uses the key word `void` as its return type in the method header. A value-returning method will use `int`, `double`, `boolean`, or any other valid data type in its header. Here is an example of a method that returns an `int` value:

```
public static int sum(int num1, int num2)
{
    int result;
```

```
result = num1 + num2;  
return result;  
}
```

The name of this method is `sum`. Notice in the method header that the return type is `int`, as shown in [Figure 5-14](#).

Return Type
↓
`public static int sum(int num1, int num2)`

Figure 5-14 Return type in the method header

This code defines a method named `sum` that accepts two `int` arguments. The arguments are passed into the parameter variables `num1` and `num2`. Inside the method, a local variable, `result`, is declared. The parameter variables `num1` and `num2` are added, and their sum is assigned to the `result` variable. The last statement in the method is as follows:

```
return result;
```

This is a `return` statement. You must have a `return` statement in a value-returning method. It causes the method to end execution and it returns a value to the statement that called the method. In a value-returning method, the general format of the `return` statement is as follows:

```
return Expression;
```

Expression is the value to be returned. It can be any expression that has a value, such as a variable, literal, or mathematical expression. In this case, the `sum` method returns the value in the `result` variable. However, we could have eliminated the `result` variable and returned the expression `num1 + num2`, as shown in the following code:

```
public static int sum(int num1, int num2)  
{  
    return num1 + num2;  
}
```



NOTE:

The `return` statement's expression must be of the same data type as the return type specified in the method header, or compatible with it. Otherwise, a compiler error will occur. Java will automatically widen the value of the `return` expression, if necessary, but it will not automatically narrow it.

Calling a Value-Returning Method

The program in [Code Listing 5-9](#) shows an example of how to call the `sum` method. Notice that the documentation comments for the `sum` method have a new tag, `@return`. This tag will be explained later.

Code Listing 5-9 (ValueReturn.java)

```
1  /**
2   * This program demonstrates a value-returning method.
3   */
4
5  public class ValueReturn
6  {
7      public static void main(String[] args)
8      {
9          int total, value1 = 20, value2 = 40;
10
11         // Call the sum method, passing the contents of
12         // value1 and value2 as arguments. Assign the
13         // return value to the total variable.
14         total = sum(value1, value2);
15
16         // Display the contents of all these variables.
17         System.out.println("The sum of " + value1 +
18                             " and " + value2 + " is " +
19                             total);
20     }
21
22     /**
23      * The sum method returns the sum of its two parameters.
24      * @param num1 The first number to be added.
25      * @param num2 The second number to be added.
26      * @return The sum of num1 and num2.
27      */
28
29     public static int sum(int num1, int num2)
30     {
31         int result; // result is a local variable
32
33         // Assign the value of num1 + num2 to result.
34         result = num1 + num2;
35
36         // Return the value in the result variable.
37         return result;
38     }
39 }
```

Program Output

The sum of 20 and 40 is 60

The statement in line 14 calls the `sum` method, passing `value1` and `value2` as arguments. It assigns the value returned by the `sum` method to the `total` variable. In this case, the method will return 60. [Figure 5-15](#) shows how the arguments are passed into the method and how a value is passed back from the method.




Figure 5-15 Arguments passed to `sum` and a value returned

[Figure 5-15 Full Alternative Text](#)

When you call a value-returning method, you usually want to do something meaningful with the value it returns. The *ValueReturn.java* program shows a method's return value being assigned to a variable. This is commonly how return values are used, but you can do many other things with them. For example, the following code shows a math expression that uses a call to the `sum` method:

```
int x = 10, y = 15;
double average;
average = sum(x, y) / 2.0;
```

In the last statement, the `sum` method is called with `x` and `y` as its arguments. The method's return value, which is 25, is divided by 2.0. The result, 12.5, is assigned to `average`. Here is another example:

```
int x = 10, y = 15;
System.out.println("The sum is " + sum(x, y));
```

This code sends the `sum` method's return value to `System.out.println`, so it can be displayed on the screen. The message "The sum is 25" will be displayed.

Remember, a value-returning method returns a value of a specific data type. You can use the method's return value anywhere that you can use a regular value of the same data type. This means that anywhere an `int` value can be used, a call to an `int` value-returning method can be used. Likewise,

anywhere a `double` value can be used, a call to a `double` value-returning method can be used. The same is true for all other data types.

Using the `@return` Tag in Documentation Comments

When writing the documentation comments for a value-returning method, you can provide a description of the return value by using a `@return` tag. When the `javadoc` utility sees a `@return` tag inside a method's documentation comments, it knows that a description of the method's return value appears next.

The general format of a `@return` tag comment is as follows:

```
@return Description
```

Description is a description of the return value. Remember the following points about `@return` tag comments:

- The `@return` tag in a method's documentation comment must appear after the general description of the method.
- The description can span several lines. It ends at the end of the documentation comment (the `*/` symbol), or at the beginning of another tag.

When a method's documentation comments contain a `@return` tag, the `javadoc` utility will create a Returns section in the method's documentation. This is where the description of the method's return value will be listed. [Figure 5-16](#) shows the documentation generated by `javadoc` for the `sum` method in the `ValueReturn.java` file.

```
sum
```

```
public static int sum(int num1,  
                     int num2)
```

```
The sum method returns the sum of its two parameters.
```

```
Parameters:
```

```
    num1 - The first number to be added.
```

```
    num2 - The second number to be added.
```

```
Returns:
```

```
    The sum of num1 and num2.
```

Figure 5-16 Documentation for the `sum`

method in `ValueReturn.java`



In the Spotlight:

Using Methods

Your friend Michael runs a catering company. Some of the ingredients that his recipes require are measured in cups. When he goes to the grocery store to buy those ingredients, however, they are sold only by the fluid ounce. He has asked you to write a simple program that converts cups to fluid ounces.

You design the following algorithm:

1. *Get the number of cups from the user.*
2. *Convert the number of cups to fluid ounces.*
3. *Display the result.*

This algorithm lists the top level of tasks that the program needs to perform, and becomes the basis of the class's `main` method. The class will also have the following methods:

- `getCups`—This method will prompt the user to enter the number of cups, and then return that value as a `double`.
- `cupsToOunces`—This method will accept the number of cups as an argument and then return the equivalent number of fluid ounces as a `double`.
- `displayResults`—This method displays a message indicating the results of the conversion.

[Code Listing 5-10](#) shows the program. [Figure 5-17](#) shows interaction with the program during execution.

Code Listing 5-10 (`CupConverter.java`)

```
1  import javax.swing.JOptionPane;
2
3  /**
4   * This program converts cups to fluid ounces.
5   */
```



```

6
7 public class CupConverter
8 {
9     public static void main(String[] args)
10    {
11        double cups;    // To hold the number of cups
12        double ounces; // To hold the number of ounces
13
14        // Get the number of cups.
15        cups = getCups();
16
17        // Convert the cups to fluid ounces.
18        ounces = cupsToOunces(cups);
19
20        // Display the results.
21        displayResults(cups, ounces);
22        System.exit(0);
23    }
24
25    /**
26     The getCups method prompts the user to enter a number
27     of cups.
28     @return The number of cups entered by the user.
29 */
30
31    public static double getCups()
32    {
33        String input;    // To hold input.
34        double numCups;  // To hold cups.
35
36        // Get the number of cups from the user.
37        input = JOptionPane.showInputDialog(
38            "This program converts measurements\n" +
39            "in cups to fluid ounces. For your\n" +
40            "reference the formula is:\n" +
41            "    1 cup = 8 fluid ounces\n\n" +
42            "Enter the number of cups.");
43
44        // Convert the input to a double.
45        numCups = Double.parseDouble(input);
46
47        // Return the number of cups.
48        return numCups;
49    }
50
51    /**
52     The cupsToOunces method converts a number of
53     cups to fluid ounces, using the formula
54     1 cup = 8 fluid ounces.
55     @param numCups The number of cups to convert.
56     @return The number of ounces.
57 */
58
59    public static double cupsToOunces(double numCups)
60    {
61        return numCups * 8.0;
62    }


```

```

63
64  /**
65     The displayResults method displays a message showing
66     the results of the conversion.
67     @param cups A number of cups.
68     @param ounces A number of ounces.
69  */
70
71  public static void displayResults(double cups, double ounces)
72  {
73      // Display the number of ounces.
74      JOptionPane.showMessageDialog(null,
75          cups + " cups equals " +
76          ounces + " fluid ounces.");
77  }
78  }

```

This input dialog appears first.
The user enters 2 and clicks OK.



This message dialog shows
the results.




Figure 5-17 Interaction with the CupConverter program

(Oracle Corporate Counsel)

[Figure 5-17 Full Alternative Text](#)

Returning a boolean Value

Frequently there is a need for a method that tests an argument and returns a `true` or `false` value indicating whether or not a condition exists. Such a method would return a `boolean` value. For example, the following method accepts an argument and returns `true` if the argument is within the range of 1 through 100, or `false` otherwise:

```

public static boolean isValid(int number)
{
    boolean status;

    if (number >= 1 && number <= 100)
        status = true;
    else
        status = false;
    return status;
}

```

The following code shows an if-else statement that uses a call to the method:

```

int value = 20;
if (isValid(value))
    System.out.println("The value is within range.");
else
    System.out.println("The value is out of range.");

```

When this code executes, the message “The value is within range.” will be displayed.

Returning a Reference to an Object

A value-returning method can also return a reference to a non-primitive type, such as a `String` object. The program in [Code Listing 5-11](#) shows such an example.

Code Listing 5-11 (ReturnString.java)

```

1  /**
2      This program demonstrates a method that
3      returns a reference to a String object.
4  */
5
6  public class ReturnString
7  {
8      public static void main(String[] args)
9      {
10         String customerName;
11
12         customerName = fullName("John", "Martin");
13         System.out.println(customerName);
14     }
15
16     /**
17         The fullName method accepts two String arguments
18         containing a first and last name. It concatenates
19         them into a single String object.
20         @param first The first name.
21         @param last The last name.
22         @return A reference to a String object containing
23                 the first and last names.

```

```

24     */
25
26     public static String fullName(String first, String last)
27     {
28         String name;
29
30         name = first + " " + last;
31         return name;
32     }
33 }

```

Program Output

John Martin

Line 12 calls the `fullName` method, passing "John" and "Martin" as arguments. The method returns a reference to a `String` object containing "John Martin". The reference is assigned to the `customerName` variable. This is illustrated in [Figure 5-18](#).




Figure 5-18 The `fullName` method returning a reference to a `String` object

[Figure 5-18 Full Alternative Text](#)



MyProgrammingLab™ www.myprogramminglab.com

1. 5.11 Look at the following method header. What type of value does the method return?

```
public static double getValue(int a, float b, String c)
```

2. 5.12 Write the header for a method named `days`. The method should return an `int` and have three `int` parameter variables: `years`, `months`, and `weeks`.
3. 5.13 Write the header for a method named `distance`. The method should return a `double` and have two `double` parameter variables: `rate` and `time`.
4. 5.14 Write the header for a method named `lightYears`. The method should return a `long` and have one `long` parameter variable: `miles`.

5.5 Problem Solving with Methods

CONCEPT: A large, complex problem can be solved a piece at a time by methods.

At the beginning of this chapter we introduced the idea of using methods to “divide and conquer” a problem. Often the best way to solve a complex problem is to break it down into smaller problems, and then solve the smaller problems. The process of breaking down a problem into smaller pieces is called *functional decomposition*.

In functional decomposition, instead of writing one long method that contains all of the statements necessary to solve a problem, small methods are written, which each solve a specific part of the problem. These small methods can then be executed in the desired order to solve the problem.

Let’s look at an example. The program in [Code Listing 5-12](#) reads 30 days of sales amounts from a file, and then displays the total sales and average daily sales. Here’s a brief pseudocode model of the algorithm:

Code Listing 5-12 (SalesReport.java)

```
1  import java.util.Scanner;           // For the Scanner class
2  import java.io.*;                  // For file I/O classes
3  import javax.swing.JOptionPane;     // For the JOptionPane class
4
5  /**
6   * This program opens a file containing the sales
7   * amounts for 30 days. It calculates and displays
8   * the total sales and average daily sales.
9   */
10
11 public class SalesReport
12 {
13     public static void main(String[] args) throws IOException
14     {
15         final int NUM_DAYS = 30;     // Number of days of sales
16         String filename;              // The name of the file to open
17         double totalSales;            // Total sales for period
18         double averageSales;         // Average daily sales
19
20         // Get the name of the file.
21         filename = getFileName();
22
23         // Get the total sales from the file.
24         totalSales = getTotalSales(filename);
25
26         // Calculate the average.
27         averageSales = totalSales / NUM_DAYS;
28
29         // Display the total and average.
```

```

30         displayResults(totalSales, averageSales);
31
32         System.exit(0);
33     }
34
35     /**
36      * The getFileName method prompts the user to enter
37      * the name of the file to open.
38      * @return A reference to a String object containing
39      *         the name of the file.
40      */
41
42     public static String getFileName()
43     {
44         String file;    // To hold the file name
45
46         // Prompt the user to enter a file name.
47         file = JOptionPane.showInputDialog(
48             "Enter the name of the file\n" +
49             "containing 30 days of sales amounts.");
50
51         // Return the name.
52         return file;
53     }
54
55     /**
56      * The getTotalSales method opens a file and
57      * reads the daily sales amounts, accumulating
58      * the total. The total is returned.
59      * @param filename The name of the file to open.
60      * @return The total of the sales amounts.
61      */
62
63     public static double getTotalSales(String filename) throws IOException
64     {
65         double total = 0.0; // Accumulator
66         double sales; // A daily sales amount
67
68         // Open the file.
69         File file = new File(filename);
70         Scanner inputFile = new Scanner(file);
71
72         // This loop processes the lines read from the file,
73         // until the end of the file is encountered.
74         while (inputFile.hasNext())
75         {
76             // Read a double from the file.
77             sales = inputFile.nextDouble();
78
79             // Add sales to the value already in total.
80             total += sales;
81         }
82
83         // Close the file.
84         inputFile.close();
85
86         // Return the total sales.

```

```

87         return total;
88     }
89
90     /**
91      * The displayResults method displays the total and
92      * average daily sales.
93      * @param total The total sales.
94      * @param avg The average daily sales.
95      */
96
97     public static void displayResults(double total, double avg)
98     {
99         // Display the formatted total and average sales.
100        JOptionPane.showMessageDialog(null,
101            String.format("The total sales for the period is $%,.2f\n" +
102                "The average daily sales were $%,.2f",
103                total, avg));
104    }
105 }

```

- *Ask the user to enter the name of the file.*
- *Get the total of the sales amounts in the file.*
- *Calculate the average daily sales.*
- *Display the total and average daily sales.*

The file *MonthlySales.txt*, in this chapter's source code (available at www.pearsonhighered.com/gaddis), is used to test the program. [Figure 5-19](#) shows interaction with the program during execution.




Figure 5-19 Interaction with the SalesReport.java program

(Oracle Corporate Counsel)

[Figure 5-19 Full Alternative Text](#)

Instead of writing the entire program in the `main` method, the algorithm was broken down into the following methods:

- `getFileName`—This method displays an input dialog box asking the user to enter the name of the file containing 30 days of sales amounts. The method returns a reference to a `String` object containing the name entered by the user.
- `getTotalSales`—This method accepts the name of a file as an argument. The file is opened, the sales amounts are read from it, and the total of the sales amounts is accumulated. The method returns the total as a `double`.
- `displayResults`—This method accepts as arguments the total sales and the average daily sales. It displays a message dialog box indicating these values.

Calling Methods That Throw Exceptions

One last thing about the *SalesReport.java* program should be discussed. Notice that the `main` method header (in line 13) and the `getTotalSales` method header (in line 63) both have a `throws IOException` clause. The `getTotalSales` method has the clause because it uses a `Scanner` object to open a file. As you know from [Chapter 4](#), any method that uses a `Scanner` object to open a file should have a `throws IOException` clause in its header. Let's quickly review why this is so.

When a `Scanner` object has a problem opening a file, it throws an exception known as `IOException`. Java requires that either (a) the exception is handled in the method that caused it to occur, or (b) the method terminates and throws the exception again. For now you must write your methods to throw the exception again because you will not learn how to handle exceptions until [Chapter 11](#). By writing a `throws IOException` clause in a method's header, you are telling the compiler that the method does not handle the exception. Instead, it throws the exception again.

That explains why the `getTotalSales` method has the `throws IOException` clause, but it doesn't explain why the `main` method has one. After all, `main` doesn't use a `Scanner` object to perform any file operations. The reason `main` has to have the clause is because `main` calls the `getTotalSales` method. If the `Scanner` object in `getTotalSales` throws an `IOException`, the `getTotalSales` method terminates and throws the `IOException` again. That means that `main` must either handle the exception, or terminate and throw it once again. When the `main` method throws the exception, the JVM displays an error message on the screen.



Until you learn how to handle exceptions in [Chapter 11](#), just remember this when writing programs that throw exceptions: If a method calls another method that has a `throws` clause in its header, then the calling method should have the same `throws` clause.

5.6 Common Errors to Avoid

- Putting a semicolon at the end of a method header. Method headers are never terminated with a semicolon.
- Writing modifiers or return types in a method call statement. Method modifiers and return types are written in method headers, but never in method calls.
- Forgetting to write the empty parentheses in a call to a method that accepts no arguments. You must always write the parentheses in a method call statement, even if the method doesn't accept arguments.
- Forgetting to pass arguments to methods that require them. If a method has parameter variables, you must provide arguments when calling the method.
- Passing an argument of a data type that cannot be automatically converted to the data type of the parameter variable. Java will automatically perform a widening conversion if the argument's data type is ranked lower than the parameter variable's data type. But Java will not automatically convert an argument to a lower-ranking data type.
- Attempting to access a parameter variable with code outside the method where the variable is declared. A parameter variable is visible only within the method it is declared in.
- Not writing the data type of each parameter variable in a method header. Each parameter variable declaration inside the parentheses of a method header must include the variable's data type.
- Changing the contents of a method's parameter variable and expecting the argument that was passed into the parameter to change as well. Method arguments are passed by value, which means that a copy of the argument is passed into a parameter variable. Changes to the parameter variable have no effect on the argument.
- Using a variable to receive a method's return value when the variable's data type is incompatible with the data type of the return value. A variable that receives a method's return value must be of a data type that is compatible with the data type of the return value.
- Not writing a `return` statement in a value-returning method. If a method's return type is anything other than `void`, it should return a value.
- Not writing a required `throws` clause in a method that calls another method. Any method that calls a method with a `throws` clause in its header must either handle the potential exception or have the same `throws` clause. You will learn how to handle exceptions in [Chapter 11](#).

Review Questions and Exercises

Multiple Choice and True/False

1. This type of method does not return a value.

1. `null`
2. `void`
3. `empty`
4. `anonymous`

2. This appears at the beginning of a method definition.

1. semicolon
2. parentheses
3. body
4. header

3. The body of a method is enclosed in .

1. curly braces { }
2. square brackets []
3. parentheses ()
4. quotation marks " "

4. A method header can contain .

1. method modifiers
2. the method return type
3. the method name
4. a list of parameter declarations

5. all of these

6. none of these

5. A value that is passed into a method when it is called is known as a(n) .

1. parameter

2. argument

3. signal

4. return value

6. A variable that receives a value that is passed into a method is known as a(n) .

1. parameter

2. argument

3. signal

4. return value

7. This `javadoc` tag is used to document a parameter variable.

1. `@parameter`

2. `@param`

3. `@paramvar`

4. `@arg`

8. This statement causes a method to end and sends a value back to the statement that called the method.

1. `end`

2. `send`

3. `exit`

4. `return`

9. This `javadoc` tag is used to document a method's return value.

1. `@methodreturn`

2. `@ret`
 3. `@return`
 4. `@returnval`
10. **True or False:** You terminate a method header with a semicolon.
 11. **True or False:** When passing an argument to a method, Java will automatically perform a widening conversion (convert the argument to a higher-ranking data type), if necessary.
 12. **True or False:** When passing an argument to a method, Java will automatically perform a narrowing conversion (convert the argument to a lower-ranking data type), if necessary.
 13. **True or False:** A parameter variable's scope is the entire program that contains the method in which the parameter is declared.
 14. **True or False:** When code in a method changes the value of a parameter, it also changes the value of the argument that was passed into the parameter.
 15. **True or False:** When an object, such as a `String`, is passed as an argument, it is actually a reference to the object that is passed.
 16. **True or False:** The contents of a `String` object cannot be changed.
 17. **True or False:** When passing multiple arguments to a method, the order in which the arguments are passed is not important.
 18. **True or False:** No two methods in the same program can have a local variable with the same name.
 19. **True or False:** It is possible for one method to access a local variable that is declared in another method.
 20. **True or False:** You must have a `return` statement in a value-returning method.

Find the Error

1. Find the error in the following method definition:

```
// This method has an error!
public static void sayHello();
{
    System.out.println("Hello");
}
```

2. Look at the following method header:

```
public static void showValue(int x)
```

The following code has a call to the `showValue` method. Find the error.

```
int x = 8;
showValue(int x);    // Error!
```

3. Find the error in the following method definition:

```
// This method has an error!
public static double timesTwo(double num)
{
    double result = num * 2;
}
```

4. Find the error in the following method definition:

```
// This method has an error!
public static int half(double num)
{
    double result = num / 2.0;
    return result;
}
```

Algorithm Workbench

1. Examine the following method header, and then write an example call to the method:

```
public static void doSomething(int x)
```

2. Here is the code for the `displayValue` method, shown earlier in this chapter:

```
public static void displayValue(int num)
{
    System.out.println("The value is " + num);
}
```

For each of the following code segments, indicate whether it will successfully compile or cause an error:

1. `displayValue(100);`

2. `displayValue(6.0);`

3. `short s = 5;`

4. `displayValue(s);`

5. `long num = 1;`
6. `displayValue(num);`
7. `displayValue(6.2f);`
8. `displayValue((int) 7.5);`

3. Look at the following method header:

```
public static void myMethod(int a, int b, int c)
```

Now look at the following call to `myMethod`:

```
myMethod(3, 2, 1);
```

When this call executes, what value will be stored in `a`? What value will be stored in `b`? What value will be stored in `c`?

4. What will the following program display?

```
public class ChangeParam
{
    public static void main(String[] args)
    {
        int x = 1;
        double y = 3.4;
        System.out.println(x + " " + y);
        changeUs(x, y);
        System.out.println(x + " " + y);

    }
    public static void changeUs(int a, double b)
    {
        a = 0;
        b = 0.0;
        System.out.println(a + " " + b);
    }
}
```

5. A program contains the following method definition:

```
public static int cube(int num)
{
    return num * num * num;
}
```

Write a statement that passes the value 4 to this method and assigns its return value to a variable named `result`.

6. A program contains the following method:

```
public static void display(int arg1, double arg2, char arg3)
{
    System.out.println("The values are " + arg1 + ", " +
                        arg2 + ", and " + arg3);
}
```

Write a statement that calls this method and passes the following variables as arguments:

```
char initial = 'T';
int age = 25;
double income = 50000.00;
```

7. Write a method named `timesTen`. The method should accept a `double` argument, and return a `double` value that is ten times the value of the argument.
8. Write a method named `square` that accepts an integer argument and returns the square of that argument.
9. Write a method named `getName` that prompts the user to enter his or her first name, and then returns the user's input.
10. Write a method named `quartersToDollars`. The method should accept an `int` argument that is a number of quarters, and return the equivalent number of dollars as a `double`. For example, if you pass 4 as an argument, the method should return 1.0; and if you pass 7 as an argument, the method should return 1.75.

Short Answer

1. What is the “divide and conquer” approach to problem solving?
2. What is the difference between a void method and a value-returning method?
3. What is the difference between an argument and a parameter variable?
4. Where do you declare a parameter variable?
5. Explain what is meant by the phrase “pass by value.”
6. Why do local variables lose their values between calls to the method in which they are declared?

Programming Challenges

MyProgrammingLab™ Visit www.myprogramminglab.com to complete many of these Programming Challenges online and get instant feedback.

1. showChar Method

Write a method named `showChar`. The method should accept two arguments: a reference to a `String` object and an integer. The integer argument is a character position within the `String`, with the first character being at position 0. When the method executes, it should display the character at that character position. Here is an example of a call to the method:

```
showChar("New York", 2);
```

In this call, the method will display the character `w` because it is in position 2. Demonstrate the method in a complete program.

2. Retail Price Calculator



The Retail Price Calculator Problem

Write a program that asks the user to enter an item's wholesale cost and its markup percentage. It should then display the item's retail price. For example:

- If an item's wholesale cost is 5.00 and its markup percentage is 100 percent, then the item's retail price is 10.00.
- If an item's wholesale cost is 5.00 and its markup percentage is 50 percent, then the item's retail price is 7.50.

The program should have a method named `calculateRetail` that receives the wholesale cost and the markup percentage as arguments, and returns the retail price of the item.

3. Rectangle Area—Complete the Program

If you have downloaded the book's source code from www.pearsonhighered.com/gaddis, you will

find a partially written program named *AreaRectangle.java* in this chapter's source code folder. Your job is to complete the program. When it is complete, the program will ask the user to enter the width and length of a rectangle, and then display the rectangle's area. The program calls the following methods, which have not been written:

- `getLength`—This method should ask the user to enter the rectangle's length, and then return that value as a `double`.
- `getWidth`—This method should ask the user to enter the rectangle's width, and then return that value as a `double`.
- `getArea`—This method should accept the rectangle's length and width as arguments, and return the rectangle's area. The area is calculated by multiplying the length by the width.
- `displayData`—This method should accept the rectangle's length, width, and area as arguments, and display them in an appropriate message on the screen.

4. Paint Job Estimator

A painting company has determined that for every 115 square feet of wall space, one gallon of paint and eight hours of labor will be required. The company charges \$18.00 per hour for labor. Write a program that allows the user to enter the number of rooms to be painted and the price of the paint per gallon. It should also ask for the square feet of wall space in each room. The program should have methods that return the following data:

- The number of gallons of paint required
- The hours of labor required
- The cost of the paint
- The labor charges
- The total cost of the paint job

Then it should display the data on the screen.

5. Falling Distance

When an object is falling because of gravity, the following formula can be used to determine the distance the object falls in a specific time period:

$$d = \frac{1}{2} g t^2$$

The variables in the formula are as follows: d is the distance in meters, g is 9.8, and t is the amount of time, in seconds, that the object has been falling.

Write a method named `fallingDistance` that accepts an object's falling time (in seconds) as an argument. The method should return the distance, in meters, that the object has fallen during that time interval. Demonstrate the method by calling it in a loop that passes the values 1 through 10 as arguments, and displays the return value.

6. Celsius Temperature Table

The formula for converting a temperature from Fahrenheit to Celsius is

$$C = 5/9 (F - 32)$$

where F is the Fahrenheit temperature and C is the Celsius temperature. Write a method named `celsius` that accepts a Fahrenheit temperature as an argument. The method should return the temperature, converted to Celsius. Demonstrate the method by calling it in a loop that displays a table of the Fahrenheit temperatures 0 through 20 and their Celsius equivalents.

7. Test Average and Grade

Write a program that asks the user to enter five test scores. The program should display a letter grade for each score and the average test score. Write the following methods in the program:

- `calcAverage`—This method should accept five test scores as arguments and return the average of the scores.
- `determineGrade`—This method should accept a test score as an argument and return a letter grade for the score, based on the following grading scale:

Score	Letter Grade
90–100	A
80–89	B
70–79	C

Below 60 F

8. Conversion Program

Write a program that asks the user to enter a distance in meters. The program will then pre-sent the following menu of selections:

1. Convert to kilometers
2. Convert to inches
3. Convert to feet
4. Quit the program

The program will convert the distance to kilometers, inches, or feet, depending on the user's selection. Here are the specific requirements:

- Write a `void` method named `showKilometers`, which accepts the number of meters as an argument. The method should display the argument converted to kilometers. Convert the meters to kilometers using the following formula:


```
kilometers = meters * 0.001
```
- Write a `void` method named `showInches`, which accepts the number of meters as an argument. The method should display the argument converted to inches. Convert the meters to inches using the following formula:


```
inches = meters * 39.37
```
- Write a `void` method named `showFeet`, which accepts the number of meters as an argument. The method should display the argument converted to feet. Convert the meters to feet using the following formula:


```
feet = meters * 3.281
```
- Write a `void` method named `menu` that displays the menu of selections. This method should not accept any arguments.
- The program should continue to display the menu until the user enters 4 to quit the program.
- The program should not accept negative numbers for the distance in meters.

- If the user selects an invalid choice from the menu, the program should display an error message.

Here is an example session with the program, using console input. The user's input is shown in bold.

```
Enter a distance in meters: 500 [Enter]
```

```
1. Convert to kilometers
2. Convert to inches
3. Convert to feet
4. Quit the program
```

```
Enter your choice: 1 [Enter]
500 meters is 0.5 kilometers.
```

```
1. Convert to kilometers
2. Convert to inches
3. Convert to feet
4. Quit the program
```

```
Enter your choice: 3 [Enter]
500 meters is 1640.5 feet.
```

```
1. Convert to kilometers
2. Convert to inches
3. Convert to feet
4. Quit the program
```

```
Enter your choice: 4 [Enter]
Bye!
```

9. Distance Traveled Modification

The distance a vehicle travels can be calculated as follows:

`Distance = Speed * Time`

Write a method named `distance` that accepts a vehicle's speed and time as arguments, and returns the distance the vehicle has traveled. Modify the "Distance Traveled" program you wrote in [Chapter 4](#) (Programming Challenge 2) to use the method.

10. Stock Profit

The profit from the sale of a stock can be calculated as follows:

`Profit = ((NS * SP) - SC) - ((NS * PP) + PC)`

where *NS* is the number of shares, *PP* is the purchase price per share, *PC* is the purchase commission paid, *SP* is the sale price per share, and *SC* is the sale commission paid. If the calculation yields a positive value, then the sale of the stock resulted in a profit. If the calculation yields a negative

number, then the sale resulted in a loss.

Write a method that accepts as arguments the number of shares, the purchase price per share, the purchase commission paid, the sale price per share, and the sale commission paid. The method should return the profit (or loss) from the sale of stock. Demonstrate the method in a program that asks the user to enter the necessary data and displays the amount of the profit or loss.

11. Multiple Stock Sales

Use the method that you wrote for Programming Challenge 10 (Stock Profit) in a program that calculates the total profit or loss from the sale of multiple stocks. The program should ask the user for the number of stock sales, and the necessary data for each stock sale. It should accumulate the profit or loss for each stock sale and then display the total.

12. Kinetic Energy

In physics, an object that is in motion is said to have kinetic energy. The following formula can be used to determine a moving object's kinetic energy:

$$KE = \frac{1}{2} m v^2$$

The variables in the formula are as follows: *KE* is the kinetic energy, *m* is the object's mass in kilograms, and *v* is the object's velocity, in meters per second.

Write a method named `kineticEnergy` that accepts an object's mass (in kilograms) and velocity (in meters per second) as arguments. The method should return the amount of kinetic energy that the object has. Demonstrate the method by calling it in a program that asks the user to enter values for mass and velocity.

13. isPrime Method

A prime number is a number that is evenly divisible only by itself and 1. For example, the number 5 is prime because it can be evenly divided only by 1 and 5. The number 6, however, is not prime because it can be divided evenly by 1, 2, 3, and 6.

Write a method named `isPrime`, which takes an integer as an argument and returns `true` if the argument is a prime number, or `false` otherwise. Demonstrate the method in a complete program.



Recall that the `%` operator divides one number by another, and returns the remainder of the division. In an expression such as `num1 % num2`, the `%` operator will return 0 if `num1` is evenly divisible by `num2`.

14. Prime Number List

Use the `isPrime` method that you wrote in Programming Challenge 13 in a program that stores a list of all the prime numbers from 1 through 100 in a file.

15. Even/Odd Counter

You can use the following logic to determine whether a number is even or odd:

```
if ((number % 2) == 0)
{
    // The number is even.
}
else
{
    // The number is odd.
}
```

Write a program with a method named `isEven` that accepts an `int` argument. The method should return `true` if the argument is even, or `false` otherwise. The program's `main` method should use a loop to generate 100 random integers. It should use the `isEven` method to determine whether each random number is even, or odd. When the loop is finished, the program should display the number of even numbers that were generated, and the number of odd numbers.

16. Present Value

Suppose you want to deposit a certain amount of money into a savings account, and then leave it alone to draw interest for the next 10 years. At the end of 10 years, you would like to have \$10,000 in the account. How much do you need to deposit today to make that happen? You can use the following formula, which is known as the present value formula, to find out:

$$P = F (1 + r)^{-n}$$

The terms in the formula are as follows:

- P is the **present value**, or the amount that you need to deposit today.
- F is the **future value** that you want in the account. (In this case, F is \$10,000.)
- r is the **annual interest rate**.

- n is the **number of years** that you plan to let the money sit in the account.

Write a method named `presentValue` that performs this calculation. The method should accept the future value, annual interest rate, and number of years as arguments. It should return the present value, which is the amount that you need to deposit today. Demonstrate the method in a program that lets the user experiment with different values for the formula's terms.

17. Rock, Paper, Scissors Game

Write a program that lets the user play the game of Rock, Paper, Scissors against the computer. The program should work as follows.

1. When the program begins, a random number in the range of 1 through 3 is generated. If the number is 1, then the computer has chosen rock. If the number is 2, then the computer has chosen paper. If the number is 3, then the computer has chosen scissors. (Don't display the computer's choice yet.)
2. The user enters his or her choice of "rock", "paper", or "scissors" at the keyboard. (You can use a menu if you prefer.)
3. The computer's choice is displayed.
4. A winner is selected according to the following rules:
 - If one player chooses rock and the other player chooses scissors, then rock wins. (The rock smashes the scissors.)
 - If one player chooses scissors and the other player chooses paper, then scissors wins. (Scissors cuts paper.)
 - If one player chooses paper and the other player chooses rock, then paper wins. (Paper wraps rock.)
 - If both players make the same choice, the game must be played again to determine the winner.

Be sure to divide the program into methods that perform each major task.

18. ESP Game

Write a program that tests your ESP (extrasensory perception). The program should randomly select the name of a color from the following list of words:

- *Red, Green, Blue, Orange, Yellow*

To select a word, the program can generate a random number. For example, if the number is 0, the selected word is *Red*; if the number is 1, the selected word is *Green*; and so forth.

Next, the program should ask the user to enter the color that the computer has selected. After the user has entered his or her guess, the program should display the name of the randomly selected color. The program should repeat this 10 times and then display the number of times the user correctly guessed the selected color. Be sure to modularize the program into methods that perform each major task.

Chapter 6 A First Look at Classes

TOPICS

1. [6.1 Objects and Classes](#)
2. [6.2 Writing a Simple Class, Step by Step](#)
3. [6.3 Instance Fields and Methods](#)
4. [6.4 Constructors](#)
5. [6.5 Passing Objects as Arguments](#)
6. [6.6 Overloading Methods and Constructors](#)
7. [6.7 Scope of Instance Fields](#)
8. [6.8 Packages and `import` Statements](#)
9. [6.9 Focus on Object-Oriented Design: Finding the Classes and Their Responsibilities](#)

On the Web: Case Study—The `Amortization` Class
10. [6.10 Common Errors to Avoid](#)

6.1 Objects and Classes

CONCEPT: An object is a software component that exists in memory and serves a specific purpose in a program. An object is created from a class that contains code describing the object.

If you have ever driven a car, you know that a car consists of a lot of components. It has a steering wheel, an accelerator pedal, a brake pedal, a gear shifter, a speedometer, and numerous other devices that the driver interacts with. There are also a lot of components under the hood, such as the engine, the battery, the radiator, and so forth. So, a car is not just one single object, but rather a collection of objects that work together.

This same notion applies to computer programming as well. Most programming languages in use today are object-oriented. With an object-oriented language, such as Java, you create programs that are made of objects. In programming, however, an object isn't a physical device, like a steering wheel or a brake pedal; it's a software component that exists in the computer's memory and performs a specific task. In software, an object has two general capabilities:

- An object can store data. The data stored in an object are commonly called *fields*.
- An object can perform operations. The operations that an object can perform are called *methods*.

Objects are very important in Java. Here are some examples of objects that you have previously learned about:

- If you need to read input from the keyboard, or from a file, you can use a `Scanner` object.
- If you need to generate random numbers, you can use a `Random` object.
- If you need to write output to a file, you can use a `PrintWriter` object.

When a program needs the services of a particular type of object, it creates that object in memory, and then calls that object's methods as necessary.

Classes: Where Objects Come From

Objects are very useful, but they don't just magically appear in your program. Before a specific type of object can be used by a program, that object has to be created in memory. And, before an object can be created in memory, you must have a class for the object.

A *class* is code that describes a particular type of object. It specifies the data that an object can hold (the object's fields), and the actions that an object can perform (the object's methods). You can think

of a class as a code “blueprint” that can be used to create a particular type of object. It serves a purpose similar to that of the blueprint for a house. The blueprint itself is not a house, but rather a detailed description of a house. When we use the blueprint to build an actual house, we could say we are building an instance of the house described by the blueprint. If we so desire, we can build several identical houses from the same blueprint. Each house is a separate instance of the house described by the blueprint. This idea is illustrated in [Figure 6-1](#).




Figure 6-1 A blueprint and houses built from the blueprint

So, a class is not an object, but a description of an object. When a program is running, it can use the class to create, in memory, as many objects of a specific type as needed. Each object that is created from a class is called an *instance* of the class.

NOTE:

Up to this chapter, you have used classes for a different purpose: as containers for a program’s methods. All of the Java programs that you have written so far have had a class containing a `main` method, and possibly other methods. In this chapter, you will learn how to write classes from which objects can be created.

Classes in the Java API

So far, the objects that you have used in your programs are created from classes in the Java API. For example, each time you create a `Scanner` object, you are creating an instance of a class named `Scanner`, which is in the Java API. Likewise, when you create a `Random` object, you are creating an instance of a class named `Random`, which is in the Java API. The same is true for `PrintWriter` objects. When you need to write data to a file, you create an instance of the `PrintWriter` class, which is in the Java API. Look at [Code Listing 6-1](#), a program that uses all of these types of objects.

Code Listing 6-1 (ObjectDemo.java)

```
1  import java.util.Scanner; // Needed for the Scanner class
2  import java.util.Random;  // Needed for the Random class
3  import java.io.*;         // Needed for file I/O classes
4
5  /**
6   * This program writes random numbers to a file.
7   */
8
9  public class ObjectDemo
10 {
11     public static void main(String[] args) throws IOException
12     {
13         int maxNumbers;    // Max number of random numbers
14         int number;        // To hold a random number
15
16         // Create a Scanner object for keyboard input.
17         Scanner keyboard = new Scanner(System.in);
18
19         // Create a Random object to generate random numbers.
20         Random rand = new Random();
21
22         // Create a PrintWriter object to open the file.
23         PrintWriter outputFile = new PrintWriter("numbers.txt");
24
25         // Get the number of random numbers to write.
26         System.out.print("How many random numbers should I write? ");
27         maxNumbers = keyboard.nextInt();
28
29         // Write the random numbers to the file.
30         for (int count = 0; count < maxNumbers; count++)
31         {
32             // Generate a random integer.
33             number = rand.nextInt();
34
35             // Write the random integer to the file.
36             outputFile.println(number);
37         }
38
39         // Close the file.
40         outputFile.close();
41         System.out.println("Done");
42     }
43 }
```

Program Output with Example Input Shown in Bold

How many random numbers should I write? **10** [Enter]
Done

In a nutshell, this program writes a specified number of random numbers to a file named `numbers.txt`. When the program runs, it asks the user for the number of random numbers to write. It then writes that many numbers to the file. To do its job, it creates three objects:

- In line 17 it creates an instance of the `Scanner` class, and assigns the object's address to a variable named `keyboard`. The object will be used to read keyboard input.
- In line 20 it creates an instance of the `Random` class, and assigns the object's address to a variable named `rand`. The object will be used to generate random numbers.
- In line 23 it creates an instance of the `PrintWriter` class, and assigns the object's address to a variable named `outputFile`. The object will be used to write output to the `numbers.txt` file.

[Figure 6-2](#) illustrates the three objects that the program creates. As the program runs, it uses these objects to accomplish certain tasks. For example:




Figure 6-2 Objects created by the ObjectDemo program

[Figure 6-2 Full Alternative Text](#)

- In line 27 the `Scanner` object's `nextInt` method is called to read the user's input (which is the number of random numbers to generate). The value that is returned from the method is assigned to the `maxNumbers` variable.

- In line 33 the `Random` object's `nextInt` method is called to get a random integer. The value that is returned from the method is assigned to the `number` variable.
- In line 36 the `PrintWriter` object's `println` method is called to write the value of the `number` variable to the file.
- In line 40 the `PrintWriter` object's `close` method is called to close the file.

This simple example demonstrates how most programs work. A program creates the various objects that it needs to complete its job. Each object has a set of methods that can be called, causing the object to perform an operation. When the program needs an object to do something, it calls the appropriate method.



NOTE:

The `import` statements that appear in lines 1 through 3 of [Code Listing 6-1](#) make the `Scanner`, `Random`, and `PrintWriter` classes available to the program. You will learn more about how the Java API is organized, and why you need these `import` statements later in this chapter.

Primitive Variables vs. Objects

Chapter 2 introduced you to the Java primitive data types: `byte`, `short`, `int`, `long`, `char`, `float`, `double`, and `boolean`. By now you have seen many programs that use both primitive data types and objects. In fact, the program in [Code Listing 6-1](#) uses two primitive variables (`maxNumbers` and `number`, both `int` variables), as well as a `Scanner` object, a `Random` object, and a `PrintWriter` object.

You've probably noticed that the steps required to create an object differ from the steps required to create a primitive variable. For example, to create an `int` variable, you simply need a declaration such as the following:

```
int wholeNumber;
```

But, to create an object, you have to write some extra code. For example, the following statement creates a `Random` object:

```
Random rand = new Rand();
```

Primitive variables, such as `ints`, `doubles`, and so forth, are simply storage locations in the computer's memory. A primitive data type is called "primitive" because a variable created with a primitive data type has no built-in capabilities other than storing a value. When you declare a primitive variable, the compiler sets aside, or allocates, a chunk of memory that is big enough for that variable. For example, look at the following variable declarations:

```
int wholeNumber;  
double realNumber;
```

Recall from Chapter 2 that an `int` uses 4 bytes of memory and a `double` uses 8 bytes of memory. These declaration statements will cause memory to be allocated as shown in [Figure 6-3](#).




Figure 6-3 Memory allocation

The memory that is allocated for a primitive variable is the actual location that will hold any value that is assigned to that variable. For example, suppose we use the following statements to assign values to the variables shown in [Figure 6-3](#):

```
wholeNumber = 99;  
realNumber = 123.45;
```

[Figure 6-4](#) shows how the assigned values are stored in each variable's memory location.




Figure 6-4 Values assigned to the variables

As you can see from these illustrations, primitive variables are very straightforward. When you are working with a primitive variable, you are using a storage location that holds a piece of data.

This is different from the way that objects work. When you are working with an object, you are typically using two things:

- The object itself, which must be created in memory
- A reference variable that refers to the object

The object that is created in memory holds data of some sort and performs operations of some sort. (Exactly what the data and operations are depends on what kind of object it is.) In order to work with the object in code, you need some way to refer to the object. That's where the reference variable comes in. The reference variable doesn't hold an actual piece of data that your program will work with. Instead, it holds the object's memory address. We say that the variable references the object. When you want to work with the object, you use the variable that references it.

Reference variables, also known as class type variables, can be used only to reference objects. [Figure 6-5](#) illustrates two objects that have been created in memory, each referenced by a variable.




Figure 6-5 Two objects referenced by variables

To understand how reference variables and objects work together, think about flying a kite. In order to fly a kite, you need a spool of string attached to it. When the kite is airborne, you use the spool of string to hold on to the kite and control it. This is similar to the relationship between an object and the variable that references the object. As shown in [Figure 6-6](#), the object is like the kite, and the variable that references the object is like the spool of string.




Figure 6-6 The kite and string metaphor

Creating an object typically requires the following two steps:

1. You declare a reference variable.
2. You create the object in memory, and assign its memory address to the reference variable.

After you have performed these steps, you can use the reference variable to work with the object. Once again, here is the familiar example of how you create an object from the `Random` class:

```
Random rand = new Random();
```

Let's look at the different parts of this statement:

- The first part of the statement, appearing on the left side of the `=` operator, reads `Random rand`. This declares a variable named `rand`, which can be used to reference an object of the `Random` type.
- The second part of the statement, appearing on the right side of the `=` operator, reads `new Random()`. The `new` operator creates an object in memory, and returns that object's memory address. So, the expression `new Random()` creates an object from the `Random` class, and returns that object's memory address.
- The `=` operator assigns the memory address that was returned from the `new` operator to the `rand` variable.

After this statement executes, the `rand` variable will reference a `Random` object, as shown in [Figure 6-7](#). The `rand` variable can then be used to perform operations with the object, such as generating random numbers.



Figure 6-7 The `rand` variable references a `Random` object



MyProgrammingLab™ www.myprogramminglab.com

1. 6.1 What does an object use its fields for?

2. 6.2 What are an object's methods?
3. 6.3 How is a class like a blueprint?
4. 6.4 You have programs that create `Scanner`, `Random`, and `PrintWriter` objects. Where are the `Scanner`, `Random`, and `PrintWriter` classes?
5. 6.5 What does the `new` operator do?
6. 6.6 What values do reference variables hold?
7. 6.7 How is the relationship between an object and a reference variable similar to a kite and a spool of string?

6.2 Writing a Simple Class, Step by Step

CONCEPT: You can write your own classes to create the objects that you need in a program. We will go through the process of writing a class in a step-by-step fashion.

The Java API provides many prewritten classes, ready for use in your programs. Sometimes, however, you will wish you had an object to perform a specific task, and no such class will exist in the Java API. This is not a problem, because you can write your own classes with the specific fields and methods that you need for any situation.

In this section we will write a class named `Rectangle`. Each object that is created from the `Rectangle` class will be able to hold data about a rectangle. Specifically, a `Rectangle` object will have the following fields:

- `length`. The `length` field will hold the rectangle's length.
- `width`. The `width` field will hold the rectangle's width.

The `Rectangle` class will also have the following methods:

- `setLength`. The `setLength` method will store a value in an object's `length` field.
- `setWidth`. The `setWidth` method will store a value in an object's `width` field.
- `getLength`. The `getLength` method will return the value in an object's `length` field.
- `getWidth`. The `getWidth` method will return the value in an object's `width` field.
- `getArea`. The `getArea` method will return the area of the rectangle, which is the result of an object's length multiplied by its width.

When designing a class it is often helpful to draw a UML diagram. UML stands for Unified Modeling Language. It provides a set of standard diagrams for graphically depicting object-oriented systems. [Figure 6-8](#) shows the general layout of a UML diagram for a class. Notice that the diagram is a box that is divided into three sections. The top section is where you write the name of the class. The middle section holds a list of the class's fields. The bottom section holds a list of the class's methods.




Figure 6-8 General layout of a UML diagram for a class

Following this layout, [Figure 6-9](#) shows a UML diagram for our `Rectangle` class. Throughout this book we frequently use UML diagrams to illustrate classes.




Figure 6-9 UML diagram for the Rectangle class

[Figure 6-9 Full Alternative Text](#)

Writing the Code for a Class

Now that we have identified the fields and methods that we want the `Rectangle` class to have, let's write the Java code. First, we use an editor to create a new file named *Rectangle.java*. In the *Rectangle.java* file we will start by writing a general class “skeleton” as follows:

```
public class Rectangle
{
}
```



Writing Classes and Creating Objects

The key word `public`, which appears in the first line, is an access specifier. An access specifier indicates how the class may be accessed. The `public` access specifier indicates that the class will be publicly available to code that is written outside the *Rectangle.java* file. Almost all of the classes

that we write in this book are `public`.

Following the access specifier is the key word `class`, followed by `Rectangle`, which is the name of the class. On the next line an opening brace appears, which is followed by a closing brace. The contents of the class, which are the fields and methods, will be written inside these braces. The general format of a class definition is as follows:

```
AccessSpecifier class Name
{
    Members
}
```

In general terms, the fields and methods that belong to a class are referred to as the class's *members*.

Writing the Code for the Class Fields

Let's continue writing our `Rectangle` class by filling in the code for some of its members. First we will write the code for the class's two fields, `length` and `width`. We will use variables of the `double` data type for the fields. The new lines of code are shown in bold, as follows:

```
public class Rectangle
{
    private double length;
    private double width;
}
```

These two lines of code that we have added declare the variables `length` and `width`. Notice that both declarations begin with the key word `private`, preceding the data type. The key word `private` is an access specifier. It indicates that these variables may not be accessed by statements outside the class.

By using the `private` access modifier, a class can hide its data from code outside the class. When a class's fields are hidden from outside code, the data is protected from accidental corruption. It is a common practice to make all of a class's fields `private` and to provide access to those fields through methods only. In other words, a class usually has `private` fields, and `public` methods that access those fields. [Table 6-1](#) summarizes the difference between the `private` and `public` access specifiers.

Table 6-1 Summary of the `private` and `public` access specifiers for class members

Access

Specifier

Description

`private` When the `private` access specifier is applied to a class member, the member cannot be accessed by code outside the class. The member can be accessed only by methods that are members of the same class.

`public` When the `public` access specifier is applied to a class member, the member can be accessed by code inside the class or outside.

Writing the `setLength` Method

Now we will begin writing the class methods. We will start with the `setLength` method. This method will allow code outside the class to store a value in the `length` field. [Code Listing 6-2](#) shows the `Rectangle` class at this stage of its development. The `setLength` method is in lines 17 through 20. (This file is in the source code folder *Chapter 06\Rectangle Class Phase 1*.)

Code Listing 6-2 (`Rectangle.java`)

```
1  /**
2      Rectangle class, phase 1
3      Under construction!
4  */
5
6  public class Rectangle
7  {
8      private double length;
9      private double width;
10
11     /**
12         The setLength method stores a value in the
13         length field.
14         @param len The value to store in length.
15     */
16
17     public void setLength(double len)
18     {
19         length = len;
20     }
21 }
```

In lines 11 through 15, we write a block comment that gives a brief description of the method. It's important always to write comments that describe a class's methods so that in the future, anyone reading the code will understand it. The definition of the method appears in lines 17 through 20. Here

is the method header:

```
public void setLength(double len)
```

The method header looks very much like any other method header that you learned to write in Chapter 5. Let's look at the parts as follows:

- `public`. The key word `public` is an access specifier. It indicates that the method may be called by statements outside the class.
- `void`. This is the method's return type. The key word `void` indicates that the method returns no data to the statement that called it.
- `setLength`. This is the name of the method.
- `(double len)`. This is the declaration of a parameter variable of the `double` data type, named `len`.

[Figure 6-10](#) labels each part of the header for the `setLength` method.




Figure 6-10 Header for the `setLength` method

[Figure 6-10 Full Alternative Text](#)

Notice that the word `static` does not appear in the method header. When a method is designed to work on an instance of a class, it is referred to as an *instance method*, and you do not write the word `static` in the header. Because this method will store a value in the `length` field of an instance of the `Rectangle` class, it is an instance method. We will discuss this in greater detail later.

After the header, the body of the method appears inside a set of braces:

```
{  
    length = len;  
}
```

The body of this method has only one statement, which assigns the value of `len` to the `length` field.

When the method executes, the `len` parameter variable will hold the value of an argument that is passed to the method. That value is assigned to the `length` field.

Before adding the other methods to the class, it might help if we demonstrate how the `setLength` method works. First, notice that the `Rectangle` class does not have a `main` method. This class is not a complete program, but is a blueprint that `Rectangle` objects may be created from. Other programs will use the `Rectangle` class to create objects. The programs that create and use these objects will have their own `main` methods. We can demonstrate the class's `setLength` method by saving the current contents of the *Rectangle.java* file and then creating the program shown in [Code Listing 6-3](#).

Code Listing 6-3 (LengthDemo.java)

```
1  /**
2   * This program demonstrates the Rectangle class's
3   * setLength method.
4   */
5
6  public class LengthDemo
7  {
8      public static void main(String[] args)
9      {
10         // Create a Rectangle object and assign its
11         // address to the box variable.
12         Rectangle box = new Rectangle();
13
14         // Indicate what we are doing.
15         System.out.println("Sending the value 10.0 " +
16                             "to the setLength method.");
17
18         // Call the box object's setLength method.
19         box.setLength(10.0);
20
21         // Indicate we are done.
22         System.out.println("Done.");
23     }
24 }
```

Program Output

```
Sending the value 10.0 to the setLength method.
Done.
```

The program in [Code Listing 6-3](#) must be saved as *LengthDemo.java* in the same folder or directory as the file *Rectangle.java*. The following command can then be used with the JDK to compile the program:

```
javac LengthDemo.java
```

When the compiler reads the source code for `LengthDemo.java` and sees that a class named `Rectangle` is being used, it looks in the current folder or directory for the file `Rectangle.class`. That file does not exist, however, because we have not yet compiled `Rectangle.java`. So, the compiler searches for the file `Rectangle.java` and compiles it. This creates the file `Rectangle.class`, which makes the `Rectangle` class available. The compiler then finishes compiling `LengthDemo.java`. The resulting `LengthDemo.class` file may be executed with the following command:

```
java LengthDemo
```

The output of the program is shown at the bottom of [Code Listing 6-3](#).

Let's look at each statement in this program's `main` method. First, the program uses the following statement, in line 12, to create a `Rectangle` object and associate it with a variable:

```
Rectangle box = new Rectangle();
```

Let's dissect the statement into two parts. The first part of the statement,

```
Rectangle box
```

declares a variable named `box`. The data type of the variable is `Rectangle`. (Because the word `Rectangle` is not the name of a primitive data type, Java assumes it to be the name of a class.) Recall that a variable of a class type is a reference variable, and it holds the memory address of an object. When a reference variable holds an object's memory address, it is said that the variable references the object. So, the variable `box` will be used to reference a `Rectangle` object. The second part of the statement is as follows:

```
= new Rectangle();
```

This part of the statement uses the key word `new`, which creates an object in memory. After the word `new`, the name of a class followed by a set of parentheses appears. This specifies the class that the object should be created from. In this case, an object of the `Rectangle` class is created. The memory address of the object is then assigned (by the `=` operator) to the variable `box`. After the statement executes, the variable `box` will reference the object that was created in memory. This is illustrated in [Figure 6-11](#).



Figure 6-11 The `box` variable references a `Rectangle` class object

Notice that [Figure 6-11](#) shows the `Rectangle` object's `length` and `width` fields set to 0. All of a class's numeric fields are initialized to 0 by default.



The parentheses in this statement are required. It would be an error to write the statement as follows:

```
Rectangle box = new Rectangle; // ERROR!!
```

The statement in lines 15 and 16 uses the `System.out.println` method to display a message on the screen. The next statement, in line 19, calls the `box` object's `setLength` method as follows:

```
box.setLength(10.0);
```

This statement passes the argument 10.0 to the `setLength` method. When the method executes, the value 10.0 is copied into the `len` parameter variable. The method assigns the value of `len` to the `length` field and then terminates. [Figure 6-12](#) shows the state of the `box` object after the method executes.




Figure 6-12 The state of the `box` object after the `setLength` method executes

Writing the `setWidth` Method

Now that we've seen how the `setLength` method works, let's add the `setWidth` method to the `Rectangle` class. The `setWidth` method is similar to `setLength`. It accepts an argument, which is assigned to the `width` field. [Code Listing 6-4](#) shows the updated `Rectangle` class. The `setWidth` method is in lines 28 through 31. (This file is stored in the source code folder *Chapter 06\Rectangle Class Phase 2*.)

Code Listing 6-4 (`Rectangle.java`)

```
1  /**
2      Rectangle class, phase 2
```

```

3     Under construction!
4     */
5
6     public class Rectangle
7     {
8         private double length;
9         private double width;
10
11        /**
12         The setLength method stores a value in the
13         length field.
14         @param len The value to store in length.
15        */
16
17        public void setLength(double len)
18        {
19            length = len;
20        }
21
22        /**
23         The setWidth method stores a value in the
24         width field.
25         @param w The value to store in width.
26        */
27
28        public void setWidth(double w)
29        {
30            width = w;
31        }
32    }

```

The `setWidth` method has a parameter variable named `w`, which is assigned to the `width` field. For example, assume that `box` references a `Rectangle` object and the following statement is executed:

```
box.setWidth(20.0);
```

After this statement executes, the `box` object's `width` field will be set to 20.0.

Writing the `getLength` and `getWidth` Methods

Because the `length` and `width` fields are private, we wrote the `setLength` and `setWidth` methods to allow code outside the `Rectangle` class to store values in the fields. We must also write methods that allow code outside the class to get the values that are stored in these fields. That's what the `getLength` and `getWidth` methods will do. The `getLength` method will return the value stored in the `length` field, and the `getWidth` method will return the value stored in the `width` field.

Here is the code for the `getLength` method:

```

public double getLength()
{
    return length;
}

```

Assume that `size` is a `double` variable and that `box` references a `Rectangle` object, and the following statement is executed:

```
size = box.getLength();
```

This statement assigns the value that is returned from the `getLength` method to the `size` variable. After this statement executes, the `size` variable will contain the same value as the `box` object's `length` field.

The `getWidth` method is similar to `getLength`. The code for the method follows:

```
public double getWidth()  
{  
    return width;  
}
```

This method returns the value that is stored in the `width` field. For example, assume that `size` is a `double` variable and that `box` references a `Rectangle` object, and the following statement is executed:

```
size = box.getWidth();
```

This statement assigns the value that is returned from the `getWidth` method to the `size` variable. After this statement executes, the `size` variable will contain the same value as the `box` object's `width` field.

[Code Listing 6-5](#) shows the `Rectangle` class with all of the members we have discussed so far. The code for the `getLength` and `getWidth` methods is shown in lines 33 through 53. (This file is stored in the source code folder *Chapter 06\Rectangle Class Phase 3*.)

Code Listing 6-5 (Rectangle.java)

```
1  /**  
2      Rectangle class, phase 3  
3      Under construction!  
4  */  
5  
6  public class Rectangle  
7  {  
8      private double length;  
9      private double width;  
10  
11     /**  
12         The setLength method stores a value in the  
13         length field.  
14         @param len The value to store in length.  
15     */  
16  
17     public void setLength(double len)  
18     {
```

```

19         length = len;
20     }
21
22     /**
23      * The setWidth method stores a value in the
24      * width field.
25      * @param w The value to store in width.
26      */
27
28     public void setWidth(double w)
29     {
30         width = w;
31     }
32
33     /**
34      * The getLength method returns a Rectangle
35      * object's length.
36      * @return The value in the length field.
37      */
38
39     public double getLength()
40     {
41         return length;
42     }
43
44     /**
45      * The getWidth method returns a Rectangle
46      * object's width.
47      * @return The value in the width field.
48      */
49
50     public double getWidth()
51     {
52         return width;
53     }
54 }

```

Before continuing we should demonstrate how these methods work. Look at the program in [Code Listing 6-6](#). (This file is also stored in the source code folder *Chapter 06\Rectangle Class Phase 3*.)

Code Listing 6-6 (LengthWidthDemo.java)

```

1  /**
2   * This program demonstrates the Rectangle class's
3   * setLength, setWidth, getLength, and getWidth methods.
4   */
5
6  public class LengthWidthDemo
7  {
8      public static void main(String[] args)
9      {
10         // Create a Rectangle object.
11         Rectangle box = new Rectangle();
12

```

```

13      // Call the object's setLength method, passing 10.0
14      // as an argument.
15      box.setLength(10.0);
16
17      // Call the object's setWidth method, passing 20.0
18      // as an argument.
19      box.setWidth(20.0);
20
21      // Display the object's length and width.
22      System.out.println("The box's length is " +
23                          box.getLength());
24      System.out.println("The box's width is " +
25                          box.getWidth());
26  }
27  }

```

Program Output

```

The box's length is 10.0
The box's width is 20.0

```

Let's take a closer look at the program. In line 11, this program creates a `Rectangle` object, which is referenced by the `box` variable. Then the following statements execute in lines 15 and 19:

```

box.setLength(10.0);
box.setWidth(20.0);

```

After these statements execute, the `box` object's `length` field is set to 10.0 and its `width` field is set to 20.0. The state of the object is shown in [Figure 6-13](#).



Figure 6-13 State of the `box` object

Next, the following statement in lines 22 and 23 executes as follows:

```

System.out.println("The box's length is " +
                    box.getLength());

```

This statement calls the `box.getLength()` method, which returns the value 10.0. The following message is displayed on the screen:

```

The box's length is 10.0

```

Then the following statement executes in lines 24 and 25:

```
System.out.println("The box's width is " +  
                    box.getWidth());
```

This statement calls the `box.getWidth()` method, which returns the value 20.0. The following message is displayed on the screen:

```
The box's width is 20.0
```

Writing the `getArea` Method

The last method we will write for the `Rectangle` class is `getArea`. This method returns the area of a rectangle, which is its length multiplied by its width. Here is the code for the `getArea` method:

```
public double getArea()  
{  
    return length * width;  
}
```

This method returns the result of the mathematical expression `length * width`. For example, assume that `area` is a `double` variable and that `box` references a `Rectangle` object, and the following code is executed:

```
box.setLength(10.0);  
box.setWidth(20.0);  
area = box.getArea();
```

The last statement assigns the value that is returned from the `getArea` method to the `area` variable. After this statement executes, the `area` variable will contain the value 200.0.

[Code Listing 6-7](#) shows the `Rectangle` class with all of the members we have discussed so far. The `getArea` method appears in lines 61 through 64. (This file is stored in the source code folder *Chapter 06\Rectangle Class Phase 4*.)

Code Listing 6-7 (`Rectangle.java`)

```
1  /**  
2      Rectangle class, phase 4  
3      Under construction!  
4  */  
5  
6  public class Rectangle  
7  {  
8      private double length;  
9      private double width;  
10  
11     /**
```



```

12         The setLength method stores a value in the
13         length field.
14         @param len The value to store in length.
15     */
16
17     public void setLength(double len)
18     {
19         length = len;
20     }
21
22     /**
23         The setWidth method stores a value in the
24         width field.
25         @param w The value to store in width.
26     */
27
28     public void setWidth(double w)
29     {
30         width = w;
31     }
32
33     /**
34         The getLength method returns a Rectangle
35         object's length.
36         @return The value in the length field.
37     */
38
39     public double getLength()
40     {
41         return length;
42     }
43
44     /**
45         The getWidth method returns a Rectangle
46         object's width.
47         @return The value in the width field.
48     */
49
50     public double getWidth()
51     {
52         return width;
53     }
54
55     /**
56         The getArea method returns a Rectangle
57         object's area.
58         @return The product of length times width.
59     */
60
61     public double getArea()
62     {
63         return length * width;
64     }
65 }

```

The program in [Code Listing 6-8](#) demonstrates all the methods of the `Rectangle` class, including `getArea`. (This file is also stored in the source code folder *Chapter 06\Rectangle Class Phase 4.*)

Code Listing 6-8 (RectangleDemo.java)

```
1  /**
2   * This program demonstrates the Rectangle class's
3   * setLength, setWidth, getLength, getWidth, and
4   * getArea methods.
5   */
6
7  public class RectangleDemo
8  {
9      public static void main(String[] args)
10     {
11         // Create a Rectangle object.
12         Rectangle box = new Rectangle();
13
14         // Set length to 10.0 and width to 20.0.
15         box.setLength(10.0);
16         box.setWidth(20.0);
17
18         // Display the length.
19         System.out.println("The box's length is " +
20                             box.getLength());
21
22         // Display the width.
23         System.out.println("The box's width is " +
24                             box.getWidth());
25
26         // Display the area.
27         System.out.println("The box's area is " +
28                             box.getArea());
29     }
30 }
```

Program Output

```
The box's length is 10.0
The box's width is 20.0
The box's area is 200.0
```

Accessor and Mutator Methods

As mentioned earlier, it is common practice to make all of a class's fields private and to provide public methods for accessing and changing those fields. This ensures that the object owning those fields is in control of all changes being made to them. A method that gets a value from a class's field but does not change it is known as an *accessor method*. A method that stores a value in a field or changes the value of a field in some other way is known as a *mutator method*. In the `Rectangle` class, the methods `getLength` and `getWidth` are accessors, and the methods `setLength` and `setWidth` are mutators.

 NOTE:

Mutator methods are sometimes called “setters” and accessor methods are sometimes called “getters.”

The Importance of Data Hiding

Data hiding is an important concept in object-oriented programming. An object hides its internal data from code that is outside the class that the object is an instance of. Only the class’s methods may directly access and make changes to the object’s internal data. You hide an object’s internal data by making the class’s fields `private`, and making the methods that access those fields `public`.

As a beginning student, you might be wondering why you would want to hide the data that is inside the classes you create. As you learn to program, you will be the user of your own classes, so it might seem that you are putting forth a great effort to hide data from yourself. If you write software in industry, however, the classes that you create will be used as components in large software systems, and programmers other than yourself will be using your classes. By hiding a class’s data, and allowing it to be accessed only through the class’s methods, you can better ensure that the class will operate as you intended it to.

Avoiding Stale Data

In the `Rectangle` class, the `getLength` and `getWidth` methods return the values stored in fields, but the `getArea` method returns the result of a calculation. You might be wondering why the area of the rectangle is not stored in a field, like the length and the width. The area is not stored in a field because it could potentially become stale. When the value of an item is dependent on other data and that item is not updated when the other data is changed, it is said that the item has become *stale*. If the area of the rectangle were stored in a field, the value of the field would become incorrect as soon as either the `length` or `width` field changed.

When designing a class, you should take care not to store in a field calculated data that can potentially become stale. Instead, provide a method that returns the result of the calculation.

Showing Access Specification in UML Diagrams

In [Figure 6-9](#) we presented a UML diagram for the `Rectangle` class. The diagram listed all of the members of the class but did not indicate which members were private and which were public. In a

UML diagram, you have the option to place a - character before a member name to indicate that it is private, or a + character to indicate that it is public. [Figure 6-14](#) shows the UML diagram modified to include this notation.




Figure 6-14 UML diagram for the Rectangle class

[Figure 6-14 Full Alternative Text](#)

Data Type and Parameter Notation in UML Diagrams

The Unified Modeling Language also provides notation that you may use to indicate the data types of fields, methods, and parameter variables. To indicate the data type of a field, place a colon followed by the name of the data type after the name of the field. For example, the `length` field in the `Rectangle` class is a `double`. It could be listed in the UML diagram as follows:

```
- length : double
```

The return type of a method can be listed in the same manner: After the method's name, place a colon followed by the return type. The `Rectangle` class's `getLength` method returns a `double`, so it could be listed in the UML diagram as follows:

```
+ getLength() : double
```

Parameter variables and their data types may be listed inside a method's parentheses. For example, the `Rectangle` class's `setLength` method has a `double` parameter named `len`, so it could be listed in the UML diagram as follows:

```
+ setLength(len : double) : void
```

[Figure 6-15](#) shows a UML diagram for the `Rectangle` class with parameter and data type notation.




Figure 6-15 UML diagram for the `Rectangle` class with parameter and data type notation

[Figure 6-15 Full Alternative Text](#)

Layout of Class Members

Notice that in the `Rectangle` class, the field variables are declared first and then the methods are defined. You are not required to write field declarations before the method definitions. In fact, some programmers prefer to write the definitions for the public methods first and write the declarations for the private fields last. Regardless of the style you use, you should be consistent. In this book, we always write the field declarations first, followed by the method definitions. [Figure 6-16](#) shows this layout.

```
public class ClassName
{
    Field declarations

    Method definitions
}
```

Figure 6-16 Typical layout of class members

[Figure 6-16 Full Alternative Text](#)



MyProgrammingLab™ www.myprogramminglab.com

1. 6.8 You hear someone make the following comment: “A blueprint is a design for a house. A carpenter can use the blueprint to build the house. If the carpenter wishes, he or she can build several identical houses from the same blueprint.” Think of this as a metaphor for classes and objects. Does the blueprint represent a class, or does it represent an object?
2. 6.9 In this chapter we used the metaphor of a kite attached to a spool of string to describe the relationship between an object and a reference variable. In this metaphor, does the kite represent an object, or a reference variable?
3. 6.10 When a variable is said to reference an object, what is actually stored in the variable?
4. 6.11 A string literal, such as "Joe", causes what type of object to be created?
5. 6.12 Look at the UML diagram in [Figure 6-17](#) and answer the following questions:




Figure 6-17 UML diagram

[Figure 6-17 Full Alternative Text](#)

1. What is the name of the class?

2. What are the fields?
3. What are the methods?
4. What are the private members?
5. What are the public members?
6. 6.13 Assume that `limo` is a variable that references an instance of the class shown in [Figure 6-17](#). Write a statement that calls `setMake` and passes the argument `"Cadillac"`.
7. 6.14 What does the key word `new` do?
8. 6.15 What is an accessor? What is a mutator?
9. 6.16 What is a stale data item?

6.3 Instance Fields and Methods

CONCEPT: Each instance of a class has its own set of fields, which are known as instance fields. You can create several instances of a class and store different values in each instance's fields. The methods that operate on an instance of a class are known as instance methods.

The program in [Code Listing 6-8](#) creates one instance of the `Rectangle` class. It is possible to create many instances of the same class, each with its own data. For example, the `RoomAreas.java` program in [Code Listing 6-9](#) creates three instances of the `Rectangle` class, referenced by the variables `kitchen`, `bedroom`, and `den`. [Figure 6-18](#) shows example interaction with the program. (The file in [Code Listing 6-9](#) is stored in the source code folder *Chapter 06\Rectangle Class Phase 4.*)

Code Listing 6-9 (`RoomAreas.java`)

```
1  import javax.swing.JOptionPane;
2
3  /**
4   * This program creates three instances of the
5   * Rectangle class.
6   */
7
8  public class RoomAreas
9  {
10     public static void main(String[] args)
11     {
12         double number;           // To hold a number
13         double totalArea;        // The total area
14         String input;            // To hold user input
15
16         // Create three Rectangle objects.
17         Rectangle kitchen = new Rectangle();
18         Rectangle bedroom = new Rectangle();
19         Rectangle den = new Rectangle();
20
21         // Get and store the dimensions of the kitchen.
22         input = JOptionPane.showInputDialog("What is the " +
23                                           "kitchen's length?");
24         number = Double.parseDouble(input);
25         kitchen.setLength(number);
26         input = JOptionPane.showInputDialog("What is the " +
27                                           "kitchen's width?");
28         number = Double.parseDouble(input);
29         kitchen.setWidth(number);
30
31         // Get and store the dimensions of the bedroom.
32         input = JOptionPane.showInputDialog("What is the " +
33                                           "bedroom's length?");
34         number = Double.parseDouble(input);
35         bedroom.setLength(number);
```



```

36     input = JOptionPane.showInputDialog("What is the " +
37                                         "bedroom's width?");
38     number = Double.parseDouble(input);
39     bedroom.setWidth(number);
40
41     // Get and store the dimensions of the den.
42     input = JOptionPane.showInputDialog("What is the " +
43                                         "den's length?");
44     number = Double.parseDouble(input);
45     den.setLength(number);
46     input = JOptionPane.showInputDialog("What is the " +
47                                         "den's width?");
48     number = Double.parseDouble(input);
49     den.setWidth(number);
50
51     // Calculate the total area of the rooms.
52     totalArea = kitchen.getArea() + bedroom.getArea()
53                + den.getArea();
54
55     // Display the total area of the rooms.
56     JOptionPane.showMessageDialog(null, "The total area " +
57                                     "of the rooms is " + totalArea);
58
59     System.exit(0);
60 }
61 }

```




Figure 6-18 Interaction with the `RoomAreas.java` program

(Oracle Corporate Counsel)

[Figure 6-18 Full Alternative Text](#)

In lines 17, 18, and 19, the following code creates three objects, each an instance of the `Rectangle` class:

```
Rectangle kitchen = new Rectangle();  
Rectangle bedroom = new Rectangle();  
Rectangle den = new Rectangle();
```

[Figure 6-19](#) illustrates how the `kitchen`, `bedroom`, and `den` variables reference the objects.




Figure 6-19 The kitchen, bedroom, and den variables reference Rectangle objects

[Figure 6-19 Full Alternative Text](#)

In the example session with the program, the user enters 10 and 14 as the length and width of the kitchen, 15 and 12 as the length and width of the bedroom, and 20 and 30 as the length and width of the den. [Figure 6-20](#) shows the states of the objects after these values are stored in them.




Figure 6-20 States of the objects after data has been stored in them

[Figure 6-20 Full Alternative Text](#)

Notice from [Figure 6-20](#) that each instance of the `Rectangle` class has its own `length` and `width` variables. For this reason, the variables are known as *instance variables*, or *instance fields*. Every instance of a class has its own set of instance fields and can store its own values in those fields.

The methods that operate on an instance of a class are known as *instance methods*. All of the methods in the `Rectangle` class are instance methods because they perform operations on specific instances of the class. For example, look at the following statement in line 25 of the `RoomAreas.java` program:

```
kitchen.setLength(number);
```

This statement calls the `setLength` method, which stores a value in the `kitchen` object's `length` field. Now look at the following statement in line 35:

```
bedroom.setLength(number);
```

This statement also calls the `setLength` method, but this time it stores a value in the `bedroom` object's `length` field. Likewise, the following statement in line 45 calls the `setLength` method to store a value in the `den` object's `length` field:

```
den.setLength(number);
```

The `setLength` method stores a value in a specific instance of the `Rectangle` class. This is true of all of the methods that are members of the `Rectangle` class.



NOTE:

As previously mentioned, instance methods do not have the key word `static` in their headers.



Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 6.17 Assume that `r1` and `r2` are variables that reference `Rectangle` objects, and the following statements are executed:

```
r1.setLength(5.0);
```

```
r2.setLength(10.0);  
r1.setWidth(20.0);  
r2.setWidth(15.0);
```

Fill in the boxes in [Figure 6-21](#) that represent each object's `length` and `width` fields.




Figure 6-21 Fill in the boxes for each field

[Figure 6-21 Full Alternative Text](#)

6.4 Constructors

CONCEPT: A constructor is a method that is automatically called when an object is created.

A constructor is a method that is automatically called when an instance of a class is created. Constructors normally perform initialization or setup operations, such as storing initial values in instance fields. They are called “constructors” because they help construct an object.



Initializing an Object with a Constructor

A constructor method has the same name as the class. For example, [Code Listing 6-10](#) shows the first few lines of a new version of the `Rectangle` class. In this version of the class, a constructor has been added. (This file is stored in the source code folder *Chapter 06\Rectangle Class Phase 5*.)

Code Listing 6-10 (Rectangle.java)

```
1  /**
2      Rectangle class, phase 5
3  */
4
5  public class Rectangle
6  {
7      private double length;
8      private double width;
9
10     /**
11         Constructor
12         @param len The length of the rectangle.
13         @param w The width of the rectangle.
14     */
15
16     public Rectangle(double len, double w)
17     {
18         length = len;
19         width = w;
20     }
```

... *The remainder of the class has not changed, and is not shown.*

This constructor accepts two arguments, which are passed into the `len` and `w` parameter variables. The parameter variables are then assigned to the `length` and `width` fields.

Notice that the constructor’s header doesn’t specify a return type—not even `void`. This is because

constructors are not executed by explicit method calls and cannot return a value. The method header for a constructor takes the following general format:

```
AccessSpecifier ClassName(Parameters...)
```

Here is an example statement that declares the variable `box`, creates a `Rectangle` object, and passes the values 7.0 and 14.0 to the constructor.

```
Rectangle box = new Rectangle(7.0, 14.0);
```

After this statement executes, `box` will reference a `Rectangle` object whose `length` field is set to 7 and whose `width` field is set to 14. The program in [Code Listing 6-11](#) demonstrates the `Rectangle` class constructor. (This file is also stored in the source code folder *Chapter 06\Rectangle Class Phase 5.*)

Code Listing 6-11 (ConstructorDemo.java)

```
1  /**
2   * This program demonstrates the Rectangle class's
3   * constructor.
4   */
5
6  public class ConstructorDemo
7  {
8      public static void main(String[] args)
9      {
10         // Create a Rectangle object, passing 5.0 and
11         // 15.0 as arguments to the constructor.
12         Rectangle box = new Rectangle(5.0, 15.0);
13
14         // Display the length.
15         System.out.println("The box's length is " +
16                             box.getLength());
17
18         // Display the width.
19         System.out.println("The box's width is " +
20                             box.getWidth());
21
22         // Display the area.
23         System.out.println("The box's area is " +
24                             box.getArea());
25     }
26 }
```

Program Output

```
The box's length is 5.0
The box's width is 15.0
The box's area is 75.0
```

Showing Constructors in a UML Diagram

There is more than one accepted way of showing a class's constructor in a UML diagram. In this book, we simply show a constructor just as any other method, except we list no return type. [Figure 6-22](#) shows a UML diagram for the `Rectangle` class with the constructor listed.




Figure 6-22 UML diagram for the `Rectangle` class showing the constructor

[Figure 6-22 Full Alternative Text](#)

Uninitialized Local Reference Variables

The program in [Code Listing 6-11](#) initializes the `box` variable with the address of a `Rectangle` object. Reference variables can also be declared without being initialized, as in the following statement:

```
Rectangle box;
```

Note that this statement does not create a `Rectangle` object. It only declares a variable named `box` that can be used to reference a `Rectangle` object. Because the `box` variable does not yet hold an object's address, it is an *uninitialized reference variable*.

After declaring the reference variable, the following statement can be used to assign it the address of an object. This statement creates a `Rectangle` object, passes the values 7.0 and 14.0 to its constructor, and assigns the object's address to the `box` variable:

```
box = new Rectangle(7.0, 14.0);
```


Once this statement executes, the `box` variable will reference a `Rectangle` object.

You need to be careful when using uninitialized reference variables. Recall from Chapter 5 that local variables *must* be initialized or assigned a value before they can be used. This is also true for local reference variables. A local reference variable must reference an object before it can be used. Otherwise a compiler error will occur.

The Default Constructor

When an object is created, its constructor is *always* called. But what if we do not write a constructor in the object's class? If you do not write a constructor in a class, Java automatically provides one when the class is compiled. The constructor that Java provides is known as the *default constructor*. The default constructor doesn't accept arguments. It sets all of the object's numeric fields to 0 and `boolean` fields to `false`. If the object has any fields that are reference variables, the default constructor sets them to the special value `null`, which means that they do not reference anything.

The *only* time that Java provides a default constructor is when you do not write your own constructor for a class. For example, at the beginning of this chapter we developed the `Rectangle` class without writing a constructor for it. When we compiled the class, the compiler generated a default constructor that set both the `length` and `width` fields to 0.0. Assume that the following code uses that version of the class to create a `Rectangle` object:

```
// We wrote no constructor for the Rectangle class.  
Rectangle r = new Rectangle(); // Calls the default constructor
```

When we created `Rectangle` objects using that version of the class, we did not pass any arguments to the default constructor because the default constructor doesn't accept arguments.

Later we added our own constructor to the class. The constructor that we added accepts arguments for the `length` and `width` fields. When we compiled the class at that point, Java did not provide a default constructor. The constructor that we added became the only constructor that the class has. When we create `Rectangle` objects with that version of the class, we *must* pass the `length` and `width` arguments to the constructor. Using that version of the class, the following statement would cause an error because we have not provided arguments for the constructor:

```
// Now we wrote our own constructor for the Rectangle class.  
Rectangle box = new Rectangle(); // Error! Must now pass arguments.
```

Because we have added our own constructor, which requires two arguments, the class no longer has a default constructor.

Writing Your Own No-Arg Constructor

A constructor that does not accept arguments is known as a *no-arg constructor*. The default

constructor doesn't accept arguments, so it is considered a no-arg constructor. In addition, you can write your own no-arg constructor. For example, suppose we wrote the following constructor for the `Rectangle` class:

```
public Rectangle()  
{  
    length = 1.0;  
    width = 1.0;  
}
```

If we were using this constructor in our `Rectangle` class, we would not pass any arguments when creating a `Rectangle` object. The following code shows an example. After this code executes, the `Rectangle` object's `length` and `width` fields would both be set to 1.0.

```
// Now we have written our own no-arg constructor.  
Rectangle r = new Rectangle(); // Calls the no-arg constructor
```

The String Class Constructor

Earlier in this chapter (in Section 6.1) we discussed the difference between creating a primitive variable and creating an object. You create primitive variables with simple declaration statements, and you create objects with the `new` operator. There is one class, however, that can be instantiated without the `new` operator: the `String` class.

Because string operations are so common, Java allows you to create `String` objects in the same way that you create primitive variables. Here is an example:

```
String name = "Joe Mahoney";
```

This statement creates a `String` object in memory, initialized with the string literal `"Joe Mahoney"`. The object is referenced by the `name` variable. If you wish, you can use the `new` operator to create a `String` object, and initialize the object by passing a string literal to the constructor, as shown here:

```
String name = new String("Joe Mahoney");
```

NOTE:

`String` objects are a special case in Java. Because they are so commonly used, Java provides numerous shortcut operations with `String` objects that are not possible with objects of other types. In addition to creating a `String` object without using the `new` operator, you can use the `=` operator to assign values to `String` objects, the `+` operator to concatenate strings, and so forth. Chapter 9 discusses several of the `String` class methods.

In the Spotlight:



Creating the `CellPhone` Class

Wireless Solutions, Inc., is a business that sells cell phones and wireless service. You are a programmer in the company's information technology (IT) department, and your team is designing a program to manage all of the cell phones that are in inventory. You have been asked to design a class that represents a cell phone. The data that should be kept as fields in the class are as follows:

- The name of the phone's manufacturer will be assigned to the `manufact` field.
- The phone's model number will be assigned to the `model` field.
- The phone's retail price will be assigned to the `retailPrice` field.

The class will also have the following methods:

- A constructor that accepts arguments for the manufacturer, model number, and retail price.
- A `setManufact` method that accepts an argument for the manufacturer. This method will allow us to change the value of the `manufact` field after the object has been created, if necessary.
- A `setModel` method that accepts an argument for the model. This method will allow us to change the value of the `model` field after the object has been created, if necessary.
- A `setRetailPrice` method that accepts an argument for the retail price. This method will allow us to change the value of the `retailPrice` field after the object has been created, if necessary.
- A `getManufact` method that returns the phone's manufacturer.
- A `getModel` method that returns the phone's model number.
- A `getRetailPrice` method that returns the phone's retail price.

[Figure 6-23](#) shows a UML diagram for the class. [Code Listing 6-12](#) shows the class definition.




Figure 6-23 UML diagram for the CellPhone class

[Figure 6-23 Full Alternative Text](#)

Code Listing 6-12 (CellPhone.java)

```
1  /**
2   The CellPhone class holds data about a cell phone.
3  */
4
5  public class CellPhone
6  {
7      // Fields
8      private String manufact;      // Manufacturer
9      private String model;         // Model
10     private double retailPrice;    // Retail price
11
12     /**
13      Constructor
14      @param man The phone's manufacturer.
15      @param mod The phone's model number.
16      @param price The phone's retail price.
17     */
18
19     public CellPhone(String man, String mod, double price)
20     {
21         manufact = man;
22         model = mod;
23         retailPrice = price;
24     }
25
```

```

26  /**
27      The setManufact method sets the phone's
28      manufacturer name.
29      @param man The phone's manufacturer.
30  */
31
32  public void setManufact(String man)
33  {
34      manufact = man;
35  }
36
37  /**
38      The setModel method sets the phone's
39      model number.
40      @param mod The phone's model number.
41  */
42
43  public void setMod(String mod)
44  {
45      model = mod;
46  }
47
48  /**
49      The setRetailPrice method sets the phone's
50      retail price.
51      @param price The phone's retail price.
52  */
53
54  public void setRetailPrice(double price)
55  {
56      retailPrice = price;
57  }
58
59  /**
60      getManufact method
61      @return The name of the phone's manufacturer.
62  */
63
64  public String getManufact()
65  {
66      return manufact;
67  }
68
69  /**
70      getModel method
71      @return The phone's model number.
72  */
73
74  public String getModel()
75  {
76      return model;
77  }
78
79  /**
80      getretailPrice method
81      @return The phone's retail price.
82  */

```

```

83
84     public double getRetailPrice()
85     {
86         return retailPrice;
87     }
88 }

```

The `CellPhone` class will be used by several programs that your team is developing. To perform a simple test of the class, you write the program shown in [Code Listing 6-13](#). This is a simple program that prompts the user for the phone's manufacturer, model number, and retail price. An instance of the `CellPhone` class is created and the data is assigned to its attributes.

Code Listing 6-13 (`CellPhoneTest.java`)

```

1  import java.util.Scanner;
2
3  /**
4   * This program runs a simple test
5   * of the CellPhone class.
6   */
7
8  public class CellPhoneTest
9  {
10     public static void main(String[] args)
11     {
12         String testMan;    // To hold a manufacturer
13         String testMod;    // To hold a model number
14         double testPrice;  // To hold a price
15
16         // Create a Scanner object for keyboard input.
17         Scanner keyboard = new Scanner(System.in);
18
19         // Get the manufacturer name.
20         System.out.print("Enter the manufacturer: ");
21         testMan = keyboard.nextLine();
22
23         // Get the model number.
24         System.out.print("Enter the model number: ");
25         testMod = keyboard.nextLine();
26
27         // Get the retail price.
28         System.out.print("Enter the retail price: ");
29         testPrice = keyboard.nextDouble();
30
31         // Create an instance of the CellPhone class,
32         // passing the data that was entered as arguments
33         // to the constructor.
34         CellPhone phone = new CellPhone(testMan, testMod, testPrice);
35
36         // Get the data from the phone and display it.
37         System.out.println();
38         System.out.println("Here is the data that you provided:");
39         System.out.println("Manufacturer: " + phone.getManufact());
40         System.out.println("Model number: " + phone.getModel());

```

```
41         System.out.println("Retail price: " + phone.getRetailPrice());
42     }
43 }
```

Program Output with Example Input Shown in Bold

Enter the manufacturer: **Acme Electronics** [Enter]
Enter the model number: **M1000** [Enter]
Enter the retail price: **199.99** [Enter]

Here is the data that you provided:
Manufacturer: Acme Electronics
Model number: M1000
Retail price: \$199.99



In the Spotlight:

Simulating Dice with Objects

Dice traditionally have six sides, representing the values 1 through 6. Some games, however, use specialized dice that have a different number of sides. For example, the fantasy role-playing game *Dungeons and Dragons*® uses dice with four, six, eight, ten, twelve, and twenty sides.

Suppose you are writing a program that needs to roll simulated dice with various numbers of sides. A simple approach would be to write a `Die` class with a constructor that accepts the number of sides as an argument. The class would also have appropriate methods for rolling the die, and getting the die's value. [Figure 6-24](#) shows the UML diagram for such a class, and [Code Listing 6-14](#) shows the code.

Die
- sides : int - value : int
+ Die(numSides : int) + roll() : void + getSides() : int + getValue() : int

Figure 6-24 UML diagram for the Die class

[Figure 6-24 Full Alternative Text](#)

Code Listing 6-14 (Die.java)

```
1  import java.util.Random;
2
3  /**
4   The Die class simulates a six-sided die.
5  */
6
7  public class Die
8  {
9   private int sides;    // Number of sides
10  private int value;    // The die's value
11
12  /**
13   The constructor performs an initial
14   roll of the die.
15   @param numSides The number of sides for this die.
16  */
17
18  public Die(int numSides)
19  {
20   sides = numSides;
21   roll();
22  }
23
24  /**
25   The roll method simulates the rolling of
26   the die.
27  */
28
29  public void roll()
30  {
31   // Create a Random object.
32   Random rand = new Random();
33
34   // Get a random value for the die.
35   value = rand.nextInt(sides) + 1;
36  }
37
38  /**
39   getSides method
40   @return The number of sides for this die.
41  */
42
43  public int getSides()
44  {
```



```

45         return sides;
46     }
47
48     /**
49         getValue method
50         @return The value of the die.
51     */
52
53     public int getValue()
54     {
55         return value;
56     }
57 }

```

Let's take a closer look at the code for the class:

Lines 9 and 10: These statements declare two `int` fields. The `sides` field will hold the number of sides that the die has, and the `value` field will hold the value of the die once it has been rolled.

Lines 18–22: This is the constructor. Notice that the constructor has a parameter for the number of sides. The parameter is assigned to the `sides` field in line 20. Line 21 calls the `roll` method, which simulates the rolling of the die.

Lines 29–36: This is the `roll` method, which simulates the rolling of the die. In line 32 a `Random` object is created, and it is referenced by the `rand` variable. Line 35 uses the `Random` object to get a random number that is in the appropriate range for this particular die. For example, if the `sides` field is set to 6, the expression `rand.nextInt(sides) + 1` will return a random integer in the range of 1 through 6. The random number is assigned to the `value` field.

Lines 43–46: This is the `getSides` method, an accessor that returns the `sides` field.

Lines 53–56: This is the `getValue` method, an accessor that returns the `value` field.

The program in [Code Listing 6-15](#) demonstrates the class. It creates two instances of the `Die` class: one with six sides, and the other with twelve sides. It then simulates five rolls of the dice.

Code Listing 6-15 (DiceDemo.java)

```
1  /**
2      This program simulates the rolling of dice.
3  */
4
5  public class DiceDemo
6  {
7      public static void main(String[] args)
8      {
9          final int DIE1_SIDES = 6;    // Number of sides for die #1
10         final int DIE2_SIDES = 12;   // Number of sides for die #2
11         final int MAX_ROLLS = 5;     // Number of times to roll
12
13         // Create two instances of the Die class.
14         Die die1 = new Die(DIE1_SIDES);
15         Die die2 = new Die(DIE2_SIDES);
16
17         // Display the initial state of the dice.
18         System.out.println("This simulates the rolling of a " +
19                             DIE1_SIDES + " sided die and a " +
20                             DIE2_SIDES + " sided die.");
21
22         System.out.println("Initial value of the dice:");
23         System.out.println(die1.getValue() + " " + die2.getValue());
24
25         // Roll the dice five times.
26         System.out.println("Rolling the dice " + MAX_ROLLS + " times.");
27
28         for (int i = 0; i < MAX_ROLLS; i++)
29         {
30             // Roll the dice.
31             die1.roll();
32             die2.roll();
33
34             // Display the values of the dice.
35             System.out.println(die1.getValue() + " " + die2.getValue());
36         }
37     }
38 }
```

Program Output

```
This simulates the rolling of a 6 sided die and a 12 sided die.
Initial value of the dice:
2 7
Rolling the dice 5 times.
3 5
5 2
2 1
4 1
5 9
```

Let's take a closer look at the program:

Lines These statements declare three constants. `DIE1_SIDES` is the number of sides for the first die
9– (6), `DIE2_SIDES` is the number of sides for the second die (12), and `MAX_ROLLS` is the number
11: of times to roll the die (5).

Lines These statements create two instances of the `Die` class. Notice that `DIE1_SIDES`, which is 6, is
14– passed to the constructor in line 14, and `DIE2_SIDES`, which is 12, is passed to the constructor
15: in line 15. As a result, `die1` will reference a `Die` object with six sides, and `die2` will
reference a `Die` object with twelve sides.

Lines This statement displays the initial value of both `Die` objects. (Recall that the `Die` class
23: constructor performs an initial roll of the die.)

Lines This `for` loop iterates five times. Each time the loop iterates, line 31 calls the `die1` object's
28– `roll` method, and line 32 calls the `die2` object's `roll` method. Line 35 displays the value of
36: both dice.

Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 6.18 How is a constructor named?
2. 6.19 What is a constructor's return type?
3. 6.20 Assume that the following is a constructor, which appears in a class:

```
ClassAct(int number)
{
    item = number;
}
```

1. What is the name of the class that this constructor appears in?
2. Write a statement that creates an object from the class and passes the value 25 as an argument to the constructor.

6.5 Passing Objects as Arguments

CONCEPT: When an object is passed as an argument to a method, the object's address is passed into the method's parameter variable. As a result, the parameter references the object.

When you are developing applications that work with objects, you will often need to write methods that accept objects as arguments. For example, suppose that a program is using the `Die` class that was previously shown in [Code Listing 6-14](#). The following code shows a method named `showDieSides` that accepts a `Die` object as an argument:

```
void showDieSides(Die d)
{
    System.out.println("This die has " + d.getSides() +
                       " sides.");
}
```

The following code sample shows how we might create a `Die` object, and then pass it as an argument to the `showDieSides` method:

```
Die myDie = new Die(6);
showDieSides(myDie)
```

When you pass an object as an argument, the thing that is passed into the parameter variable is the object's memory address. As a result, the parameter variable references the object, and the method has access to the object.

The program shown in [Code Listing 6-16](#) gives a complete demonstration. It creates two `Die` objects: one with six sides, and the other with twenty sides. It passes each object to a method named `rollDie` that displays the die's sides, rolls the die, and displays the die's value.

Code Listing 6-16 (`DieArgument.java`)

```
1  /**
2   * This program rolls a 6-sided die and
3   * a 20-sided die.
4   */
5
6  public class DieArgument
7  {
8      public static void main(String[] args)
9      {
10         final int SIX_SIDES = 6;
11         final int TWENTY_SIDES = 20;
12
13         // Create a 6-sided die.
14         Die sixDie = new Die(SIX_SIDES);
15
```

```

16      // Create a 20-sided die.
17      Die twentyDie = new Die(TWENTY_SIDES);
18
19      // Roll the dice.
20      rollDie(sixDie);
21      rollDie(twentyDie);
22  }
23
24  /**
25   * This method simulates a die roll, displaying
26   * the die's number of sides and value.
27   * @param d The Die object to roll.
28   */
29
30  public static void rollDie(Die d)
31  {
32      // Display the number of sides.
33      System.out.println("Rolling a " + d.getSides() +
34                          " sided die.");
35
36      // Roll the die.
37      d.roll();
38
39      // Display the die's value.
40      System.out.println("The die's value: " + d.getValue());
41  }
42  }

```

Program Output

```

Rolling a 6 sided die.
The die's value: 3
Rolling a 20 sided die.
The die's value: 19

```

In the Spotlight:



Simulating the Game of Cho-Han

Cho-Han is a traditional Japanese gambling game in which a dealer uses a cup to roll two six-sided dice. The cup is placed upside down on a table so that the value of the dice is concealed. Players then wager on whether the sum of the dice values is even (Cho) or odd (Han). The winner or winners take all of the wagers, or the house takes them if there are no winners.

We will develop a program that simulates a simplified variation of the game. The simulated game will have a dealer and two players. The players will not wager money, but will simply guess whether the sum of the dice values is even (Cho) or odd (Han). One point will be awarded to the player, or players, correctly guessing the outcome. The game will play for five rounds, and the player with the most points is the grand winner.

In the program, we will use the `Die` class that was introduced in [Code Listing 6-14](#). We will create two instances of the class to represent two six-sided dice. In addition to the `Die` class, we will write the following classes:

- `Dealer` class: We will create an instance of this class to represent the dealer. It will have the ability to roll the dice, report the value of the dice, and report whether the total dice value is Cho or Han.
- `Player` class: We will create two instances of this class to represent the players. Instances of the `Player` class can store the player's name, make a guess between Cho and Han, and be awarded points.

First, let's look at the `Dealer` class. [Figure 6-25](#) shows a UML diagram for the class, and [Code Listing 6-17](#) shows the code.




Figure 6-25 UML diagram for the `Dealer` class

[Figure 6-25 Full Alternative Text](#)

Code Listing 6-17 (`Dealer.java`)

```
1  /**
2      Dealer class for the game of Cho-Han
3  */
4
5  public class Dealer
6  {
```

```

7     private int die1Value;    // The value of die #1
8     private int die2Value;    // The value of die #2
9
10    /**
11        Constructor
12    */
13
14    public Dealer()
15    {
16        die1Value = 0;
17        die2Value = 0;
18    }
19
20    /**
21        The rollDice method rolls the dice and saves
22        their values.
23    */
24
25    public void rollDice()
26    {
27        final int SIDES = 6; // Number of sides for the dice
28
29        // Create the two dice. (This also rolls them.)
30        Die die1 = new Die(SIDES);
31        Die die2 = new Die(SIDES);
32
33        // Record their values.
34        die1Value = die1.getValue();
35        die2Value = die2.getValue();
36    }
37
38    /**
39        The getChoOrHan method returns the result of
40        the dice roll, Cho or Han.
41        @return Either "Cho (even)" or "Han (odd)"
42    */
43
44    public String getChoOrHan()
45    {
46        String result; // To hold the result
47
48        // Get the sum of the dice.
49        int sum = die1Value + die2Value;
50
51        // Determine even or odd.
52        if (sum % 2 == 0)
53            result = "Cho (even)";
54        else
55            result = "Han (odd)";
56
57        // Return the result.
58        return result;
59    }
60
61    /**
62        The getDie1Value method returns the value of
63        die #1.

```

```

64         @return The die1Value field
65     */
66
67     public int getDie1Value()
68     {
69         return die1Value;
70     }
71
72     /**
73         The getDie2Value method returns the value of
74         die #2.
75         @return The die2Value field
76     */
77
78     public int getDie2Value()
79     {
80         return die2Value;
81     }
82 }

```

Let's take a closer look at the code for the `Dealer` class:

- Lines 7 and 8 declare the fields `die1Value` and `die2Value`. These fields will hold the value of the two dice after they have been rolled.
- The constructor, in lines 14 through 18, initializes the `die1Value` and `die2Value` fields to 0.
- The `rollDice` method, in lines 25 through 36, simulates the rolling of the dice. Lines 30 and 31 create two `Die` objects. Recall that the `Die` class constructor performs an initial roll of the die, so there is no need to call the `Die` objects' `roll` method. Lines 34 and 35 save the value of the dice in the `die1Value` and `die2Value` fields.
- The `getChoOrHan` method, in lines 44 through 59, returns a string indicating whether the sum of the dice is Cho (even) or Han (odd).
- The `getDie1Value` method, in lines 67 through 70, returns the value of the first die (stored in the `die1Value` field).
- The `getDie2Value` method, in lines 78 through 81, returns the value of the second die (stored in the `die2Value` field).

Now let's look at the `Player` class. [Figure 6-26](#) shows a UML diagram for the class, and [Code Listing 6-18](#) shows the code.




Figure 6-26 UML diagram for the Player class

[Figure 6-26 Full Alternative Text](#)

Code Listing 6-18 (Player.java)

```

1  import java.util.Random;
2
3  /**
4   * Player class for the game of Cho-Han
5   */
6
7  public class Player
8  {
9      private String name;        // The player's name
10     private String guess;       // The player's guess
11     private int points;         // The player's points
12
13     /**
14      * Constructor
15      * @param playerName The player's name.
16      */
17
18     public Player(String playerName)
19     {
20         name = playerName;
21         guess = "";
22         points = 0;
23     }
24
25     /**
26      * The makeGuess method causes the player to guess
27      * either "Cho (even)" or "Han (odd)".
28      */
29

```

```

30 public void makeGuess()
31 {
32     // Create a Random object.
33     Random rand = new Random();
34
35     // Get a random number, either 0 or 1.
36     int guessNumber = rand.nextInt(2);
37
38     // Convert the random number to a guess of
39     // either "Cho (even)" or "Han (odd)".
40     if (guessNumber == 0)
41         guess = "Cho (even)";
42     else
43         guess = "Han (odd)";
44 }
45
46 /**
47     The addPoints method adds a specified number of
48     points to the player's current balance.
49     @newPoints The points to add.
50 */
51
52 public void addPoints(int newPoints)
53 {
54     points += newPoints;
55 }
56
57 /**
58     The getName method returns the player's name.
59     @return The value of the name field.
60 */
61
62 public String getName()
63 {
64     return name;
65 }
66
67 /**
68     The getGuess method returns the player's guess.
69     @return The value of the guess field.
70 */
71
72 public String getGuess()
73 {
74     return guess;
75 }
76
77 /**
78     The getPoints method returns the player's points
79     @return The value of the points field.
80 */
81
82 public int getPoints()
83 {
84     return points;
85 }
86 }

```

Here's a summary of the code for the `Player` class:

- Lines 9 through 11 declare the fields `name`, `guess`, and `points`. These fields will hold the player's name, the player's guess, and the number of points the player has earned.
- The constructor, in lines 18 through 23, accepts an argument for the player's name, which is assigned to the `name` field. The `guess` field is assigned an empty string, and the `points` field is set to 0.
- The `makeGuess` method, in lines 30 through 44, causes the player to make a guess. The method generates a random number that is either a 0 or a 1. The `if` statement that begins at line 40 assigns the string "Cho (even)" to the `guess` field if the random number is 0, or it assigns the string "Han (odd)" to the `guess` field if the random number is 1.
- The `addPoints` method, in lines 52 through 55, adds the number of points specified by the argument to the player's `point` field.
- The `getName` method, in lines 62 through 65, returns the player's name.
- The `getGuess` method, in lines 72 through 75, returns the player's guess.
- The `getPoints` method, in lines 82 through 85, returns the player's points.

[Code Listing 6-19](#) shows the program that uses these classes to simulate the game. The `main` method simulates five rounds of the game, displaying the results of each round, and then displays the overall game results.

Code Listing 6-19 (ChoHan.java)

```
1  import java.util.Scanner;
2
3  public class ChoHan
4  {
5      public static void main(String[] args)
6      {
7          final int MAX_ROUNDS = 5;    // Number of rounds
8          String player1Name;          // First player's name
9          String player2Name;          // Second player's name
10
11          // Create a Scanner object for keyboard input.
12          Scanner keyboard = new Scanner(System.in);
13
14          // Get the players' names.
15          System.out.print("Enter the first player's name: ");
16          player1Name = keyboard.nextLine();
17          System.out.print("Enter the second player's name: ");
18          player2Name = keyboard.nextLine();
19
20          // Create the dealer.
```

```

21     Dealer dealer = new Dealer();
22
23     // Create the two players.
24     Player player1 = new Player(player1Name);
25     Player player2 = new Player(player2Name);
26
27     // Play the rounds.
28     for (int round = 0; round < MAX_ROUNDS; round++)
29     {
30         System.out.println("-----");
31         System.out.printf("Now playing round %d.\n", round + 1);
32
33         // Roll the dice.
34         dealer.rollDice();
35
36         // The players make their guesses.
37         player1.makeGuess();
38         player2.makeGuess();
39
40         // Determine the winner of this round.
41         roundResults(dealer, player1, player2);
42     }
43
44     // Display the grand winner.
45     displayGrandWinner(player1, player2);
46 }
47
48 /**
49     The roundResults method determines the results of
50     the current round.
51     @param dealer The Dealer object
52     @param player1 Player #1 object
53     @param player2 Player #2 object
54 */
55
56 public static void roundResults(Dealer dealer, Player player1,
57                                Player player2)
58 {
59     // Show the dice values.
60     System.out.printf("The dealer rolled %d and %d.\n",
61                       dealer.getDie1Value(), dealer.getDie2Value());
62     System.out.printf("Result: %s\n", dealer.getChoOrHan());
63
64     // Check each player's guess and award points.
65     checkGuess(player1, dealer);
66     checkGuess(player2, dealer);
67 }
68
69 /**
70     The checkGuess method checks a player's guess against
71     the dealer's result.
72     @param player The Player object to check.
73     @param dealer The Dealer object.
74 */
75
76 public static void checkGuess(Player player, Dealer dealer)
77 {
78     final int POINTS_TO_ADD = 1; // Points to award winner

```

```

79         String guess = player.getGuess(); // Player's guess
80         String choHanResult = dealer.getChoOrHan(); // Cho or Han
81
82         // Display the player's guess.
83         System.out.printf("The player %s guessed %s.\n",
84                             player.getName(), player.getGuess());
85
86         // Award points if the player guessed correctly.
87         if (guess.equalsIgnoreCase(choHanResult))
88         {
89             player.addPoints(POINTS_TO_ADD);
90             System.out.printf("Awarding %d point(s) to %s.\n",
91                             POINTS_TO_ADD, player.getName());
92         }
93     }
94
95     /**
96      * The displayGrandWinner method displays the game's grand winner.
97      * @param player1 Player #1
98      * @param player2 Player #2
99      */
100
101     public static void displayGrandWinner(Player player1, Player player2)
102     {
103         System.out.println("-----");
104         System.out.println("Game over. Here are the results:");
105         System.out.printf("%s: %d points.\n", player1.getName(),
106                             player1.getPoints());
107         System.out.printf("%s: %d points.\n", player2.getName(),
108                             player2.getPoints());
109
110         if (player1.getPoints() > player2.getPoints())
111             System.out.println(player1.getName() + " is the grand winner!");
112         else if (player2.getPoints() > player1.getPoints())
113             System.out.println(player2.getName() + " is the grand winner!");
114         else
115             System.out.println("Both players are tied!");
116     }
117 }

```

Program Output with Example Input Shown in Bold

```

Enter the first player's name: Chelsea [Enter]
Enter the second player's name: Chris [Enter]
-----
Now playing round 1.
The dealer rolled 3 and 6.
Result: Han (odd)
The player Chelsea guessed Han (odd).
Awarding 1 point(s) to Chelsea.
The player Chris guessed Han (odd).
Awarding 1 point(s) to Chris.

```

```

-----
Now playing round 2.
The dealer rolled 4 and 5.
Result: Han (odd)
The player Chelsea guessed Cho (even).
The player Chris guessed Cho (even).
-----
Now playing round 3.
The dealer rolled 5 and 6.
Result: Han (odd)
The player Chelsea guessed Cho (even).
The player Chris guessed Han (odd).
Awarding 1 point(s) to Chris.
-----
Now playing round 4.
The dealer rolled 1 and 6.
Result: Han (odd)
The player Chelsea guessed Cho (even).
The player Chris guessed Cho (even).
-----
Now playing round 5.
The dealer rolled 6 and 6.
Result: Cho (even)
The player Chelsea guessed Han (odd).
The player Chris guessed Cho (even).
Awarding 1 point(s) to Chris.
-----
Game over. Here are the results:
Chelsea: 1 points.
Chris: 3 points.
Chris is the grand winner!

```

Let's look at the code. Here is a summary of the `main` method:

- Lines 7 through 9 make the following declarations: `MAX_ROUNDS`—the number of rounds to play, `player1Name`—to hold the name of player #1, and `player2Name`—to hold the name of player #2.
- Lines 15 through 18 prompt the user to enter the players' names.
- Line 21 creates an instance of the `Dealer` class. The object represents the dealer, and is referenced by the `dealer` variable.
- Line 24 creates an instance of the `Player` class. The object represents player #1, and is referenced by the `player1` variable. Notice that `player1Name` is passed as an argument to the constructor.
- Line 25 creates another instance of the `Player` class. The object represents player #2, and is referenced by the `player2` variable. Notice that `player2Name` is passed as an argument to the constructor.
- The `for` loop that begins in line 28 iterates five times, causing the simulation of five rounds of the game. The loop performs the following actions:

- Line 34 causes the dealer to roll the dice.
- Line 37 causes player #1 to make a guess (Cho or Han).
- Line 38 causes player #2 to make a guess (Cho or Han).
- Line 41 passes the `dealer`, `player1`, and `player2` objects to the `roundResults` method. The method displays the results of this round.
- Line 45 passes the `player1` and `player2` objects to the `displayGrandWinner` method, which displays the grand winner of the game.

The `roundResults` method, which displays the results of a round, appears in lines 56 through 67. Here is a summary of the method:

- The method accepts references to the `dealer`, `player1`, and `player2` objects as arguments.
- The statement in lines 60 and 61 displays the value of the two dice.
- Line 62 calls the `dealer` object's `getChoOrHan` method to display the results, Cho or Han.
- Line 65 calls the `checkGuess` method, passing the `player1` and `dealer` objects as arguments. The `checkGuess` method compares a player's guess to the dealer's result (Cho or Han), and awards points to the player, if the guess is correct.
- Line 66 calls the `checkGuess` method, passing the `player2` and `dealer` objects as arguments.

The `checkGuess` method, which compares a player's guess to the dealer's result, awarding points to the player for a correct guess, appears in lines 76 through 93. Here is a summary of the method:

- The method accepts references to a `Player` object and the `Dealer` object as arguments.
- Line 78 declares the constant `POINTS_TO_ADD`, set to the value 1, which is the number of points to add to the player's balance if the player's guess is correct.
- Line 79 assigns the player's guess to the `String` object `guess`.
- Line 80 assigns the dealer's results (Cho or Han) to the `String` object `choHanResult`.
- The statement in lines 83 and 84 displays the player's name and guess.
- The `if` statement in line 87 compares the player's guess to the dealer's result. If they match, then the player guessed correctly, and line 89 awards points to the player.

The `displayGrandWinner` method, which displays the grand winner of the game, appears in lines 101 through 116. Here is a summary of the method:

- The method accepts references to the `player1` and `player2` objects.
- The statements in lines 105 through 108 display both players' names and points.
- The `if-else-if` statement that begins in line 110 determines which of the two players has the highest score, and displays that player's name as the grand winner. If both players have the same score, a tie is declared.

6.6 Overloading Methods and Constructors

CONCEPT: Two or more methods in a class may have the same name as long as their parameter lists are different. This also applies to constructors.

Method overloading is an important part of object-oriented programming. When a method is *overloaded*, it means that multiple methods in the same class have the same name, but use different types of parameters. Method overloading is important because sometimes you need several different ways to perform the same operation. For example, suppose a class has the following two methods:

```
public int add(int num1, int num2)
{
    int sum = num1 + num2;
    return sum;
}
public String add(String str1, String str2)
{
    String combined = str1 + str2;
    return combined;
}
```

Both of these methods are named `add`. They both take two arguments, which are added together. The first one accepts two `int` arguments and returns their sum. The second accepts two `String` references and returns a reference to a `String` that is a concatenation of the two arguments. When we write a call to the `add` method, the compiler must determine which one of the overloaded methods we intended to call.

The process of matching a method call with the correct method is known as *binding*. When an overloaded method is being called, Java uses the method's name and parameter list to determine which method to bind the call to. If two `int` arguments are passed to the `add` method, the version of the method with two `int` parameters is called. Likewise, when two `String` arguments are passed to `add`, the version with two `String` parameters is called.

Java uses a method's signature to distinguish it from other methods of the same name. A method's *signature* consists of the method's name and the data types of the method's parameters, in the order that they appear. For example, here are the signatures of the `add` methods that were previously shown:

```
add(int, int)
add(String, String)
```

Note that the method's return type is *not* part of the signature. For this reason, the following `add` method cannot be added to the same class with the previous ones:

```
public int add(String str1, String str2)
{
    int sum = Integer.parseInt(str1) + Integer.parseInt(str2);
    return sum;
}
```

```
}
```

Because the return type is not part of the signature, this method's signature is the same as that of the other `add` method that takes two `String` arguments. For this reason, an error message will be issued when a class containing all of these methods is compiled.

Constructors can also be overloaded, which means that a class can have more than one constructor. The rules for overloading constructors are the same for overloading other methods: Each version of the constructor must have a different parameter list. As long as each constructor has a unique signature, the compiler can tell them apart. For example, the `Rectangle` class that we discussed earlier could have the following two constructors:

```
public Rectangle()  
{  
    length = 0.0;  
    width = 0.0;  
}  
public Rectangle(double len, double w)  
{  
    length = len;  
    width = w;  
}
```

The first constructor shown here accepts no arguments, and assigns 0.0 to the `length` and `width` fields. The second constructor accepts two arguments, which are assigned to the `length` and `width` fields. The following code shows an example of how each constructor is called:

```
Rectangle box1 = new Rectangle();  
Rectangle box2 = new Rectangle(5.0, 10.0);
```

The first statement creates a `Rectangle` object, referenced by the `box1` variable, and executes the no-arg constructor. Its `length` and `width` fields will be set to 0.0. The second statement creates another `Rectangle` object, referenced by the `box2` variable, and executes the second constructor. Its `length` and `width` fields will be set to 5.0 and 10.0, respectively.

Recall that Java provides a default constructor only when you do not write any constructors for a class. If a class has a constructor that accepts arguments, but it does not have a no-arg constructor, you cannot create an instance of the class without passing arguments to the constructor. Therefore, any time you write a constructor for a class, and that constructor accepts arguments, you should also write a no-arg constructor if you want to be able to create instances of the class without passing arguments to the constructor.

The BankAccount Class

Now we will look at the `BankAccount` class. Objects that are created from this class will simulate bank accounts, allowing us to have a starting balance, make deposits, make withdrawals, and get the current balance. A UML diagram for the `BankAccount` class is shown in [Figure 6-27](#). In the figure,

the overloaded constructors and overloaded methods are pointed out. Note that the extra annotation is not part of the UML diagram. It is there to draw attention to the items that are overloaded.




Figure 6-27 UML diagram for the BankAccount class

[Figure 6-27 Full Alternative Text](#)

As you can see from the diagram, the class has three overloaded constructors. Also, the class has two overloaded methods named `deposit`, two overloaded methods named `withdraw`, and two overloaded methods named `setBalance`. The last method, `getBalance`, is not overloaded. [Code Listing 6-20](#) shows the code for the class.

Code Listing 6-20 (`BankAccount.java`)

```
1  /**
2   * The BankAccount class simulates a bank account.
3   */
4
5  public class BankAccount
6  {
7      private double balance;          // Account balance
8
9      /**
10     * This constructor sets the starting balance
11     * at 0.0.
12     */
13
14     public BankAccount()
15     {
```

```
16         balance = 0.0;
17     }
18
19     /**
20      * This constructor sets the starting balance
21      * to the value passed as an argument.
22      * @param startBalance The starting balance.
23      */
24
25     public BankAccount(double startBalance)
26     {
27         balance = startBalance;
28     }
29
30     /**
31      * This constructor sets the starting balance
32      * to the value in the String argument.
33      * @param str The starting balance, as a String.
34      */
35
36     public BankAccount(String str)
37     {
38         balance = Double.parseDouble(str);
39     }
40
41     /**
42      * The deposit method makes a deposit into
43      * the account.
44      * @param amount The amount to add to the
45      *                balance field.
46      */
47
48     public void deposit(double amount)
49     {
50         balance += amount;
51     }
52
53     /**
54      * The deposit method makes a deposit into
55      * the account.
56      * @param str The amount to add to the
57      *            balance field, as a String.
58      */
59
60     public void deposit(String str)
61     {
62         balance += Double.parseDouble(str);
63     }
64
65     /**
66      * The withdraw method withdraws an amount
67      * from the account.
68      * @param amount The amount to subtract from
69      *                the balance field.
70      */
71
72     public void withdraw(double amount)
```

```

73     {
74         balance -= amount;
75     }
76
77     /**
78      * The withdraw method withdraws an amount
79      * from the account.
80      * @param str The amount to subtract from
81      *           the balance field, as a String.
82      */
83
84     public void withdraw(String str)
85     {
86         balance -= Double.parseDouble(str);
87     }
88
89     /**
90      * The setBalance method sets the account balance.
91      * @param b The value to store in the balance field.
92      */
93
94     public void setBalance(double b)
95     {
96         balance = b;
97     }
98
99     /**
100    * The setBalance method sets the account balance.
101    * @param str The value, as a String, to store in
102    *           the balance field.
103    */
104
105    public void setBalance(String str)
106    {
107        balance = Double.parseDouble(str);
108    }
109
110    /**
111    * The getBalance method returns the
112    * account balance.
113    * @return The value in the balance field.
114    */
115
116    public double getBalance()
117    {
118        return balance;
119    }
120 }

```

The class has one field, `balance`, which is a `double`. This field holds an account's current balance. Here is a summary of the class's overloaded constructors:

- The first constructor is a no-arg constructor. It sets the `balance` field to 0.0. If we wish to execute this constructor when we create an instance of the class, we simply pass no constructor arguments. Here is an example:

```
BankAccount account = new BankAccount();
```

- The second constructor has a `double` parameter variable, `startBalance`, which is assigned to the `balance` field. If we wish to execute this constructor when we create an instance of the class, we pass a `double` value as a constructor argument. Here is an example:

```
BankAccount account = new BankAccount(1000.0);
```

- The third constructor has a `String` parameter variable, `str`. It is assumed that the `String` contains a string representation of the account's balance. The method uses the `Double.parseDouble` method to convert the string to a `double`, and then assigns it to the `balance` field. If we wish to execute this constructor when we create an instance of the class, we pass a reference to a `String` as a constructor argument. Here is an example:

```
BankAccount account = new BankAccount("1000.0");
```

This constructor is provided as a convenience. If the class is used in a program that reads the account balance from a dialog box, or from a text file, the amount does not have to be converted from a string before it is passed to the constructor.

Here is a summary of the overloaded `deposit` methods:

- The first `deposit` method has a parameter, `amount`, which is a `double`. When the method is called, an amount that is to be deposited into the account is passed into this parameter. The value of the parameter is then added to value in the `balance` field.
- The second `deposit` method has a parameter, `str`, which is a reference to a `String`. It is assumed that the `String` contains a string representation of the amount to be deposited. The method uses the `Double.parseDouble` method to convert the string to a `double`, and then adds it to the `balance` field. For example, if we call the method and pass "500.0" as the argument, it will add 500.0 to the `balance` field. As with the overloaded constructors, this method is provided as a convenience for programs that read the amount to be deposited from a dialog box or a text file.

Here is a summary of the overloaded `withdraw` methods:

- The first `withdraw` method has a parameter, `amount`, which is a `double`. When the method is called, an amount that is to be withdrawn from the account is passed into this parameter. The value of the parameter is then subtracted from the value in the `balance` field.
- The second `withdraw` method has a parameter, `str`, which is a reference to a `String`. It is assumed that the `String` contains a string representation of the amount to be withdrawn. This amount is converted to a `double`, and then subtracted from the `balance` field. As with the overloaded constructors and `deposit` methods, this method is provided as a convenience.

Here is a summary of the overloaded `setBalance` methods:

- The first `setBalance` method accepts a `double` argument, which is assigned to the `balance` field.
- The second `setBalance` method accepts a `String` reference as an argument. It is assumed that the `String` contains a string representation of the account's balance. The `String` is converted to a `double` and then assigned to the `balance` field. As with many of the other overloaded methods, this method is provided as a convenience.

The remaining method is `getBalance`. It returns the value in the `balance` field, which is the current account balance. The `AccountTest.java` program, shown in [Code Listing 6-21](#), demonstrates the `BankAccount` class. Its output is shown in [Figure 6-28](#).

Code Listing 6-21 (AccountTest.java)

```

1  import javax.swing.JOptionPane; // For the JOptionPane class
2
3  /**
4   * This program demonstrates the BankAccount class.
5   */
6
7  public class AccountTest
8  {
9      public static void main(String[] args)
10     {
11         String input;    // To hold user input
12
13         // Get the starting balance.
14         input = JOptionPane.showInputDialog(
15             "What is your account's starting balance?");
16
17         // Create a BankAccount object.
18         BankAccount account = new BankAccount(input);
19
20         // Get the amount of pay.
21         input = JOptionPane.showInputDialog(
22             "How much were you paid this month?");
23
24         // Deposit the user's pay into the account.
25         account.deposit(input);
26
27         // Display the new balance.
28         JOptionPane.showMessageDialog(null,
29             String.format("Your pay has been deposited.\n" +
30                 "Your current balance is $%,.2f",
31                 account.getBalance()));
32
33         // Withdraw some cash from the account.
34         input = JOptionPane.showInputDialog(
35             "How much would you like to withdraw?");
36         account.withdraw(input);
37
38         // Display the new balance

```

```

39         JOptionPane.showMessageDialog(null,
40             String.format("Now your balance is $%,.2f",
41                 account.getBalance()));
42
43     System.exit(0);
44 }
45 }

```




Figure 6-28 Interaction with the AccountTest program

(Oracle Corporate Counsel)

[Figure 6-28 Full Alternative Text](#)

Overloaded Methods Make Classes More Useful

You might be wondering why all those overloaded methods appear in the `BankAccount` class, especially because many of them weren't used by the demonstration program in [Code Listing 6-21](#). After all, wouldn't it be simpler if the class had only the methods we were going to use?

An object's purpose is to provide a specific service. The service provided by the `BankAccount` class is that it simulates a bank account. Any program that needs a simulated bank account can simply create a `BankAccount` object and then use its methods to put the simulation into action. Because the `BankAccount` class has numerous overloaded methods, it is much more flexible than it would be if it provided only one way to perform every operation. By providing overloaded constructors, `deposit` methods, `withdraw` methods, and `setBalance` methods, we made the `BankAccount` class useful to programs other than our simple demonstration program. This is an important consideration to keep in mind when you design classes of your own.

6.7 Scope of Instance Fields

CONCEPT: Instance fields are visible to all of the class's instance methods.

Recall from Chapter 2 that a variable's scope is the part of a program where the variable may be accessed by its name. A variable's name is visible only to statements inside the variable's scope. The location of a variable's declaration determines the variable's scope.

In this chapter you have seen variables declared as instance fields in a class. An instance field can be accessed by any instance method in the same class as the field. If an instance field is declared with the `public` access specifier, it can also be accessed by code outside the class.

Shadowing

In Chapter 2 you saw that you cannot have two local variables with the same name in the same scope. This applies to parameter variables as well. A parameter variable is, in essence, a local variable. So, you cannot give a parameter variable and a local variable in the same method the same name.

However, you can have a local variable or a parameter variable with the same name as a field. When you do, the name of the local or parameter variable *shadows* the name of the field. This means that the field name is hidden by the name of the local or parameter variable.

For example, assume that the `Rectangle` class's `setLength` method had been written in the following manner:

```
public void setLength(double len)
{
    int length;           // Local variable
    length = len;
}
```

In this code a local variable is given the same name as a field. Therefore, the local variable's name shadows the field's name. When the statement `length = len;` is executed, the value of `len` is assigned to the local variable `length`, not to the field. The unintentional shadowing of field names can cause elusive bugs, so you need to be careful not to give local variables the same names as fields.



MyProgrammingLab™ www.myprogramminglab.com

1. 6.21 Is it required that overloaded methods have different return values, different parameter

lists, or both?

2. 6.22 What is a method's signature?

3. 6.23 Look at the following class:

```
public class CheckPoint
{
    public void message(int x)
    {
        System.out.print("This is the first version ");
        System.out.println("of the method.");
    }
    public void message(String x)
    {
        System.out.print("This is the second version ");
        System.out.println("of the method.");
    }
}
```

What will the following code display?

```
CheckPoint cp = new CheckPoint();
cp.message("1");
cp.message(1);
```

4. 6.24 How many default constructors may a class have?

6.8 Packages and `import` Statements

CONCEPT: The classes in the Java API are organized into packages. An `import` statement tells the compiler which package a class is located in.

In Chapter 2 you were introduced to the Java API, which is a standard library of prewritten classes. Each class in the Java API is designed for a specific purpose, and you can use the classes in your own programs. You've already used a few classes from the API, such as the `String` class, the `Scanner` class, the `JOptionPane` class, and the `Random` class.

All of the classes in the Java API are organized into packages. A *package* is simply a group of related classes. Each package also has a name. For example, the `Scanner` class is in the `java.util` package.

Many of the classes in the Java API are not automatically available to your program. Quite often, you have to *import* an API class in order to use it. You use the `import` key word to import a class. For example, the following statement is required to import the `Scanner` class:

```
import java.util.Scanner;
```

This statement tells the compiler that the `Scanner` class is located in the `java.util` package. Without this statement, the compiler will not be able to locate the `Scanner` class, and the program will not compile.

Explicit and Wildcard `import` Statements

There are two types of `import` statements: explicit and wildcard. An *explicit* `import` statement identifies the package location of a single class. For example, the following statement explicitly identifies the location of the `Scanner` class:

```
import java.util.Scanner;
```

The `java.util` package contains several other classes as well as the `Scanner` class. For example, the `Random` class is also part of the `java.util` package. If a program needs to use the `Scanner` class and the `Random` class, it will have to import both of these classes. One way to do this is to write explicit `import` statements for each class, as follows:

```
import java.util.Scanner;
import java.util.Random;
```

Another way to import both of these classes is to use a wildcard `import` statement. A *wildcard* `import` statement tells the compiler to import all of the classes in a package. Here is an example:

```
import java.util.*;
```

The `.*` that follows the package name tells the compiler to import all the classes that are part of the `java.util` package. Using a wildcard `import` statement does not affect the performance or the size of your program. It merely tells the compiler that you want to make every class in a particular package available to your program.

The `java.lang` Package

The Java API does have one package, `java.lang`, which is automatically imported into every Java program. This package contains general classes, such as `String` and `System`, that are fundamental to the Java programming language. You do not have to write an `import` statement for any class that is part of the `java.lang` package.

Other API Packages

There are numerous packages in the Java API. [Table 6-2](#) lists a few of them.

Table 6-2 A few of the standard Java packages

Package	Description
<code>java.io</code>	Provides classes that perform various types of input and output.
<code>java.lang</code>	Provides general classes for the Java language. This package is automatically imported.
<code>java.net</code>	Provides classes for network communications.
<code>java.security</code>	Provides classes that implement security features.
<code>java.sql</code>	Provides classes for accessing databases using structured query language.
<code>java.text</code>	Provides various classes for formatting text.

`java.util` Provides various utility classes.

See Appendix G, available on the book's companion Web site at www.pearsonhighered.com/gaddis, for a more detailed look at packages.

6.9 Focus on Object-Oriented Design: Finding the Classes and Their Responsibilities

So far you have learned the basics of writing a class, creating an object from the class, and using the object to perform operations. Although this knowledge is necessary to create an object-oriented application, it is not the first step. The first step is to analyze the problem that you are trying to solve and determine the classes that you will need. In this section, we will discuss a simple technique for finding the classes in a problem and determining their responsibilities.

Finding the Classes

When developing an object-oriented application, one of your first tasks is to identify the classes that you will need to create. Typically, your goal is to identify the different types of real-world objects that are present in the problem, and then create classes for those types of objects within your application.

Over the years, software professionals have developed numerous techniques for finding the classes in a given problem. One simple and popular technique involves the following steps:

1. Get a written description of the problem domain.
2. Identify all the nouns (including pronouns and noun phrases) in the description. Each of these is a potential class.
3. Refine the list to include only the classes that are relevant to the problem.

Let's take a closer look at each of these steps.

Writing a Description of the Problem Domain

The *problem domain* is the set of real-world objects, parties, and major events related to the problem. If you adequately understand the nature of the problem you are trying to solve, you can write a description of the problem domain yourself. If you do not thoroughly understand the nature of the problem, you should have an expert write the description for you.

For example, suppose we are programming an application that the manager of Joe's Automotive Shop will use to print service quotes for customers. Here is a description that an expert, perhaps Joe himself, might have written:

- Joe’s Automotive Shop services foreign cars, and specializes in servicing cars made by Mercedes, Porsche, and BMW. When a customer brings a car to the shop, the manager gets the customer’s name, address, and telephone number. Then the manager determines the make, model, and year of the car, and gives the customer a service quote. The service quote shows the estimated parts charges, estimated labor charges, sales tax, and total estimated charges.

The problem domain description should include any of the following:

- Physical objects such vehicles, machines, or products
- Any role played by a person, such as manager, employee, customer, teacher, student, and so forth
- The results of a business event, such as a customer order, or in this case a service quote
- Recordkeeping items, such as customer histories and payroll records

Identifying All of the Nouns

The next step is to identify all of the nouns and noun phrases. (If the description contains pronouns, include them too.) Here’s another look at the previous problem domain description. This time the nouns and noun phrases appear in bold.

- Joe’s Automotive Shop services foreign cars, and specializes in servicing cars made by Mercedes, Porsche, and BMW. When a customer brings a car to the shop, the manager gets the customer’s name, address, and telephone number. Then the manager determines the make, model, and year of the car, and gives the customer a service quote. The service quote shows the estimated parts charges, estimated labor charges, sales tax, and total estimated charges.

Notice that some of the nouns are repeated. The following list shows all of the nouns without duplication.

address	foreign cars	Porsche
BMW	Joe’s Automotive Shop	sales tax
car	make	service quote
cars	manager	shop
customer	Mercedes	telephone number

estimated labor charges modeltotal estimated charges

estimated parts charges nameyear

Refining the List of Nouns

The nouns that appear in the problem description are merely candidates to become classes. It might not be necessary to make classes for them all. The next step is to refine the list to include only the classes that are necessary to solve the particular problem at hand. We will look at the common reasons that a noun can be eliminated from the list of potential classes.

1. Some of the nouns really mean the same thing.

In this example, the following sets of nouns refer to the same thing:

- **car, cars, and foreign cars**

These all refer to the general concept of a car.

- **Joe’s Automotive Shop and shop**

Both of these refer to the company “Joe’s Automotive Shop.”

We can settle on a single class for each of these. In this example we will arbitrarily eliminate **cars** and **foreign cars** from the list, and use the word **car**. Likewise, we will eliminate **Joe’s Automotive Shop** from the list and use the word **shop**. The updated list of potential classes is as follows:

addressforeign carsPorsche

BMWJoe’s Automotive Shop sales tax

carmake service quote

cars	manager	shop
customer	Mercedes	telephone number
estimated labor charges model		total estimated charges
estimated parts charges name		year

Because **car**, **cars**, and **foreign cars** mean the same thing in this problem, we have eliminated **cars** and **foreign cars**. Also, because **Joe's Automotive Shop** and **shop** mean the same thing, we have eliminated **Joe's Automotive Shop**.

2. Some nouns might represent items that we do not need to be concerned with in order to solve the problem.

A quick review of the problem description reminds us of what our application should do: print a service quote. In this example we can eliminate two unnecessary classes from the list:

- We can cross **shop** off the list because our application only needs to be concerned with individual service quotes. It doesn't need to work with or determine any company-wide information. If the problem description asked us to keep a total of all the service quotes, then it would make sense to have a class for the shop.
- We will not need a class for the **manager** because the problem statement does not direct us to process any information about the manager. If there were multiple shop managers, and the problem description had asked us to record which manager generated each service quote, then it would make sense to have a class for the manager.

At this point the updated list of potential classes is as follows:

address	foreign cars	Porsche
BMW	Joe's Automotive Shop sales tax	

car	make	service quote
cars	manager	shop
customer	Mercedes	telephone number
estimated labor charges model		total estimated charges
estimated parts charges name		year

Our problem description does not direct us to process any information about the **shop**, or any information about the **manager**, so we have eliminated those from the list.

3. Some of the nouns might represent objects, not classes.

We can eliminate **Mercedes**, **Porsche**, and **BMW** as classes because, in this example, they all represent specific cars, and can be considered instances of a **car** class. In the description it refers to a specific car brought to the shop by a customer. Therefore, it would also represent an instance of a **cars** class. At this point the updated list of potential classes is as follows:

address	foreign cars	Porsche
BMW	Joe’s Automotive Shop	sales tax
car	make	service quote
cars	manager	shop
customer	Mercedes	telephone number
estimated labor charges model		total estimated charges

car	make	service quote
cars	manager	shop
customer	Mercedes	telephone number
estimated labor charges model		total estimated charges
estimated parts charges name		year

We have eliminated **address**, **estimated labor charges**, **estimated parts charges**, **make**, **model**, **name**, **sales tax**, **telephone number**, **total estimated charges**, and **year** as classes because they represent simple values that can be stored in primitive variables.

As you can see from the list, we have eliminated everything except **cars**, **customer**, and **service quote**. This means that in our application, we will need classes to represent cars, customers, and service quotes. Ultimately, we will write a `Car` class, a `Customer` class, and a `ServiceQuote` class.

Identifying a Class's Responsibilities

Once the classes have been identified, the next task is to identify each class's responsibilities. A class's *responsibilities* are as follows:

- The things that the class is responsible for knowing
- The actions that the class is responsible for doing

When you have identified the things that a class is responsible for knowing, you have identified the class's attributes. These values will be stored in fields. Likewise, when you have identified the actions that a class is responsible for doing, you have identified its methods.

It is often helpful to ask the questions “In the context of this problem, what must the class know? What must the class do?” The first place to look for the answers is in the description of the problem domain. Many of the things that a class must know and do will be mentioned. Some class responsibilities, however, might not be directly mentioned in the problem domain, so brainstorming is often required. Let's apply this methodology to the classes we previously identified from our problem domain.

The Customer class

In the context of our problem domain, what must the `Customer` class know? The description directly mentions the following items, which are all attributes of a customer:

- The customer's name
- The customer's address
- The customer's telephone number

These are all values that can be represented as strings and stored in the class's fields. The `Customer` class can potentially know many other things. One mistake that can be made at this point is to identify too many things that an object is responsible for knowing. In some applications, a `Customer` class might know the customer's email address. This particular problem domain does not mention that the customer's email address is used for any purpose, so we should not include it as a responsibility.

Now let's identify the class's methods. In the context of our problem domain, what must the `Customer` class do? The only obvious actions are as follows:

- Create an object of the `Customer` class
- Set and get the customer's name
- Set and get the customer's address
- Set and get the customer's telephone number

From this list we can see that the `Customer` class will have a constructor, as well as accessor and mutator methods for each of its fields. [Figure 6-29](#) shows a UML diagram for the `Customer` class.




Figure 6-29 UML diagram for the Customer class

[Figure 6-29 Full Alternative Text](#)

The car Class

In the context of our problem domain, what must an object of the `Car` class know? The following items are all attributes of a car, and are mentioned in the problem domain:

- The car's make
- The car's model
- The car's year

Now let's identify the class's methods. In the context of our problem domain, what must the `Car` class do? Once again, the only obvious actions are the standard set of methods that we will find in most classes (constructors, accessors, and mutators). Specifically, the actions are:

- Create an object of the `Car` class
- Set and get the car's make
- Set and get the car's model
- Set and get the car's year

[Figure 6-30](#) shows a UML diagram for the `Car` class at this point.

Car
– make : String – model : String – year : int
+ Car() + setMake(m : String) : void + setModel(m : String) : void + setYear(y : int) : void + getMake() : String + getModel() : String + getYear() : int

Figure 6-30 UML diagram for the Car class

[Figure 6-30 Full Alternative Text](#)

The ServiceQuote Class

In the context of our problem domain, what must an object of the `ServiceQuote` class know? The problem domain mentions the following items:

- The estimated parts charges
- The estimated labor charges
- The sales tax
- The total estimated charges

Careful thought and a little brainstorming will reveal that two of these items are the results of calculations: sales tax and total estimated charges. These items are dependent on the values of the estimated parts and labor charges. In order to avoid the risk of holding stale data, we will not store these values in fields. Rather, we will provide methods that calculate these values and return them. The other methods that we will need for this class are a constructor and the accessors and mutators for the estimated parts charges and estimated labor charges fields. [Figure 6-31](#) shows a UML diagram for the `ServiceQuote` class.

ServiceQuote
– partsCharges : double – laborCharges : double
+ ServiceQuote() + setPartsCharges(c : double): void + setLaborCharges(c : double): void + getPartsCharges() : double + getLaborCharges() : double + salesTax() : double + totalCharges() : double

Figure 6-31 UML diagram for the

ServiceQuote class

[Figure 6-31 Full Alternative Text](#)

This Is Only the Beginning

You should consider the process that we have discussed in this section merely as a starting point. It's important to realize that designing an object-oriented application is an iterative process. It may take you several attempts to identify all of the classes that you will need, and determine all of their responsibilities. As the design process unfolds, you will gain a deeper understanding of the problem, and consequently you will see ways to improve the design.

Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 6.25 What is a problem domain?
2. 6.26 When designing an object-oriented application, who should write a description of the problem domain?
3. 6.27 How do you identify the potential classes in a problem domain description?
4. 6.28 What are a class's responsibilities?
5. 6.29 What two questions should you ask to determine a class's responsibilities?
6. 6.30 Will all of a class's actions always be directly mentioned in the problem domain description?

See the `Amortization` Class Case Study, available on the book's companion Web site at www.pearsonhighered.com/gaddis, for an in-depth example using this chapter's topics.

6.10 Common Errors to Avoid

- **Putting a semicolon at the end of a method header.** A semicolon never appears at the end of a method header.
- **Declaring a variable to reference an object, but forgetting to use the `new` key word to create the object.** Declaring a variable to reference an object does not create an object. You must use the `new` key word to create the object.
- **Forgetting the parentheses that must appear after the class name, which appears after the `new` key word.** The name of a class appears after the `new` key word, and a set of parentheses appears after the class name. You must write the parentheses even if no arguments are passed to the constructor.
- **Forgetting to provide arguments when a constructor requires them.** When using a constructor that has parameter variables, you must provide arguments for them.
- **Trying to overload methods by giving them different return types.** Overloaded methods must have unique parameter lists.
- **Forgetting to write a no-arg constructor for a class that you want to be able to create instances of without passing arguments to the constructor.** If you write a constructor that accepts arguments, you must also write a no-arg constructor for the same class if you want to be able to create instances of the class without passing arguments to the constructor.
- **Unintentionally declaring a local variable with the same name as a field of the same class in a method.** When a method's local variable has the same name as a field in the same class, the local variable's name shadows the field's name.

Review Questions and Exercises

Multiple Choice and True/False

1. This is a collection of programming statements that specify the fields and methods that a particular type of object may have.
 1. class
 2. method
 3. parameter
 4. instance
2. A class is analogous to a(n) .
 1. house
 2. blueprint
 3. drafting table
 4. architect
3. An object is a(n) .
 1. blueprint
 2. primitive data type
 3. variable
 4. instance of a class
4. This is a class member that holds data.
 1. method
 2. instance
 3. field
 4. constructor

5. This key word causes an object to be created in memory.

1. `create`
2. `new`
3. `object`
4. `construct`

6. This is a method that gets a value from a class's field, but does not change it.

1. accessor
2. constructor
3. void
4. mutator

7. This is a method that stores a value in a field or in some other way changes the value of a field.

1. accessor
2. constructor
3. void
4. mutator

8. When the value of an item is dependent on other data, and that item is not updated when the other data is changed, what has the value become?

1. bitter
2. stale
3. asynchronous
4. moldy

9. This is a method that is automatically called when an instance of a class is created.

1. accessor
2. constructor
3. void

4. mutator
10. When a local variable has the same name as a field, the local variable's name does this to the field's name.
 1. shadows
 2. complements
 3. deletes
 4. merges with
11. This is automatically provided for a class if you do not write one yourself.
 1. accessor method
 2. default instance
 3. default constructor
 4. variable declaration
12. Two or more methods in a class may have the same name, as long as this is different.
 1. their return values
 2. their access specifier
 3. their parameter lists
 4. their memory address
13. The process of matching a method call with the correct method is known as .
 1. matching
 2. binding
 3. linking
 4. connecting
14. A class's responsibilities are .
 1. the objects created from the class
 2. things the class knows

- 3. actions the class performs
 - 4. both b and c
15. **True or False:** The `new` operator creates an instance of a class.
16. **True or False:** Each instance of a class has its own set of instance fields.
17. **True or False:** When you write a constructor for a class, it still has the default constructor that Java automatically provides.
18. **True or False:** A class may not have more than one constructor.
19. **True or False:** To find the classes needed for an object-oriented application, you identify all of the verbs in a description of the problem domain.

Find the Error

1. Find the error in the following class:

```
public class MyClass
{
    private int x;
    private double y;
    public void MyClass(int a, double b)
    {
        x = a;
        y = b;
    }
}
```

2. Assume that the following method is a member of a class. Find the error.

```
public void total(int value1, value2, value3)
{
    return value1 + value2 + value3;
}
```

3. The following statement attempts to create a `Rectangle` object. Find the error.

```
Rectangle box = new Rectangle;
```

4. Find the error in the following class:

```
public class TwoValues
{
    private int x, y;
    public TwoValues()
    {
        x = 0;
```

```

    }
    public TwoValues()
    {
        x = 0;
        y = 0;
    }
}

```

5. Find the error in the following class:

```

public class FindTheError
{
    public int square(int number)
    {
        return number * number;
    }
    public double square(int number)
    {
        return number * number;
    }
}

```

Algorithm Workbench

1. Design a class named `Pet`, which should have the following fields:

- `name`. The `name` field holds the name of a pet.
- `animal`. The `animal` field holds the type of animal that a pet is. Example values are “Dog”, “Cat”, and “Bird”.
- `age`. The `age` field holds the pet’s age.

The `Pet` class should also have the following methods:

- `setName`. The `setName` method stores a value in the `name` field.
- `setAnimal`. The `setAnimal` method stores a value in the `animal` field.
- `setAge`. The `setAge` method stores a value in the `age` field.
- `getName`. The `getName` method returns the value of the `name` field.
- `getAnimal`. The `getAnimal` method returns the value of the `animal` field.
- `getAge`. The `getAge` method returns the value of the `age` field.

1. Draw a UML diagram of the class. Be sure to include notation showing each field and method’s access specification and data type. Also include notation showing any

method parameters and their data types.

2. Write the Java code for the `Pet` class.

2. Look at the following partial class definition, and then respond to the questions that follow it:

```
public class Book
{
    private String title;
    private String author;
    private String publisher;
    private int copiesSold;
}
```

1. Write a constructor for this class. The constructor should accept an argument for each of the fields.
2. Write accessor and mutator methods for each field.
3. Draw a UML diagram for the class, including the methods you have written.

3. Consider the following class declaration:

```
public class Square
{
    private double sideLength;
    public double getArea()
    {
        return sideLength * sideLength;
    }
    public double getSideLength()
    {
        return sideLength;
    }
}
```

1. Write a no-arg constructor for this class. It should assign the `sideLength` field the value 0.0.
2. Write an overloaded constructor for this class. It should accept an argument that is copied into the `sideLength` field.

4. Look at the following description of a problem domain:

5. The bank offers the following types of accounts to its customers: savings accounts, checking accounts, and money market accounts. Customers are allowed to deposit money into an account (thereby increasing its balance), withdraw money from an account (thereby decreasing its balance), and earn interest on the account. Each account has an interest rate.
6. Assume that you are writing an application that will calculate the amount of interest earned for a bank account.

1. Identify the potential classes in this problem domain.
2. Refine the list to include only the necessary class or classes for this problem.
3. Identify the responsibilities of the class or classes.

Short Answer

1. What is the difference between a class and an instance of a class?
2. A contractor uses a blueprint to build a set of identical houses. Are classes analogous to the blueprint or the houses?
3. What is an accessor method? What is a mutator method?
4. Is it a good idea to make fields private? Why or why not?
5. If a class has a private field, what has access to the field?
6. What is the purpose of the `new` key word?
7. Assume a program named `MailList.java` is stored in the `DataBase` folder on your hard drive. The program creates objects of the `Customer` and `Account` classes. Describe the steps that the compiler goes through in locating and compiling the `Customer` and `Account` classes.
8. Why are constructors useful for performing “start-up” operations?
9. Under what circumstances does Java automatically provide a default constructor for a class?
10. What do you call a constructor that accepts no arguments?
11. When the same name is used for two or more methods in the same class, how does Java tell them apart?
12. How does method overloading improve the usefulness of a class?

Programming Challenges

MyProgrammingLab™ Visit www.myprogramminglab.com to complete many of these Programming Challenges online and get instant feedback.

1. Employee Class

Write a class named `Employee` that has the following fields:

- `name`. The `name` field references a `String` object that holds the employee's name.
- `idNumber`. The `idNumber` is an `int` variable that holds the employee's ID number.
- `department`. The `department` field references a `String` object that holds the name of the department where the employee works.
- `position`. The `position` field references a `String` object that holds the employee's job title.

The class should have the following constructors:

- A constructor that accepts the following values as arguments and assigns them to the appropriate fields: employee's name, employee's ID number, department, and position.
- A constructor that accepts the following values as arguments and assigns them to the appropriate fields: employee's name and ID number. The `department` and `position` fields should be assigned an empty string (`""`).
- A no-arg constructor that assigns empty strings (`""`) to the `name`, `department`, and `position` fields, and 0 to the `idNumber` field.

Write appropriate mutator methods that store values in these fields and accessor methods that return the values in these fields. Once you have written the class, write a separate program that creates three `Employee` objects to hold the following data:

Name	ID Number	Department	Position
Susan Meyers	47899	Accounting	Vice President
Mark Jones	39119	IT	Programmer

The program should store this data in the three objects and then display the data for each employee on the screen.

2. Car Class

Write a class named `Car` that has the following fields:

- **yearModel.** The `yearModel` field is an `int` that holds the car's year model.
- **make.** The `make` field references a `String` object that holds the make of the car.
- **speed.** The `speed` field is an `int` that holds the car's current speed.

In addition, the class should have the following constructor and other methods.

- **Constructor.** The constructor should accept the car's year model and make as arguments. These values should be assigned to the object's `yearModel` and `make` fields. The constructor should also assign 0 to the `speed` field.
- **Accessors.** Appropriate accessor methods should get the values stored in an object's `yearModel`, `make`, and `speed` fields.
- **accelerate.** The `accelerate` method should add 5 to the `speed` field each time it is called.
- **brake.** The `brake` method should subtract 5 from the `speed` field each time it is called.

Demonstrate the class in a program that creates a `Car` object, and then calls the `accelerate` method five times. After each call to the `accelerate` method, get the current speed of the car and display it. Then call the `brake` method five times. After each call to the `brake` method, get the current speed of the car and display it.

3. Personal Information Class



The Personal Information Class Problem

Design a class that holds the following personal data: name, address, age, and phone number. Write

appropriate accessor and mutator methods. Demonstrate the class by writing a program that creates three instances of it. One instance should hold your information, and the other two should hold your friends' or family members' information.

4. RetailItem Class

Write a class named `RetailItem` that holds data about an item in a retail store. The class should have the following fields:

- `description`. The `description` field references a `String` object that holds a brief description of the item.
- `unitsOnHand`. The `unitsOnHand` field is an `int` variable that holds the number of units currently in inventory.
- `price`. The `price` field is a `double` that holds the item's retail price.

Write a constructor that accepts arguments for each field, appropriate mutator methods that store values in these fields, and accessor methods that return the values in these fields. Once you have written the class, write a separate program that creates three `RetailItem` objects and stores the following data in them:

Description Units on Hand Price		
Item #1 Jacket	12	59.95
Item #2 Designer Jeans	40	34.95
Item #3 Shirt	20	24.95

5. Payroll Class

Design a `Payroll` class that has fields for an employee's name, ID number, hourly pay rate, and number of hours worked. Write the appropriate accessor and mutator methods and a constructor that accepts the employee's name and ID number as arguments. The class should also have a method that returns the employee's gross pay, which is calculated as the number of hours worked multiplied by the hourly pay rate. Write a program that demonstrates the class by creating a `Payroll` object, then asking the user to enter the data for an employee. The program should display the amount of gross pay

earned.

6. TestScores Class

Design a `TestScores` class that has fields to hold three test scores. The class should have a constructor, accessor and mutator methods for the test score fields, and a method that returns the average of the test scores. Demonstrate the class by writing a separate program that creates an instance of the class. The program should ask the user to enter three test scores, which are stored in the `TestScores` object. Then the program should display the average of the scores, as reported by the `TestScores` object.

7. Circle Class

Write a `Circle` class that has the following fields:

- **radius:** a `double`
- **PI:** a `final double` initialized with the value 3.14159

The class should have the following methods:

- **Constructor.** Accepts the radius of the circle as an argument.
- **Constructor.** A no-arg constructor that sets the `radius` field to 0.0.
- **setRadius.** A mutator method for the radius field.
- **getRadius.** An accessor method for the radius field.
- **area.** Returns the area of the circle, which is calculated as

```
area = PI * radius * radius
```

- **diameter.** Returns the diameter of the circle, which is calculated as

```
diameter = radius * 2
```

- **circumference.** Returns the circumference of the circle, which is calculated as

```
circumference = 2 * PI * radius
```

Write a program that demonstrates the `Circle` class by asking the user for the circle's radius, creating a `Circle` object, and then reporting the circle's area, diameter, and circumference.

8. Temperature Class

Write a `Temperature` class that will hold a temperature in Fahrenheit, and provide methods to get the temperature in Fahrenheit, Celsius, and Kelvin. The class should have the following field:

- `fahrenheit` – A `double` that holds a Fahrenheit temperature.

The class should have the following methods:

- **Constructor** – The constructor accepts a Fahrenheit temperature (as a `double`) and stores it in the `ftemp` field.
- `setFahrenheit` – The `setFahrenheit` method accepts a Fahrenheit temperature (as a `double`) and stores it in the `ftemp` field.
- `getFahrenheit` – Returns the value of the `ftemp` field, as a Fahrenheit temperature (no conversion required).
- `celsius` – Returns the value of the `ftemp` field converted to Celsius.
- `kelvin` – Returns the value of the `ftemp` field converted to Kelvin.

Use the following formula to convert the Fahrenheit temperature to Celsius:

$$\text{Celsius} = (5 / 9) \times (\text{Fahrenheit} - 32)$$

Use the following formula to convert the Fahrenheit temperature to Kelvin:

$$\text{Kelvin} = ((5 / 9) \times (\text{Fahrenheit} - 32)) + 273$$

Demonstrate the `Temperature` class by writing a separate program that asks the user for a Fahrenheit temperature. The program should create an instance of the `Temperature` class, with the value entered by the user passed to the constructor. The program should then call the object's methods to display the temperature in Celsius and Kelvin.

9. Days in a Month

Write a class named `MonthDays`. The class's constructor should accept two arguments:

- An integer for the month (1=January, 2 February, etc.).
- An integer for the year

The class should have a method named `numberOfDays` that returns the number of days in the specified

month. The method should use the following criteria to identify leap years:

1. Determine whether the year is divisible by 100. If it is, then it is a leap year if and only if it is divisible by 400. For example, 2000 is a leap year but 2100 is not.
2. If the year is not divisible by 100, then it is a leap year if and only if it is divisible by 4. For example, 2008 is a leap year but 2009 is not.

Demonstrate the class in a program that asks the user to enter the month (letting the user enter an integer in the range of 1 through 12) and the year. The program should then display the number of days in that month. Here is a sample run of the program:

```
Enter a month (1-12): 2 [Enter]
Enter a year: 2008 [Enter]
29 days
```

10. A Game of Twenty-One

For this assignment, you will write a program that lets the user play against the computer in a variation of the popular blackjack card game. In this variation of the game, two six-sided dice are used instead of cards. The dice are rolled, and the player tries to beat the computer's hidden total without going over 21.

Here are some suggestions for the game's design:

- Each round of the game is performed as an iteration of a loop that repeats as long as the player agrees to roll the dice, and the player's total does not exceed 21.
- At the beginning of each round, the program will ask the user whether or not he or she wants to roll the dice to accumulate points.
- During each round, the program simulates the rolling of two six-sided dice. It rolls the dice first for the computer, and then it asks the user whether he or she wants to roll. (Use the `Die` class that was shown in [Code Listing 6-14](#) to simulate the dice.)
- The loop keeps a running total of both the computer's and the user's points.
- The computer's total should remain hidden until the loop has finished.
- After the loop has finished, the computer's total is revealed, and the player with the most points, without going over 21, wins.

11. Freezing and Boiling Points

The following table lists the freezing and boiling points of several substances.

Substance	Freezing Point	Boiling Point
Ethyl Alcohol	−173	172
Oxygen	−362	−306
Water	32	212

Design a class that stores a temperature in a `temperature` field and has the appropriate accessor and mutator methods for the field. In addition to appropriate constructors, the class should have the following methods:

- **`isEthylFreezing`**. This method should return the `boolean` value `true` if the temperature stored in the `temperature` field is at or below the freezing point of ethyl alcohol. Otherwise, the method should return `false`.
- **`isEthylBoiling`**. This method should return the `boolean` value `true` if the temperature stored in the `temperature` field is at or above the boiling point of ethyl alcohol. Otherwise, the method should return `false`.
- **`isOxygenFreezing`**. This method should return the `boolean` value `true` if the temperature stored in the `temperature` field is at or below the freezing point of oxygen. Otherwise, the method should return `false`.
- **`isOxygenBoiling`**. This method should return the `boolean` value `true` if the temperature stored in the `temperature` field is at or above the boiling point of oxygen. Otherwise, the method should return `false`.
- **`isWaterFreezing`**. This method should return the `boolean` value `true` if the temperature stored in the `temperature` field is at or below the freezing point of water. Otherwise, the method should return `false`.
- **`isWaterBoiling`**. This method should return the `boolean` value `true` if the temperature stored in the `temperature` field is at or above the boiling point of water. Otherwise, the method should return `false`.

Write a program that demonstrates the class. The program should ask the user to enter a temperature, and then display a list of the substances that will freeze at that temperature and those that will boil at that temperature. For example, if the temperature is `−20` the class should report that water will freeze

and oxygen will boil at that temperature.

12. SavingsAccount Class

Design a `SavingsAccount` class that stores a savings account's annual interest rate and balance. The class constructor should accept the amount of the savings account's starting balance. The class should also have methods for subtracting the amount of a withdrawal, adding the amount of a deposit, and adding the amount of monthly interest to the balance. The monthly interest rate is the annual interest rate divided by twelve. To add the monthly interest to the balance, multiply the monthly interest rate by the balance, and add the result to the balance.

Test the class in a program that calculates the balance of a savings account at the end of a period of time. It should ask the user for the annual interest rate, the starting balance, and the number of months that have passed since the account was established. A loop should then iterate once for every month, performing the following:

1. Ask the user for the amount deposited into the account during the month. Use the class method to add this amount to the account balance.
2. Ask the user for the amount withdrawn from the account during the month. Use the class method to subtract this amount from the account balance.
3. Use the class method to calculate the monthly interest.

After the last iteration, the program should display the ending balance, the total amount of deposits, the total amount of withdrawals, and the total interest earned.

13. Deposit and Withdrawal Files

Use Notepad or another text editor to create a text file named `Deposits.txt`. The file should contain the following numbers, one per line:

```
100.00
124.00
78.92
37.55
```

Next, create a text file named `Withdrawals.txt`. The file should contain the following numbers, one per line:

```
29.88
110.00
27.52
50.00
12.90
```

The numbers in the `Deposits.txt` file are the amounts of deposits that were made to a savings account during the month, and the numbers in the `Withdrawals.txt` file are the amounts of withdrawals that were made during the month. Write a program that creates an instance of the `SavingsAccount` class that you wrote in Programming Challenge 12. The starting balance for the object is 500.00. The program should read the values from the `Deposits.txt` file and use the object's method to add them to the account balance. The program should read the values from the `Withdrawals.txt` file and use the object's method to subtract them from the account balance. The program should call the class method to calculate the monthly interest, and then display the ending balance and the total interest earned.

14. Patient Charges

Write a class named `Patient` that has fields for the following data:

- First name, middle name, and last name
- Address, city, state, and ZIP code
- Phone number
- Name and phone number of emergency contact

The `Patient` class should have a constructor, accessor, and mutator that accept an argument for each field.

Next, write a class named `Procedure` that represents a medical procedure that has been performed on a patient. The `Procedure` class should have fields for the following data:

- Name of the procedure
- Date of the procedure
- Name of the practitioner who performed the procedure
- Charges for the procedure

The `Procedure` class should have a constructor, accessor, and mutator that accept an argument for each field.

Next, write a program that creates an instance of the `Patient` class, initialized with sample data. Then, create three instances of the `Procedure` class, initialized with the following data:

Procedure #1:

Procedure #2:

Procedure #3:

Procedure name: Physical Exam Procedure name: X-ray Procedure name: Blood test

Date: Today’s date Date: Today’s date Date: Today’s date

Practitioner: Dr. Irvine Practitioner: Dr. Jamison Practitioner: Dr. Smith

Charge: 250.00 Charge: 500.00 Charge: 200.00

The program should display the patient’s information, information about all three of the procedures, and the total charges of the three procedures.

15. Dice Game

Write a program that uses the `Die` class that was presented in this chapter to play a simple dice game between the computer and the user. The program should create two instances of the `Die` class (each a 6-sided die). One `Die` object is the computer’s die, and the other `Die` object is the user’s die.

The program should have a loop that iterates 10 times. Each time the loop iterates, it should roll both dice. The die with the highest value wins. (In case of a tie, there is no winner for that particular roll of the dice.)

As the loop iterates, the program should keep count of the number of times the computer wins, and the number of times that the user wins. After the loop performs all of its iterations, the program should display who was the grand winner, the computer or the user.

16. Roulette Wheel Colors

On a roulette wheel, the pockets are numbered from 0 to 36. The colors of the pockets are as follows:

- Pocket 0 is green.
- For pockets 1 through 10, the odd numbered pockets are red and the even numbered pockets are black.
- For pockets 11 through 18, the odd numbered pockets are black and the even numbered pockets are red.
- For pockets 19 through 28, the odd numbered pockets are red and the even numbered pockets are black.
- For pockets 29 through 36, the odd numbered pockets are black and the even numbered pockets are red.

Write a class named `RoulettePocket`. The class's constructor should accept a pocket number. The class should have a method named `getPocketColor` that returns the pocket's color, as a string.

Demonstrate the class in a program that asks the user to enter a pocket number, and displays whether the pocket is green, red, or black. The program should display an error message if the user enters a number that is outside the range of 0 through 36.

17. Coin Toss Simulator

Write a class named `Coin`. The `Coin` class should have the following field:

- A `String` named `sideUp`. The `sideUp` field will hold either “heads” or “tails” indicating the side of the coin that is facing up.

The `Coin` class should have the following methods:

- A no-arg constructor that randomly determines the side of the coin that is facing up (“heads” or “tails”) and initializes the `sideUp` field accordingly.
- A `void` method named `toss` that simulates the tossing of the coin. When the `toss` method is called, it randomly determines the side of the coin that is facing up (“heads” or “tails”) and sets the `sideUp` field accordingly.
- A method named `getSideUp` that returns the value of the `sideUp` field.

Write a program that demonstrates the `Coin` class. The program should create an instance of the class and display the side that is initially facing up. Then, use a loop to toss the coin 20 times. Each time the coin is tossed, display the side that is facing up. The program should keep count of the number of times heads is facing up and the number of times tails is facing up, and display those values after the loop finishes.

18. Tossing Coins for a Dollar

For this assignment you will create a game program using the `Coin` class from Programming Challenge 16. The program should have three instances of the `Coin` class: one representing a quarter, one representing a dime, and one representing a nickel.

When the game begins, your starting balance is \$0. During each round of the game, the program will toss the simulated coins. When a coin is tossed, the value of the coin is added to your balance if it lands heads-up. For example, if the quarter lands heads-up, 25 cents is added to your balance. Nothing is added to your balance for coins that land tails-up. The game is over when your balance reaches one dollar or more. If your balance is exactly one dollar, you win the game. You lose if your balance exceeds one dollar.

19. Fishing Game Simulation

For this assignment, you will write a program that simulates a fishing game. In this game, a six-sided die is rolled to determine what the user has caught. Each possible item is worth a certain number of fishing points. The points will remain hidden until the user is finished fishing, and then a message is displayed congratulating the user, depending on the number of fishing points gained.

Here are some suggestions for the game's design:

- Each round of the game is performed as an iteration of a loop that repeats as long as the player wants to fish for more items.
- At the beginning of each round, the program will ask the user whether or not he or she wants to continue fishing.
- The program simulates the rolling of a six-sided die (use the `Die` class that was shown in [Code Listing 6-14](#)).
- Each item that can be caught is represented by a number generated from the die; for example, 1 for “a huge fish”, 2 for “an old shoe”, 3 for “a little fish”, and so on.
- Each item the user catches is worth a different amount of points.
- The loop keeps a running total of the user's fishing points.
- After the loop has finished, the total number of fishing points is displayed, along with a message that varies depending on the number of points earned.

Chapter 7 Arrays and the ArrayList Class

TOPICS

1. [7.1 Introduction to Arrays](#)
2. [7.2 Processing Array Elements](#)
3. [7.3 Passing Arrays as Arguments to Methods](#)
4. [7.4 Some Useful Array Algorithms and Operations](#)
5. [7.5 Returning Arrays from Methods](#)
6. [7.6 String Arrays](#)
7. [7.7 Arrays of Objects](#)
8. [7.8 The Sequential Search Algorithm](#)
9. [7.9 Two-Dimensional Arrays](#)
10. [7.10 Arrays with Three or More Dimensions](#)
11. [7.11 The Selection Sort and the Binary Search Algorithms](#)
12. [7.12 Command-Line Arguments and Variable-Length Argument Lists](#)
13. [7.13 The ArrayList Class](#)
14. [7.14 Common Errors to Avoid](#)

7.1 Introduction to Arrays

CONCEPT: An array can hold multiple values of the same data type simultaneously.

The primitive variables you have worked with so far are designed to hold one value at a time. Each of the variable declarations in [Figure 7-1](#) causes only enough memory to be reserved to hold one value of the specified data type.

`int count;` Enough memory to
 hold one `int`.

1234

`double number;` Enough memory to
 hold one `double`.

1234.55

`char letter;` Enough memory to
 hold one `char`.

A

Figure 7-1 Variable declarations and their memory allocations

[Figure 7-1 Full Alternative Text](#)

An array, however, is an object that can store a group of values, all of the same type. Creating and using an array in Java is similar to creating and using any other type of object: You declare a reference variable and use the `new` key word to create an instance of the array in memory. Here is an example of a statement that declares an array reference variable:

```
int[] numbers;
```

This statement declares `numbers` as an array reference variable. The `numbers` variable can reference an array of `int` values. Notice that this statement looks like a regular `int` variable declaration except for the set of brackets that appears after the key word `int`. The brackets indicate that this variable is a

reference to an `int` array. Declaring an array reference variable does not create an array. The next step in the process is to use the `new` key word to create an array and assign its address to the `numbers` variable. The following statement shows an example:

```
numbers = new int[6];
```

The number inside the brackets is the array's *size declarator*. It indicates the number of elements, or values, the array can hold. When this statement is executed, `numbers` will reference an array that can hold six elements, each one an `int`. This is shown in [Figure 7-2](#).




Figure 7-2 The `numbers` array

As with any other type of object, it is possible to declare a reference variable and create an instance of an array with one statement. Here is an example:

```
int[] numbers = new int[6];
```

Arrays of any data type can be declared. The following are all valid array declarations:

```
float[] temperatures = new float[100];  
char[] letters = new char[41];  
long[] units = new long[50];  
double[] sizes = new double[1200];
```

An array's size declarator must be a non-negative integer expression. It can be a literal value, as shown in the previous examples, or a variable. It is a common practice to use a `final` variable as a size declarator. Here is an example:

```
final int NUM_ELEMENTS = 6;  
int[] numbers = new int[NUM_ELEMENTS];
```


This practice can make programs easier to maintain. When we store the size of an array in a variable, we can use the variable instead of a literal number when we refer to the size of the array. If we ever need to change the array's size, we need only to change the value of the variable. The variable should be `final` so its contents cannot be changed during the program's execution.

 **NOTE:**

Once an array is created, its size cannot be changed.

Accessing Array Elements

Although an array has only one name, the elements in the array may be accessed and used as individual variables. This is possible because each element is assigned a number known as a *subscript*. A subscript is used as an index to pinpoint a specific element within an array. The first element is assigned the subscript 0, the second element is assigned 1, and so forth. The six elements in the `numbers` array (described earlier) would have the subscripts 0 through 5. This is shown in [Figure 7-3](#).



The `numbers` array has six elements, numbered 0 through 5.

Figure 7-3 Subscripts for the `numbers` array

Subscript numbering always starts at zero. The subscript of the last element in an array is one less than the total number of elements in the array. This means that for the `numbers` array, which has six elements, 5 is the subscript for the last element.

Each element in the `numbers` array, when accessed by its subscript, can be used as an `int` variable. For example, look at the following code. The first statement stores 20 in the first element of the array (element 0), and the second statement stores 30 in the fourth element (element 3).

```
numbers[0] = 20;  
numbers[3] = 30;
```

NOTE:

The expression `numbers[0]` is pronounced “numbers sub zero.” You read these assignment statements as “numbers sub zero is assigned twenty” and “numbers sub three is assigned thirty.”

[Figure 7-4](#) illustrates the contents of the array after these statements execute.




Figure 7-4 Contents of the array after 20 is assigned to `numbers[0]` and 30 is assigned to `numbers[3]`

[Figure 7-4 Full Alternative Text](#)

NOTE:

By default, Java initializes array elements with 0. In [Figure 7-4](#), values have not been stored in elements 1, 2, 4, and 5, so they are shown as 0s.

By this point you should understand the difference between the array size declarator and a subscript. When you use the `new` key word to create an array object, the number inside the brackets is the size declarator. It indicates the number of elements in the array. The number inside the brackets in an assignment statement or any statement that works with the contents of an array is a subscript. It is used to access a specific element in the array.

Inputting and Outputting Array Contents

You can read values from the keyboard and store them in an array element just as you can a regular variable. You can also output the contents of an array element with `print` and `println`. [Code Listing 7-1](#) shows an array being used to store and display values entered by the user. [Figure 7-5](#) shows the contents of the `hours` array with the values entered by the user in the example output.

Code Listing 7-1 (`ArrayDemo1.java`)

```
1 import java.util.Scanner;  
2
```

```

3  /**
4      This program shows values being stored in an array's
5      elements and displayed.
6  */
7
8  public class ArrayDemol
9  {
10     public static void main(String[] args)
11     {
12         final int EMPLOYEES = 3;           // Number of employees
13         int[] hours = new int[EMPLOYEES]; // Array of hours
14
15         // Create a Scanner object for keyboard input.
16         Scanner keyboard = new Scanner(System.in);
17
18         System.out.println("Enter the hours worked by " +
19                             EMPLOYEES + " employees.");
20
21         // Get the hours worked by employee 1.
22         System.out.print("Employee 1: ");
23         hours[0] = keyboard.nextInt();
24
25         // Get the hours worked by employee 2.
26         System.out.print("Employee 2: ");
27         hours[1] = keyboard.nextInt();
28
29         // Get the hours worked by employee 3.
30         System.out.print("Employee 3: ");
31         hours[2] = keyboard.nextInt();
32
33         // Display the values entered.
34         System.out.println("The hours you entered are:");
35         System.out.println(hours[0]);
36         System.out.println(hours[1]);
37         System.out.println(hours[2]);
38     }
39 }

```

Program Output with Example Input Shown in Bold

```

Enter the hours worked by 3 employees.
Employee 1: 40 [Enter]
Employee 2: 20 [Enter]
Employee 3: 15 [Enter]
The hours you entered are:
40
20
15

```




Figure 7-5 Contents of the hours array



Accessing Array Elements in a Loop

Subscript numbers can be stored in variables. This makes it possible to use a loop to “cycle through” an entire array, performing the same operation on each element. For example, [Code Listing 7-1](#) could be simplified by using two `for` loops: one for inputting the values into the array and the other for displaying the contents of the array. This is shown in [Code Listing 7-2](#).

Code Listing 7-2 (ArrayDemo2.java)

```
1  import java.util.Scanner;
2
3  /**
4   * This program shows an array being processed with loops.
5   */
6
7  public class ArrayDemo2
8  {
9      public static void main(String[] args)
10     {
11         final int EMPLOYEES = 3;                // Number of employees
12         int[] hours = new int[EMPLOYEES];        // Array of hours
13
14         // Create a Scanner object for keyboard input.
15         Scanner keyboard = new Scanner(System.in);
16
17         System.out.println("Enter the hours worked by " +
18                             EMPLOYEES + " employees.");
19
20         // Get the hours for each employee.
21         for (int index = 0; index < EMPLOYEES; index++)
22         {
23             System.out.print("Employee " + (index + 1) + ": ");
24             hours[index] = keyboard.nextInt();
25         }
26
27         System.out.println("The hours you entered are:");
```

```

28
29 // Display the values entered.
30 for (int index = 0; index < EMPLOYEES; index++)
31     System.out.println(hours[index]);
32 }
33 }

```

Program Output with Example Input Shown in Bold

```

Enter the hours worked by 3 employees.
Employee 1: 40 [Enter]
Employee 2: 20 [Enter]
Employee 3: 15 [Enter]
The hours you entered are:
40
20
15

```

Let's take a closer look at the first loop in this program, which appears in lines 21 through 25. Notice that the loop's control variable, `index`, is used as a subscript in line 24:

```
hours[index] = keyboard.nextInt();
```

The variable `index` starts at 0. During the loop's first iteration, the user's input is stored in `hours[0]`. Then, `index` is incremented, so its value becomes 1. During the next iteration, the user's input is stored in `hours[1]`. This continues until values have been stored in all of the elements of the array. Notice that the loop correctly starts and ends the control variable with valid subscript values (0 through 2), as illustrated in [Figure 7-6](#). This ensures that only valid subscripts are used.




Figure 7-6 Annotated loop

Java Performs Bounds Checking

Java performs array bounds checking, which means that it does not allow a statement to use a subscript that is outside the range of valid subscripts for an array. For example, the following statement creates an array with 10 elements. The valid subscripts for the array are 0 through 9.

```
int[] values = new int[10];
```

Java will not allow a statement to use a subscript that is less than 0 or greater than 9 with this array. Bounds checking occurs at runtime. The Java compiler does not display an error message when it processes a statement that uses an invalid subscript. Instead, when the statement executes, the program throws an exception and immediately terminates. For instance, the program in [Code Listing 7-3](#) declares a three-element array, but attempts to store four values in the array. In line 17, when the program attempts to store a value in `values[3]`, it halts and an error message is displayed.

Code Listing 7-3 (InvalidSubscript.java)

```
1  /**
2   * This program uses an invalid subscript with an array.
3   */
4
5  public class InvalidSubscript
6  {
7      public static void main(String[] args)
8      {
9          int[] values = new int[3];
10
11          System.out.println("I will attempt to store four " +
12                           "numbers in a three-element array.");
13
14          for (int index = 0; index < 4; index++)
15          {
16              System.out.println("Now processing element " + index);
17              values[index] = 10;
18          }
19      }
20  }
```

Program Output

```
I will attempt to store four numbers in a three-element array.
Now processing element 0
Now processing element 1
Now processing element 2
Now processing element 3
Exception in thread "main"
java.lang.ArrayIndexOutOfBoundsException: 3
```



NOTE:

The error message you see may be different, depending on your system.

Watch Out for Off-by-One Errors

Because array subscripts start at 0 rather than 1, you have to be careful not to perform an off-by-one error. For example, look at the following code:

```
// This code has an off-by-one error.
final int SIZE = 100;
int[] numbers = new int[SIZE];
for (int index = 1; index <= SIZE; index++)
    numbers[index] = 0;
```

The intent of this code is to create an array of integers with 100 elements, and store the value 0 in each element. However, this code has an off-by-one error. The loop uses its control variable, `index`, as a subscript with the `numbers` array. During the loop's execution, the variable `index` takes on the values 1 through 100, when it should take on the values 0 through 99. As a result, the first element, which is at subscript 0, is skipped. In addition, the loop attempts to use 100 as a subscript during the last iteration. Because 100 is an invalid subscript, the program will throw an exception and halt.

Array Initialization

Like regular variables, Java allows you to initialize an array's elements when you create the array. Here is an example:

```
int[] days = {31, 28, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31};
```

This statement declares the reference variable `days`, creates an array in memory, and stores initial values in the array. The series of values inside the braces and separated with commas is called an initialization list. These values are stored in the array elements in the order they appear in the list. (The first value, 31, is stored in `days[0]`, the second value, 28, is stored in `days[1]`, and so forth.) Note that you do not use the `new` key word when you use an initialization list. Java automatically creates the array and stores the values in the initialization list in it.

The Java compiler determines the size of the array by the number of items in the initialization list. Because there are 12 items in the example statement's initialization list, the array will have 12 elements. The program in [Code Listing 7-4](#) demonstrates an array being initialized.

Code Listing 7-4 (ArrayInitialization.java)

```
1  /**
2   * This program shows an array being initialized.
3   */
4
5  public class ArrayInitialization
6  {
7      public static void main(String[] args)
8      {
9          int[] days = { 31, 28, 31, 30, 31, 30,
10                       31, 31, 30, 31, 30, 31 };
11
12          for (int index = 0; index < 12; index++)
13          {
14              System.out.println("Month " + (index + 1) +
15                                " has " + days[index] +
16                                " days.");
17          }
18      }
19  }
```

Program Output

```
Month 1 has 31 days.
Month 2 has 28 days.
Month 3 has 31 days.
Month 4 has 30 days.
Month 5 has 31 days.
Month 6 has 30 days.
Month 7 has 31 days.
Month 8 has 31 days.
Month 9 has 30 days.
Month 10 has 31 days.
Month 11 has 30 days.
Month 12 has 31 days.
```

Java allows you to spread the initialization list across multiple lines. Both of the following array declarations are equivalent:

```
double[] coins = { 0.05, 0.1, 0.25 };
double[] coins = { 0.05,
                   0.1,
                   0.25 };
```

Alternate Array Declaration Notation

Java allows you to use two different styles when declaring array reference variables. The first style is the one used in this book, with the brackets immediately following the data type, as shown here:


```
int[] numbers;
```

In the second style the brackets are placed after the variable name, as shown here:

```
int numbers[];
```

Both of these statements accomplish the same thing: They declare that `numbers` is a reference to an `int` array. The difference between the two styles is noticed when more than one variable is declared in the same statement. For example, look at the following statement:

```
int[] numbers, codes, scores;
```

This statement declares three variables: `numbers`, `codes`, and `scores`. All three are references to `int` arrays. This makes perfect sense because `int[]` is the data type for all the variables declared in the statement. Now look at the following statement, which uses the alternate notation:

```
int numbers[], codes, scores;
```

This statement declares the same three variables, but only `numbers` is a reference to an `int` array. The `codes` and `scores` variables are regular `int` variables. This is because `int` is the data type for all the variables declared in the statement, and only `numbers` is followed by the brackets. To declare all three of these variables as references to `int` arrays using the alternate notation, you need to write a set of brackets after each variable name. Here is an example:

```
int numbers[], codes[], scores[];
```

The first style is the standard notation for most Java programmers, so that is the style used in this book.



Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 7.1 Write statements that create the following arrays:

1. A 100-element `int` array referenced by the variable `employeeNumbers`.
2. A 25-element `double` array referenced by the variable `payRates`.
3. A 14-element `float` array referenced by the variable `miles`.
4. A 1000-element `char` array referenced by the variable `letters`.

2. 7.2 What's wrong with the following array declarations?

```
int[] readings = new int[-1];  
double[] measurements = new double[4.5];
```

3. 7.3 What would the valid subscript values be in a four-element array of `doubles`?
4. 7.4 What is the difference between an array's size declarator and a subscript?
5. 7.5 What does it mean for a subscript to be out-of-bounds?
6. 7.6 What happens in Java when a program tries to use a subscript that is out-of-bounds?
7. 7.7 What is the output of the following code?


```
int[] values = new int[5];  
for (int count = 0; count < 5; count++)  
    values[count] = count + 1;  
for (int count = 0; count < 5; count++)  
    System.out.println(values[count]);
```
8. 7.8 Write a statement that creates and initializes a `double` array with the following values: 1.7, 6.4, 8.9, 3.1, and 9.2. How many elements are in the array?

7.2 Processing Array Elements

CONCEPT: Individual array elements are processed like any other type of variable.

Processing array elements is no different from processing other variables. For example, the following statement multiplies `hours[3]` by the variable `payRate`:

```
grossPay = hours[3] * payRate;
```

The following are examples of pre-increment and post-increment operations on array elements:

```
int[] score = {7, 8, 9, 10, 11};
++score[2];      // Pre-increment operation
score[4]++;      // Post-increment operation
```

When using increment and decrement operators, be careful not to use the operator on the subscript when you intend to use it on the array element. For example, the following statement decrements the variable `count`, but does nothing to the value stored in the array element `amount[count]`:

```
amount[count--];
```

[Code Listing 7-5](#) demonstrates the use of array elements in a simple mathematical statement. A loop steps through each element of the array, using the elements to calculate the gross pay of five employees.

Code Listing 7-5 (PayArray.java)

```
1  import java.util.Scanner;
2
3  /**
4   * This program stores in an array the hours worked by
5   * five employees who all make the same hourly wage.
6   */
7
8  public class PayArray
9  {
10     public static void main(String[] args)
11     {
12         final int EMPLOYEES = 5;    // Number of employees
13         double payRate;             // Hourly pay rate
14         double grossPay;            // Gross pay
15
16         // Create an array to hold employee hours.
17         int[] hours = new int[EMPLOYEES];
18
19         // Create a Scanner object for keyboard input.
20         Scanner keyboard = new Scanner(System.in);
21
```

```

22 // Get the hours worked by each employee.
23 System.out.println("Enter the hours worked by " +
24                     EMPLOYEES + " employees who all earn " +
25                     "the same hourly rate.");
26
27 for (int index = 0; index < EMPLOYEES; index++)
28 {
29     System.out.print( "Employee #" + (index + 1) + ": ");
30     hours[index] = keyboard.nextInt();
31 }
32
33 // Get the hourly pay rate.
34 System.out.print("Enter the hourly rate for each employee: ");
35 payRate = keyboard.nextDouble();
36
37 // Display each employee's gross pay.
38 System.out.println( "Here is each employee's gross pay:");
39 for (int index = 0; index < EMPLOYEES; index++)
40 {
41     grossPay = hours[index] * payRate;
42     System.out.println("Employee #" + (index + 1) +
43                       ": $" + grossPay);
44 }
45 }
46 }

```

Program Output with Example Input Shown in Bold

```

Enter the hours worked by 5 employees who all earn the same hourly rate.
Employee #1: 10 [Enter]
Employee #2: 20 [Enter]
Employee #3: 30 [Enter]
Employee #4: 40 [Enter]
Employee #5: 50 [Enter]
Enter the hourly rate for each employee: 10 [Enter]
Here is each employee's gross pay:
Employee #1: $100.0
Employee #2: $200.0
Employee #3: $300.0
Employee #4: $400.0
Employee #5: $500.0

```

In line 41, the following statement assigns the value of `hours[index]` times `payRate` to the `grossPay` variable:

```
grossPay = hours[index] * payRate;
```

Array elements may also be used in relational expressions. For example, the following `if` statement determines whether `cost[20]` is less than `cost[0]`:

```
if (cost[20] < cost[0])
```

And the following `while` loop iterates as long as `value[count]` does not equal 0:

```
while (value[count] != 0)
{
    Statements
}
```

In this chapter's source code (available at www.pearsonhighered.com/gaddis), you will find the file *Overtime.java*, which is a modification of the *PayArray.java* program in [Code Listing 7-5](#). The *Overtime.java* program includes overtime wages in the gross pay. If an employee works more than 40 hours, an overtime pay rate of 1.5 times the regular pay rate is used for the excess hours.

Array Length

Each array in Java has a public field named `length`. This field contains the number of elements in the array. For example, consider an array created by the following statement:

```
double[] temperatures = new double[25];
```

Because the `temperatures` array has 25 elements, the following statement would assign 25 to the variable `size`:

```
size = temperatures.length;
```

The `length` field can be useful when processing the entire contents of an array. For example, the following loop steps through an array and displays the contents of each element. The array's `length` field is used in the test expression as the upper limit for the loop control variable:

```
for (int i = 0; i < temperatures.length; i++)
    System.out.println(temperatures[i]);
```



Be careful not to cause an off-by-one error when using the `length` field as the upper limit of a subscript. The `length` field contains the number of elements in an array. The largest subscript in an array is `length - 1`.



You cannot change the value of an array's `length` field.

The Enhanced for Loop

Java provides a specialized version of the `for` loop that, in many circumstances, simplifies array processing. It is known as the *enhanced for loop*. Here is the general format of the enhanced `for` loop:

```
for (dataType elementVariable : array)
    statement;
```

The enhanced `for` loop is designed to iterate once for every element in an array. Each time the loop iterates, it copies an array element to a variable. Let's look at the syntax more closely as follows:

- `dataType elementVariable` is a variable declaration. This variable will receive the value of a different array element during each loop iteration. During the first loop iteration, it receives the value of the first element; during the second iteration, it receives the value of the second element, and so on. This variable must be of the same data type as the array elements, or a type that the elements can automatically be converted to.
- `array` is the name of an array on which you wish the loop to operate. The loop will iterate once for every element in the array.
- `statement` is a statement that executes during a loop iteration.

For example, assume that we have the following array declaration:

```
int[] numbers = { 3, 6, 9 };
```

We can use the following enhanced `for` loop to display the contents of the `numbers` array:

```
for (int val : numbers)
    System.out.println(val);
```

Because the `numbers` array has three elements, this loop will iterate three times. The first time it iterates, the `val` variable will receive the value in `numbers[0]`. During the second iteration, `val` will receive the value in `numbers[1]`. During the third iteration, `val` will receive the value in `numbers[2]`. The code's output will be as follows:

```
3
6
9
```

If you need to execute more than one statement in the enhanced `for` loop, simply enclose the block of statements in a set of braces. Here is an example:

```
int[] numbers = { 3, 6, 9 };
for (int val : numbers)
{
    System.out.print("The next value is ");
```

```
    System.out.println(val);  
}
```

This code will produce the following output:

```
The next value is 3  
The next value is 6  
The next value is 9
```

The Enhanced `for` Loop versus the Traditional `for` Loop

When you need to access the values that are stored in an array, from the first element to the last element, the enhanced `for` loop is simpler to use than the traditional `for` loop. With the enhanced `for` loop you do not have to be concerned about the size of the array, and you do not have to create an “index” variable to hold subscripts. However, there are circumstances in which the enhanced `for` loop is not adequate. You cannot use the enhanced `for` loop as follows:

- if you need to change the contents of an array element
- if you need to work through the array elements in reverse order
- if you need to access some of the array elements, but not all of them
- if you need to simultaneously work with two or more arrays within the loop
- if you need to refer to the subscript number of a particular element

In any of these circumstances, you should use the traditional `for` loop to process the array.

Letting the User Specify an Array’s Size

Java allows you to use an integer variable to specify an array’s size declarator. This makes it possible to allow the user to specify an array’s size. [Code Listing 7-6](#) demonstrates this, as well as the use of the `length` field. It stores a number of test scores in an array and then displays them.

Code Listing 7-6 (`DisplayTestScores.java`)

```
1  import java.util.Scanner;  
2  
3  /**  
4      This program demonstrates how the user may specify an  
5      array's size.
```

```

6  */
7
8  public class DisplayTestScores
9  {
10     public static void main(String[] args)
11     {
12         int numTests;        // The number of tests
13         int[] tests;         // Array of test scores
14
15         // Create a Scanner object for keyboard input.
16         Scanner keyboard = new Scanner(System.in);
17
18         // Get the number of test scores.
19         System.out.print("How many tests do you have? ");
20         numTests = keyboard.nextInt();
21
22         // Create an array to hold that number of scores.
23         tests = new int[numTests];
24
25         // Get the individual test scores.
26         for (int index = 0; index < tests.length; index++)
27         {
28             System.out.print("Enter test score " +
29                             (index + 1) + ": ");
30             tests[index] = keyboard.nextInt();
31         }
32
33         // Display the test scores.
34         System.out.println();
35         System.out.println("Here are the scores you entered:");
36         for (int index = 0; index < tests.length; index++)
37             System.out.print(tests[index] + " ");
38     }
39 }

```

Program Output with Example Input Shown in Bold

```

How many tests do you have? 5 [Enter]
Enter test score 1: 72 [Enter]
Enter test score 2: 85 [Enter]
Enter test score 3: 81 [Enter]
Enter test score 4: 94 [Enter]
Enter test score 5: 99 [Enter]

```

```

Here are the scores you entered:
72 85 81 94 99

```

This program allows the user to determine the size of the array. In line 23 the following statement creates the array, using the `numTests` variable to determine its size:

```
tests = new int[numTests];
```


The program then uses two `for` loops. The first, in lines 26 through 31, allows the user to input each test score. The second, in lines 36 and 37, displays all of the test scores. Both loops use the `length` member to control their number of iterations as follows:

```
for (int index = 0; index < tests.length; index++)
```

Reassigning Array Reference Variables

It is possible to reassign an array reference variable to a different array, as demonstrated by the following code:

```
// Create an array referenced by the numbers variable.  
int[] numbers = new int[10];  
// Reassign numbers to a new array.  
numbers = new int[5];
```

The first statement creates a ten-element integer array and assigns its address to the `numbers` variable. This is illustrated in [Figure 7-7](#).




Figure 7-7 The numbers variable references a ten-element array

The second statement allocates a five-element integer array and assigns its address to the `numbers` variable. The address of the five-element array takes the place of the address of the ten-element array. After this statement executes, the `numbers` variable references the five-element array instead of the ten-element array. This is illustrated in [Figure 7-8](#). Because the ten-element array is no longer referenced, it cannot be accessed.




Figure 7-8 The `numbers` variable references a five-element array

Copying Arrays

Because an array is an object, there is a distinction between an array and the variable that references it. The array and the reference variable are two separate entities. This is important to remember when you wish to copy the contents of one array to another. You might be tempted to write something like the following code, thinking that you are copying an array:

```
int[] array1 = { 2, 4, 6, 8, 10 };
int[] array2 = array1;    // This does not copy array1.
```

The first statement creates an array and assigns its address to the `array1` variable. The second statement assigns `array1` to `array2`. This does not make a copy of the array referenced by `array1`. Rather, it makes a copy of the address that is stored in `array1` and stores it in `array2`. After this statement executes, both the `array1` and `array2` variables will reference the same array. This type of assignment operation is called a *reference copy*. Only the address of the array object is copied, not the contents of the array object. This is illustrated in [Figure 7-9](#).




Figure 7-9 Both `array1` and `array2` reference the same array

[Code Listing 7-7](#) demonstrates the assigning of an array's address to two reference variables. Regardless of which variable the program uses, it is working with the same array.

Code Listing 7-7 (`SameArray.java`)

```
1  /**
2   * This program demonstrates that two variables can
3   * reference the same array.
```

```

4  */
5
6  public class SameArray
7  {
8      public static void main(String[] args)
9      {
10         int[] array1 = { 2, 4, 6, 8, 10 };
11         int[] array2 = array1;
12
13         // Change one of the elements using array1.
14         array1[0] = 200;
15
16         // Change one of the elements using array2.
17         array2[4] = 1000;
18
19         // Display all the elements using array1
20         System.out.println("The contents of array1:");
21         for (int value : array1)
22             System.out.print(value + " ");
23         System.out.println();
24
25         // Display all the elements using array2
26         System.out.println("The contents of array2:");
27         for (int value : array2)
28             System.out.print(value + " ");
29         System.out.println();
30     }
31 }

```

Program Output

```

The contents of array1:
200 4 6 8 1000
The contents of array2:
200 4 6 8 1000

```

The program in [Code Listing 7-7](#) illustrates that you cannot copy an array by merely assigning one array reference variable to another. Instead, to copy an array you need to copy the individual elements of one array to another. Usually, this is best done with a loop, such as the following:

```

int[] firstArray = { 5, 10, 15, 20, 25 };
int[] secondArray = new int[5];

for (int index = 0; index < firstArray.length; index++)
    secondArray[index] = firstArray[index];

```

The loop in this code copies each element of `firstArray` to the corresponding element of `secondArray`.



1. 7.9 Look at the following statements:

```
int[] numbers1 = { 1, 3, 6, 9 };  
int[] numbers2 = { 2, 4, 6, 8 };  
int result;
```

Write a statement that multiplies element 0 of the `numbers1` array by element 3 of the `numbers2` array and assigns the result to the `result` variable.

2. 7.10 A program uses a variable named `array` that references an array of integers. You do not know the number of elements in the array. Write a `for` loop that stores `-1` in each element of the array.
3. 7.11 A program has the following declaration:

```
double[] values;
```

Write code that asks the user for the size of the array and then creates an array of the specified size, referenced by the `values` variable.

4. 7.12 Look at the following statements:

```
int[] a = { 1, 2, 3, 4, 5, 6, 7 };  
int[] b = new int[7];
```

Write code that copies the `a` array to the `b` array.

7.3 Passing Arrays as Arguments to Methods

CONCEPT: An array can be passed as an argument to a method. To pass an array, you pass the value in the variable that references the array.



Passing an Array to a Method

Quite often you'll want to write methods that process the data in arrays. As you will see, methods can be written to store values in an array, display an array's contents, total all of an array's elements, calculate their average, and so forth. Usually, such methods accept an array as an argument.

When a single element of an array is passed to a method, it is handled like any other variable. For example, [Code Listing 7-8](#) shows a loop that passes each element of the array `numbers` to the method `showValue`.

Code Listing 7-8 (PassElements.java)

```
1  /**
2   * This program demonstrates passing individual array
3   * elements as arguments to a method.
4   */
5
6  public class PassElements
7  {
8      public static void main(String[] args)
9      {
10         int[] numbers = {5, 10, 15, 20, 25, 30, 35, 40};
11
12         for (int index = 0; index < numbers.length; index++)
13             showValue(numbers[index]);
14     }
15
16     /**
17      * The showValue method displays its argument.
18      * @param n The value to display.
19      */
20
21     public static void showValue(int n)
22     {
23         System.out.print(n + " ");
24     }
25 }
```

Program Output

5 10 15 20 25 30 35 40


Each time `showValue` is called in this program, an array element is passed to the method. The `showValue` method has an `int` parameter variable named `n`, which receives the argument. The method simply displays the contents of `n`. If the method were written to accept the entire array as an argument, however, the parameter would have to be set up differently. For example, consider the following method definition. The parameter `array` is declared as an array reference variable. This indicates that the argument will be an array, not a single value.

```
public static void showArray(int[] array)
{
    for (int i = 0; i < array.length; i++)
        System.out.print(array[i] + " ");
}
```

When you pass an array as an argument, you simply pass the value in the variable that references the array, as shown here:

```
showArray(numbers);
```

When an entire array is passed into a method, it is passed just as an object is passed: The actual array itself is not passed, but a reference to the array is passed into the parameter. Consequently, this means the method has direct access to the original array. This is illustrated in [Figure 7-10](#).




```

41
42     // Read values into the array
43     for (int index = 0; index < array.length; index++)
44     {
45         System.out.print("Enter number " +
46                         (index + 1) + ": ");
47         array[index] = keyboard.nextInt();
48     }
49 }
50
51 /**
52     The showArray method accepts an array as
53     an argument and displays its contents.
54     @param array A reference to the array.
55 */
56
57 public static void showArray(int[] array)
58 {
59     // Display the array elements.
60     for (int index = 0; index < array.length; index++)
61         System.out.print(array[index] + " ");
62 }
63 }

```

Program Output with Example Input Shown in Bold

```

Enter a series of 4 numbers.
Enter number 1: 2 [Enter]
Enter number 2: 4 [Enter]
Enter number 3: 6 [Enter]
Enter number 4: 8 [Enter]
Here are the numbers that you entered:
2 4 6 8

```



MyProgrammingLab™ www.myprogramminglab.com

1. 7.13 Look at the following method header:

```
public static void myMethod(double[] array)
```

Here is an array declaration:

```
double[] numbers = new double[100];
```

Write a statement that passes the `numbers` array to the `myMethod` method.

2. 7.14 Write a method named `zero`, which accepts an `int` array as an argument and stores the value 0 in each element.

Chapter 15 Recursion

TOPICS

1. [15.1 Introduction to Recursion](#)
2. [15.2 Solving Problems with Recursion](#)
3. [15.3 Examples of Recursive Methods](#)
4. [15.4 A Recursive Binary Search Method](#)
5. [15.5 The Towers of Hanoi](#)
6. [15.6 Common Errors to Avoid](#)

15.1 Introduction to Recursion

CONCEPT: A recursive method is a method that calls itself.

You have seen instances of methods calling other methods. Method A can call method B, which can then call method C. It's also possible for a method to call itself. A method that calls itself is a recursive method. Look at the `message` method in [Code Listing 15-1](#).

Code Listing 15-1 (EndlessRecursion.java)

```
1  /**
2      This class has a recursive method.
3  */
4
5  public class EndlessRecursion
6  {
7      public static void message()
8      {
9          System.out.println("This is a recursive method.");
10         message();
11     }
12 }
```

This method displays the string "This is a recursive method." and then calls itself. Each time it calls itself, the cycle is repeated. Can you see a problem with the method? There's no way to stop the recursive calls. This method is like an infinite loop because there is no code to stop it from repeating.

Like a loop, a recursive method must have some way to control the number of times it repeats. The class in [Code Listing 15-2](#) has a modified version of the `message` method. It passes an integer argument, which holds the number of times the method should call itself.

Code Listing 15-2 (Recursive.java)

```
1  /**
2      This class has a recursive method, message,
3      which displays a message n times.
4  */
5
6  public class Recursive
7  {
8      public static void message(int n)
9      {
10         if (n > 0)
11         {
12             System.out.println("This is a recursive method.");
```

```

13         message(n - 1);
14     }
15 }
16 }

```

This method contains an `if` statement that controls the repetition. As long as the `n` parameter is greater than zero, the method displays the message and calls itself again. Each time it calls itself, it passes `n - 1` as the argument. For example, look at the program in [Code Listing 15-3](#).

Code Listing 15-3 (RecursionDemo.java)

```

1  /**
2   * This class demonstrates the Recursive.message method.
3   */
4
5  public class RecursionDemo
6  {
7      public static void main(String[] args)
8      {
9          Recursive.message(5);
10     }
11 }
12

```

Program Output

```

This is a recursive method.
This is a recursive method.
This is a recursive method.
This is a recursive method.
This is a recursive method.

```

The `main` method in this class calls the `Recursive.message` method with the argument 5, which causes the method to call itself five times. The first time the method is called, the `if` statement displays the message and then calls itself with 4 as the argument. [Figure 15-1](#) illustrates this.




Figure 15-1 First two calls of the method

[Figure 15-1 Full Alternative Text](#)

The diagram in [Figure 15-1](#) illustrates two separate calls of the `message` method. Each time the method is called, a new instance of the `n` parameter is created in memory. The first time the method is called, the `n` parameter is set to 5. When the method calls itself, a new instance of `n` is created, and the value 4 is passed into it. This cycle repeats until finally, zero is passed to the method. This is illustrated in [Figure 15-2](#).




Figure 15-2 Total of six calls to the message method

[Figure 15-2 Full Alternative Text](#)

As you can see from [Figure 15-2](#), the method is called a total of six times. The first time it is called from the `main` method of the `RecursionDemo` class, and the other five times it calls itself. The number of times that a method calls itself is known as the *depth of recursion*. In this example, the depth of recursion is five. When the method reaches its sixth call, the `n` parameter is set to 0. At that point, the `if` statement's conditional expression is `false`, so the method returns. Control of the program returns from the sixth instance of the method to the point in the fifth instance directly after the recursive method call. This is illustrated in [Figure 15-3](#).




Figure 15-3 Control returns to the point after the recursive method call

[Figure 15-3 Full Alternative Text](#)

Because there are no more statements to be executed after the method call, the fifth instance of the method returns control of the program back to the fourth instance. This repeats until all instances of the method return.

15.2 Solving Problems with Recursion

CONCEPT: A problem can be solved with recursion if it can be broken down into successive smaller problems that are identical to the overall problem.

The `Recursive` and `RecursionDemo` classes shown in the previous section demonstrate the mechanics of a recursive method. Recursion can be a powerful tool for solving repetitive problems and is an important topic in upper level computer science courses. What might not be clear to you yet is how to use recursion to solve a problem. First, it should be noted that recursion is never absolutely required to solve a problem. Any problem that can be solved recursively can also be solved iteratively, with a loop. In fact, recursive algorithms are usually less efficient than iterative algorithms. This is because a method call requires several actions to be performed by the JVM. These actions include allocating memory for parameters and local variables, and storing the address of the program location where control returns after the method terminates. These actions, which are sometimes referred to as overhead, take place with each method call. Such overhead is not necessary with a loop.

Some repetitive problems, however, are more easily solved with recursion than with iteration. Where an iterative algorithm might result in faster execution time, the programmer might be able to design a recursive algorithm faster.

In general, a recursive method works like this:

- If the problem can be solved now, without recursion, then the method solves it and returns.
- If the problem cannot be solved now, then the method reduces it to a smaller but similar problem and calls itself to solve the smaller problem.



Reducing a Problem with Recursion

In order to apply this approach, first we identify at least one case in which the problem can be solved without recursion. This is known as the *base case*. Second, we determine a way to solve the problem in all other circumstances using recursion. This is called the *recursive case*. In the recursive case, we must always reduce the problem to a smaller version of the original problem. By reducing the problem with each recursive call, the base case will eventually be reached and the recursion will stop.

Let's take an example from mathematics to examine an application of recursion. In mathematics, the notation $n!$ represents the factorial of the number n . The factorial of a nonnegative number can be defined by the following rules:

If $n=0$ then $n!=1$

If $n>0$ then $n!=1\times 2\times 3\times \dots\times n$

Let's replace the notation $n!$ with $\text{factorial}(n)$, which looks a bit more like computer code, and rewrite these rules as follows:

If $n=0$ then $\text{factorial}(n)=1$

If $n>0$ then $\text{factorial}(n)=1\times 2\times 3\times \dots\times n$

These rules state that when n is 0, its factorial is 1. When n is greater than 0, its factorial is the product of all the positive integers from 1 up to n . For instance, $\text{factorial}(6)$ is calculated as $1\times 2\times 3\times 4\times 5\times 6$.

When designing a recursive algorithm to calculate the factorial of any number, first we identify the base case, which is the part of the calculation that we can solve without recursion. That is the case where n is equal to 0 as follows:

If $n=0$ then $\text{factorial}(n)=1$

This tells how to solve the problem when n is equal to 0, but what do we do when n is greater than 0? That is the recursive case, or the part of the problem that we use recursion to solve. This is how we express it:

If $n>0$ then $\text{factorial}(n)=n\times \text{factorial}(n-1)$

This states that if n is greater than 0, the factorial of n is n times the factorial of $n-1$. Notice how the recursive call works on a reduced version of the problem, $n-1$. So, our recursive rule for calculating the factorial of a number might look like this:

If $n=0$ then $\text{factorial}(n)=1$

If $n>0$ then $\text{factorial}(n)=n\times \text{factorial}(n-1)$

The following code shows how this might be implemented in a Java method:

```
private static int factorial(int n)
{
    if (n == 0)
        return 1;    // Base case
    else
        return n * factorial(n - 1);
}
```

The program in [Code Listing 15-4](#) demonstrates the method. [Figure 15-4](#) shows example interaction with the program.



Figure 15-4 Interaction with the FactorialDemo.java program

(Oracle Corporate Counsel)

Code Listing 15-4 (FactorialDemo.java)

[illegible]

```

18     number = Integer.parseInt(input);
19
20     // Display the factorial of the number.
21     JOptionPane.showMessageDialog(null,
22         number + "! is " + factorial(number));
23
24     System.exit(0);
25 }
26
27 /**
28     The factorial method uses recursion to calculate
29     the factorial of its argument, which is assumed
30     to be a nonnegative number.
31     @param n The number to use in the calculation.
32     @return The factorial of n.
33 */
34
35 private static int factorial(int n)
36 {
37     if (n == 0)
38         return 1;    // Base case
39     else
40         return n * factorial(n - 1);
41 }
42 }


```

In the example run of the program, the factorial method is called with the argument 4 passed into `n`. Because `n` is not equal to 0, the `if` statement's `else` clause executes the following statement:

```
return n * factorial(n - 1);
```

Although this is a `return` statement, it does not immediately return. Before the return value can be determined, the value of `factorial(n - 1)` must be determined. The `factorial` method is called recursively until the fifth call, in which the `n` parameter will be set to zero. The diagram in [Figure 15-5](#) illustrates the value of `n` and the return value during each call of the method.

The method is first called from the `main` method. →



The second through fifth calls are recursive.







Figure 15-5 Recursive calls to the factorial method (Oracle Corporate Counsel)

[Figure 15-5 Full Alternative Text](#)

This diagram illustrates why a recursive algorithm must reduce the problem with each recursive call.

Eventually the recursion has to stop in order for a solution to be reached. If each recursive call works on a smaller version of the problem, then the recursive calls work toward the base case. The base case does not require recursion, so it stops the chain of recursive calls.

Usually, a problem is reduced by making the value of one or more parameters smaller with each recursive call. In our `factorial` method, the value of the parameter `n` gets closer to 0 with each recursive call. When the parameter reaches 0, the method returns a value without making another recursive call.

Direct and Indirect Recursion

The examples we have discussed so far show recursive methods that directly call themselves. This is known as direct recursion. There is also the possibility of creating indirect recursion in a program. This occurs when method A calls method B, which in turn calls method A. There can even be several methods involved in the recursion. For example, method A could call method B, which could call method C, which calls method A.

Checkpoint

MyProgrammingLab™ www.myprogramminglab.com

1. 15.1 It is said that a recursive algorithm has more overhead than an iterative algorithm. What does this mean?
2. 15.2 What is a base case?
3. 15.3 What is a recursive case?
4. 15.4 What causes a recursive algorithm to stop calling itself?
5. 15.5 What is direct recursion? What is indirect recursion?

15.3 Examples of Recursive Methods

Summing a Range of Array Elements with Recursion

In this example we look at a method, `rangeSum`, that uses recursion to sum a range of array elements. The method takes the following arguments: an `int` array that contains the range of elements to be summed, an `int` specifying the starting element of the range, and an `int` specifying the ending element of the range. Here is an example of how the method might be used:

```
int[] numbers = {1, 2, 3, 4, 5, 6, 7, 8, 9};
int sum;
sum = rangeSum(numbers, 3, 7);
```

This code specifies that `rangeSum` should return the sum of elements 3 through 7 in the `numbers` array. The return value, which in this case would be 30, is stored in `sum`. Here is the definition of the `rangeSum` method:

```
public static int rangeSum(int[] array, int start, int end)
{
    if (start > end)
        return 0;
    else
        return array[start] + rangeSum(array, start + 1, end);
}
```

This method's base case is when the `start` parameter is greater than the `end` parameter. If this is true, the method returns the value 0. Otherwise, the method executes the following statement:

```
return array[start] + rangeSum(array, start + 1, end);
```

This statement returns the sum of `array[start]` plus the return value of a recursive call. Notice that in the recursive call, the starting element in the range is `start + 1`. In essence, this statement says “return the value of the first element in the range plus the sum of the rest of the elements in the range.” The program in [Code Listing 15-5](#) demonstrates the method.

Code Listing 15-5 (RangeSum.java)

```
1  /**
2   * This program demonstrates the recursive rangeSum method.
3   */
4
5  public class RangeSum
```

```

6  {
7
8      public static void main(String[] args)
9      {
10         int[] numbers = { 1, 2, 3, 4, 5, 6, 7, 8, 9 };
11
12         System.out.print("The sum of elements 2 through " +
13                          "5 is "+ rangeSum(numbers, 2, 5));
14     }
15
16     /**
17         The rangeSum method calculates the sum of a specified
18         range of elements in array.
19         @param start Specifies the starting element.
20         @param end Specifies the ending element.
21         @return The sum of the range.
22     */
23
24     public static int rangeSum(int[] array, int start, int end)
25     {
26         if (start > end)
27             return 0;
28         else
29             return array[start] +
30                    rangeSum(array, start + 1, end);
31     }
32 }

```

Program Output

The sum of elements 2 through 5 is 18

Drawing Concentric Circles

In this example we look at the `Circles` application, which uses recursion to draw concentric circles. Concentric circles are circles that share a common center point. [Figure 15-6](#) shows the application's output. The code is shown in [Code Listing 15-6](#).




Figure 15-6 circles applet

(Oracle Corporate Counsel)

Code Listing 15-6 (Circles.java)

```
1  import javafx.application.Application;
2  import javafx.stage.Stage;
3  import javafx.scene.Scene;
4  import javafx.scene.layout.Pane;
5  import javafx.scene.shape.Circle;
6  import javafx.scene.paint.Color;
7
8  public class Circles extends Application
9  {
10     public static void main(String[] args)
11     {
12         launch(args);
13     }
14
15     @Override
16     public void start(Stage primaryStage)
17     {
18         // Constants for the scene size
19         final double SCENE_WIDTH = 440.0;
```

```

20     final double SCENE_HEIGHT = 440.0;
21
22     // Constants for the starting values
23     final double CENTER_X = 220.0;
24     final double CENTER_Y = 220.0;
25     final double RAD = 20.0;
26     final int NUM_CIRCLES = 10;
27
28     // Create an empty Pane.
29     Pane pane = new Pane();
30
31     // Recursively add 10 circles to the Pane.
32     drawCircles(pane, NUM_CIRCLES, CENTER_X, CENTER_Y, RAD);
33
34     // Create a Scene and display it.
35     Scene scene = new Scene(pane, SCENE_WIDTH, SCENE_HEIGHT);
36     primaryStage.setScene(scene);
37     primaryStage.show();
38 }
39
40 /**
41  * The drawCircles method draws concentric circles.
42  * It accepts the following arguments:
43  * p -- a Pane object to add the circles to
44  * n -- the number of circles to draw
45  * x -- x coordinate of the circle's center point
46  * y -- y coordinate of the circle's center point
47  * rad -- the circle's radius
48  */
49
50 private void drawCircles(Pane p, int n, double x, double y, double rad)
51 {
52     if (n > 0)
53     {
54         Circle circle = new Circle(x, y, rad); // Create the circle
55         circle.setStroke(Color.BLACK); // Line color is black
56         circle.setFill(null); // No fill color
57         p.getChildren().addAll(circle); // Add the it to the Pane
58         drawCircles(p, n - 1, x, y, rad + 20); // Draw the next circle
59     }
60 }
61 }

```

The `drawCircles` method, which is called from the application's `start` method, uses recursion to draw the concentric circles. The `n` parameter holds the number of circles to draw. If this parameter is set to 0, the method has reached its base case. Otherwise, line 54 creates a `Circle` object at the specified center point, using the specified radius. Line 55 sets the `Circle`'s stroke color to black, and line 56 sets the `Circle`'s fill color to `null` (the circle will not be filled with a color). Line 57 adds the `Circle` object to the `Pane`. Line 58 makes a recursive call to the `drawCircles` method with the parameters adjusted for the next circle.

The Fibonacci Series

Some mathematical problems are designed to be solved recursively. One well-known example is the calculation of Fibonacci numbers. The Fibonacci numbers, named after the Italian mathematician Leonardo Fibonacci (born circa 1170), are the following sequence:

- 0, 1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 144, 233, . . .

Notice that after the second number, each number in the series is the sum of the two previous numbers. The Fibonacci series can be defined as follows:

If $n=0$ then $\text{Fib}(n)=0$

If $n=1$ then $\text{Fib}(n)=1$

If $n \geq 2$ then $\text{Fib}(n)=\text{Fib}(n-1)+\text{Fib}(n-2)$

A recursive Java method to calculate the n th number in the Fibonacci series is shown here:

```
public static int fib(int n)
{
    if (n == 0)
        return 0;
    else if (n == 1)
        return 1;
    else
        return fib(n - 1) + fib(n - 2);
}
```

Notice that this method actually has two base cases: when n is equal to 0 and when n is equal to 1. In either case, the method returns a value without making a recursive call. The program in [Code Listing 15-7](#) demonstrates this method by displaying the first 10 numbers in the Fibonacci series.

Code Listing 15-7 (FibNumbers.java)

```
1  /**
2   * This program demonstrates the recursive fib method.
3   */
4
5  public class FibNumbers
6  {
7      public static void main(String[] args)
8      {
9          System.out.println("The first 10 numbers in " +
10                           "the Fibonacci series are:");
11
12         for (int i = 0; i < 10; i++)
```

```

13         System.out.print(fib(i) + " ");
14
15     System.out.println();
16 }
17
18 /**
19     The fib method calculates the nth
20     number in the Fibonacci series.
21     @param n The nth number to calculate.
22     @return The nth number.
23 */
24
25 public static int fib(int n)
26 {
27     if (n == 0)
28         return 0;
29     else if (n == 1)
30         return 1;
31     else
32         return fib(n - 1) + fib(n - 2);
33 }
34 }

```

Program Output

The first 10 numbers in the Fibonacci series are:
0 1 1 2 3 5 8 13 21 34

Finding the Greatest Common Divisor

Our next example of recursion is the calculation of the greatest common divisor, or GCD, of two numbers. The GCD of two positive integers, x and y , is as follows:

- if y divides x evenly, then $\text{gcd}(x,y)=y$
- Otherwise, $\text{gcd}(x,y)=\text{gcd}(y, \text{remainder of } x/y)$

This definition states that the GCD of x and y is y if x/y has no remainder. This is the base case. Otherwise, the answer is the GCD of y and the remainder of x/y . The program in [Code Listing 15-8](#) shows a recursive method for calculating the GCD.

Code Listing 15-8 (GCDdemo.java)

```

1  import java.util.Scanner;
2
3  /**
4      This program demonstrates the recursive gcd method.

```

```

5  */
6
7  public class GCDdemo
8  {
9      public static void main(String[] args)
10     {
11         int num1, num2; // Two numbers for GCD calculation
12
13         // Create a Scanner object for keyboard input.
14         Scanner keyboard = new Scanner(System.in);
15
16         // Get the first number from the user.
17         System.out.print("Enter an integer: ");
18         num1 = keyboard.nextInt();
19
20         // Get the second number from the user.
21         System.out.print("Enter another integer: ");
22         num2 = keyboard.nextInt();
23
24         // Display the GCD.
25         System.out.println("The greatest common divisor " +
26                             "of these two numbers is " +
27                             gcd(num1, num2));
28     }
29
30     /**
31      * The gcd method calculates the greatest common
32      * divisor of the arguments passed into x and y.
33      * @param x A number.
34      * @param y Another number.
35      * @returns The greatest common divisor of x and y.
36     */
37
38     public static int gcd(int x, int y)
39     {
40         if (x % y == 0)
41             return y;
42         else
43             return gcd(y, x % y);
44     }
45 }

```

Program Output with Example Input Shown in Bold

```

Enter an integer: 49 [Enter]
Enter another integer: 28 [Enter]
The greatest common divisor of these two numbers is 7

```

15.4 A Recursive Binary Search Method

CONCEPT: The recursive binary search algorithm is more elegant and easier to understand than its iterative version.

In [Chapter 7](#) you learned about the binary search algorithm and saw an iterative example written in Java. The binary search algorithm can also be implemented recursively. For example, the procedure can be expressed as:

*If array[middle] equals the search value, then the value is found.
Else if array[middle] is less than the search value, perform
 a binary search on the upper half of the array.
Else if array[middle] is greater than the search value, perform
 a binary search on the lower half of the array.*

When you compare the recursive algorithm to its iterative counterpart, it becomes evident that the recursive version is much more elegant and easier to understand. The recursive binary search algorithm is also a good example of repeatedly breaking a problem down into smaller pieces until it is solved. Here is the code for the method:

```
public static int binarySearch(int[] array, int first,
                               int last, int value)
{
    int middle;        // Mid point of search
    // Test for the base case where the
    // value is not found.
    if (first > last)
        return -1;
    // Calculate the middle position.
    middle = (first + last) / 2;
    // Search for the value.
    if (array[middle] == value)
        return middle;
    else if (array[middle] < value)
        return binarySearch(array, middle + 1,
                              last, value);
    else
        return binarySearch(array, first,
                              middle - 1, value);
}
```

The first parameter, `array`, is the array to be searched. The next parameter, `first`, holds the subscript of the first element in the search range (the portion of the array to be searched). The next parameter, `last`, holds the subscript of the last element in the search range. The last parameter, `value`, holds the value to be searched for. Like the iterative version, this method returns the subscript of the value if it is found, or `-1` if the value is not found. [Code Listing 15-9](#) demonstrates the method.

Code Listing 15-9

(RecursiveBinarySearch.java)

```
1  import java.util.Scanner;
2
3  /**
4   * This program demonstrates the recursive
5   * binary search method.
6   */
7
8  public class RecursiveBinarySearch
9  {
10     public static void main(String [] args)
11     {
12         int searchValue;    // The value to search for
13         int result;         // The search result
14         String input;       // A line of input
15         char again;         // To hold a single character
16
17         // The values in the following array are sorted
18         // in ascending order.
19         int numbers[] = {101, 142, 147, 189, 199, 207, 222,
20                        234, 289, 296, 310, 319, 388, 394,
21                        417, 429, 447, 521, 536, 600};
22
23         // Create a Scanner object for keyboard input.
24         Scanner keyboard = new Scanner(System.in);
25
26         do
27         {
28             // Get a value to search for.
29             System.out.print("Enter a value to " +
30                             "search for: ");
31             searchValue = keyboard.nextInt();
32
33             // Search for the value
34             result = binarySearch(numbers, 0,
35                                   (numbers.length - 1),
36                                   searchValue);
37
38             // Display the results.
39             if (result == -1)
40             {
41                 System.out.println(searchValue +
42                                     " was not found.");
43             }
44             else
45             {
46                 System.out.println(searchValue +
47                                     " was found at " +
48                                     "element " + result);
49             }
50
51             // Does the user want to search again?
```

```

52         System.out.print("Do you want to search again? " +
53                             "(Y or N): ");
54         // Consume the remaining newline.
55         keyboard.nextLine();
56         // Read a line of input.
57         input = keyboard.nextLine();
58
59     } while (input.charAt(0) == 'y' ||
60             input.charAt(0) == 'Y');
61 }
62
63 /**
64  The binarySearch method performs a binary search
65  on an integer array.
66  @param array The array to search.
67  @param first The first element in the search range.
68  @param last The last element in the search range.
69  @param value The value to search for.
70  @return The subscript of the value if found,
71          otherwise -1.
72 */
73
74 public static int binarySearch(int[] array, int first,
75                               int last, int value)
76 {
77     int middle;        // Mid point of search
78
79     // Test for the base case where the
80     // value is not found.
81     if (first > last)
82         return -1;
83
84     // Calculate the middle position.
85     middle = (first + last) / 2;
86
87     // Search for the value.
88     if (array[middle] == value)
89         return middle;
90     else if (array[middle] < value)
91         return binarySearch(array, middle + 1,
92                             last, value);
93     else
94         return binarySearch(array, first,
95                             middle - 1, value);
96 }
97 }

```

Program Output with Example Input Shown in Bold

```

Enter a value to search for: 289 [Enter]
289 was found at element 8
Do you want to search again? (Y or N): y [Enter]

```

Enter a value to search for: **388** [Enter]
388 was found at element 12
Do you want to search again? (Y or N): **y** [Enter]
Enter a value to search for: **101** [Enter]
101 was found at element 0
Do you want to search again? (Y or N): **y** [Enter]
Enter a value to search for: **999** [Enter]
999 was not found.
Do you want to search again? (Y or N): **n** [Enter]

15.5 The Towers of Hanoi

CONCEPT: The repetitive steps involved in solving the Towers of Hanoi game can be easily implemented in a recursive algorithm.

The Towers of Hanoi is a mathematical game that is often used in computer science textbooks to illustrate the power of recursion. The game uses three pegs and a set of discs with holes through their centers. The discs are stacked on one of the pegs as shown in [Figure 15-7](#).




Figure 15-7 The pegs and discs in the Towers of Hanoi game

[Figure 15-7 Full Alternative Text](#)

Notice that the discs are stacked on the leftmost peg, in order of size with the largest disc at the bottom. The game is based on a legend where a group of monks in a temple in Hanoi have a similar set of pegs with 64 discs. The job of the monks is to move the discs from the first peg to the third peg. The middle peg can be used as a temporary holder. Furthermore, the monks must follow these rules while moving the discs:

- Only one disk may be moved at a time.
- A disk cannot be placed on top of a smaller disc.
- All discs must be stored on a peg except while being moved.

According to the legend, when the monks have moved all of the discs from the first peg to the last peg, the world will come to an end.

To play the game, you must move all of the discs from the first peg to the third peg, following the same rules as the monks. Let's look at some example solutions to this game, for different numbers of discs. If you have only one disc, the solution to the game is simple: move the disc from peg 1 to peg 3.

If you have two discs, the solution requires three moves:

- Move disc 1 to peg 2.
- Move disc 2 to peg 3.
- Move disc 1 to peg 3.

Notice that this approach uses peg 2 as a temporary location. The complexity of the moves continues to increase as the number of discs increases. To move three discs requires the seven moves shown in [Figure 15-8](#).

The following statement describes the overall solution to the problem:


Move n discs from peg 1 to peg 3 using peg 2 as a temporary peg.

The following algorithm can be used as the basis of a recursive method that simulates the solution to the game. Notice that in this algorithm we use the variables A, B, and C to hold peg numbers.


*To move n discs from peg A to peg C, using peg B as a temporary peg:
If n > 0 then
 Move n - 1 discs from peg A to peg B, using peg C as a temporary peg.
 Move the remaining disc from peg A to peg C.
 Move n - 1 discs from peg B to peg C, using peg A as a temporary peg.
End if*

The base case for the algorithm is reached when there are no more discs to move. The following code is for a method that implements this algorithm. Note that the method does not actually move anything, but displays instructions indicating all of the disc moves to make.


```
private void moveDiscs(int num, int fromPeg, int toPeg, int tempPeg)
{
    if (num > 0)
    {
        moveDiscs(num - 1, fromPeg, tempPeg, toPeg);
        System.out.println("Move a disc from peg " + fromPeg +
                           " to peg " + toPeg);
        moveDiscs(num - 1, tempPeg, toPeg, fromPeg);
    }
}
```




Original setup.




First move: Move disc 1 to peg 3.




Second move: Move disc 2 to peg 2.




Third move: Move disc 1 to peg 2.




Fourth move: Move disc 3 to peg 3.



Fifth move: Move disc 1 to peg 1.



Sixth move: Move disc 2 to peg 3.



Seventh move: Move disc 1 to peg 3.

Figure 15-8 Steps for moving three pegs

[Figure 15-8 Full Alternative Text](#)

This method accepts arguments into the following four parameters:

`num` The number of discs to move.

`fromPeg` The peg to move the discs from.

toPeg The peg to move the discs to.

tempPeg The peg to use as a temporary peg.

If num is greater than 0, then there are discs to move. The first recursive call is as follows:

```
moveDiscs(num - 1, fromPeg, tempPeg, toPeg);
```

This statement is an instruction to move all but one disc from fromPeg to tempPeg, using toPeg as a temporary peg. The next statement is as follows:

```
System.out.println("Move a disc from peg " + fromPeg +  
                   " to peg " + toPeg);
```

This simply displays a message indicating that a disc should be moved from fromPeg to toPeg. Next, another recursive call is executed as follows:

```
moveDiscs(num - 1, tempPeg, toPeg, fromPeg);
```

This statement is an instruction to move all but one disc from tempPeg to toPeg, using fromPeg as a temporary peg. [Code Listing 15-10](#) shows the Hanoi class, which uses this method.

Code Listing 15-10 (Hanoi.java)

```
1  /**  
2      This class displays a solution to the Towers of  
3      Hanoi game.  
4  */  
5  
6  public class Hanoi  
7  {  
8      private int numDiscs;    // Number of discs  
9  
10     /**  
11         Constructor.  
12         @param n The number of discs to use.  
13     */  
14  
15     public Hanoi(int n)  
16     {  
17         // Assign the number of discs.  
18         numDiscs = n;  
19  
20         // Move the number of discs from peg 1 to peg 3  
21         // using peg 2 as a temporary storage location.  
22         moveDiscs(numDiscs, 1, 3, 2);  
23     }  
24  
25     /**
```

```

26         The moveDiscs method displays a disc move.
27         @param num The number of discs to move.
28         @param fromPeg The peg to move from.
29         @param toPeg The peg to move to.
30         @param tempPeg The temporary peg.
31     */
32
33     private void moveDiscs(int num, int fromPeg,
34                           int toPeg, int tempPeg)
35     {
36         if (num > 0)
37         {
38             moveDiscs(num - 1, fromPeg, tempPeg, toPeg);
39             System.out.println("Move a disc from peg " +
40                               fromPeg + " to peg " + toPeg);
41             moveDiscs(num - 1, tempPeg, toPeg, fromPeg);
42         }
43     }
44 }

```

The class constructor accepts an argument, which is the number of discs to use in the game. It assigns this value to the `numDiscs` field, and then calls the `moveDiscs` method in line 22. In a nutshell, this statement is an instruction to move all the discs from peg 1 to peg 3, using peg 2 as a temporary peg. The program in [Code Listing 15-11](#) demonstrates the class. It displays the instructions for moving three discs.

Code Listing 15-11 (HanoiDemo.java)

```

1  /**
2   * This class demonstrates the Hanoi class, which
3   * displays the steps necessary to solve the Towers
4   * of Hanoi game.
5   */
6
7  public class HanoiDemo
8  {
9      static public void main(String[] args)
10     {
11         Hanoi towersOfHanoi = new Hanoi(3);
12     }
13 }

```

Program Output

```

Move a disc from peg 1 to peg 3
Move a disc from peg 1 to peg 2
Move a disc from peg 3 to peg 2
Move a disc from peg 1 to peg 3
Move a disc from peg 2 to peg 1
Move a disc from peg 2 to peg 3
Move a disc from peg 1 to peg 3

```


15.6 Common Errors to Avoid

- Not coding a base case. When the base case is reached, a recursive method stops calling itself. Without a base case, the method will continue to call itself infinitely.
- Not reducing the problem with each recursive call. Unless the problem is reduced (which usually means that the value of one or more critical parameters is reduced) with each recursive call, the method will not reach the base case. If the base case is not reached, the method will call itself infinitely.
- Writing the recursive call in such a way that the base case is never reached. You might have a base case and a recursive case that reduces the problem, but if the calculations are not performed in such a way that the base case is ultimately reached, the method will call itself infinitely.

Review Questions and Exercises

Multiple Choice and True/False

1. A method is called once from a program's `main` method, and then it calls itself four times. The depth of recursion is .
 1. one
 2. four
 3. five
 4. nine
2. This is the part of a problem that can be solved without recursion.
 1. base case
 2. solvable case
 3. known case
 4. iterative case
3. This is the part of a problem that is solved with recursion.
 1. base case
 2. iterative case
 3. unknown case
 4. recursion case
4. This is when a method explicitly calls itself.
 1. explicit recursion
 2. modal recursion
 3. direct recursion
 4. indirect recursion

5. This is when method A calls method B, which calls method A.
 1. implicit recursion
 2. modal recursion
 3. direct recursion
 4. indirect recursion
6. This refers to the actions taken internally by the JVM when a method is called.
 1. overhead
 2. set up
 3. clean up
 4. synchronization
7. **True or False:** An iterative algorithm will usually run faster than an equivalent recursive algorithm.
8. **True or False:** Some problems can be solved through recursion only.
9. **True or False:** It is not necessary to have a base case in all recursive algorithms.
10. **True or False:** In the base case, a recursive method calls itself with a smaller version of the original problem.

Find the Error

1. Find the error in the following program:

```
public class FindTheError
{
    public static void main(String[] args)
    {
        myMethod(0);
    }
    public static void myMethod(int num)
    {
        System.out.print(num + " ");
        myMethod(num + 1);
    }
}
```


Algorithm Workbench

1. Write a method that accepts a `String` as an argument. The method should use recursion to display each individual character in the `String`.
2. Modify the method you wrote in Algorithm Workbench 1 so it displays the `String` backwards.
3. What will the following program display?

```
public class Checkpoint
{
    public static void main(String[] args)
    {
        int num = 0;
        showMe(num);
    }
    public static void showMe(int arg)
    {
        if (arg < 10)
            showMe(arg + 1);
        else
            System.out.println(arg);
    }
}
```

4. What will the following program display?

```
public class Checkpoint
{
    public static void main(String[] args)
    {
        int num = 0;
        showMe(num);
    }
    public static void showMe(int arg)
    {
        System.out.println(arg);
        if (arg < 10)
            showMe(arg + 1);
    }
}
```

5. What will the following program display?

```
public class ReviewQuestion5
{
    public static void main(String[] args)
    {
        int x = 10;
        System.out.println(myMethod(x));
    }
    public static int myMethod(int num)
    {
```

```

        if (num <= 0)
            return 0;
        else
            return myMethod(num - 1) + num;
    }
}

```

6. Convert the following iterative method to one that uses recursion:

```

public static void sign(int n)
{
    while (n > 0)
    {
        System.out.println("No Parking");
        n--;
    }
}

```

7. Write an iterative version (using a loop instead of recursion) of the `factorial` method shown in this chapter.

Short Answer

1. What is the difference between an iterative algorithm and a recursive algorithm?
2. What is a recursive algorithm's base case? What is the recursive case?
3. What is the base case of each of the recursive methods listed in Algorithm Workbench 3, 4, and 5?
4. What type of recursive method do you think would be more difficult to debug: one that uses direct recursion or one that uses indirect recursion? Why?
5. Which repetition approach is less efficient: a loop or a recursive method? Why?
6. When recursion is used to solve a problem, why must the recursive method call itself to solve a smaller version of the original problem?
7. How is a problem usually reduced with a recursive method?

Programming Challenges

MyProgrammingLab™ Visit www.myprogramminglab.com to complete many of these Programming Challenges online and get instant feedback.

1. Recursive Multiplication

Write a recursive function that accepts two arguments into the parameters `x` and `y`. The function should return the value of `x` times `y`. Remember, multiplication can be performed as repeated addition as follows:

$$7 * 4 = 4 + 4 + 4 + 4 + 4 + 4 + 4$$

2. isMember Method

Write a recursive `boolean` method named `isMember`. The method should search an array for a specified value, and return `true` if the value is found in the array, or `false` if the value is not found in the array. Demonstrate the method in a program.

3. String Reverser

Write a recursive method that accepts a string as its argument and prints the string in reverse order. Demonstrate the method in a program.

4. maxElement Method

Write a method named `maxElement`, which returns the largest value in an array that is passed as an argument. The method should use recursion to find the largest element. Demonstrate the method in a program.

5. Palindrome Detector

A palindrome is any word, phrase, or sentence that reads the same forward and backward. Here are some well-known palindromes:

- Able was I, ere I saw Elba

- A man, a plan, a canal, Panama
- Desserts, I stressed
- Kayak

Write a `boolean` method that uses recursion to determine whether a `String` argument is a palindrome. The method should return `true` if the argument reads the same forward and backward. Demonstrate the method in a program.

6. Character Counter

Write a method that uses recursion to count the number of times a specific character occurs in an array of characters. Demonstrate the method in a program.

7. Recursive Power Method

Write a method that uses recursion to raise a number to a power. The method should accept two arguments: the number to be raised and the exponent. Assume that the exponent is a nonnegative integer. Demonstrate the method in a program.



The Recursive Power Problem

8. Sum of Numbers

Write a method that accepts an integer argument and returns the sum of all the integers from 1 up to the number passed as an argument. For example, if 50 is passed as an argument, the method will return the sum of 1, 2, 3, 4, . . . 50. Use recursion to calculate the sum. Demonstrate the method in a program.

9. Ackermann's Function

Ackermann's function is a recursive mathematical algorithm that can be used to test how well a computer performs recursion. Write a method `ackermann(m, n)`, which solves Ackermann's function. Use the following logic in your method:

- If $m = 0$ then return $n+1$

- If $n = 0$ then return `ackermann(m - 1, 1)`
- Otherwise, return `ackermann(m - 1, ackermann(m, n - 1))`

Test your method in a program that displays the return values of the following method calls:

- `ackermann(0, 0)` `ackermann(0, 1)` `ackermann(1, 1)` `ackermann(1, 2)`
- `ackermann(1, 3)` `ackermann(2, 2)` `ackermann(3, 2)`

10. Recursive Population Class

In Programming Challenge 9 of [Chapter 4](#) you wrote a population class that predicts the size of a population of organisms after a number of days. Modify the class so it uses a recursive method instead of a loop to calculate the number of organisms.

Chapter 16 Sorting, Searching, and Algorithm Analysis

TOPICS

1. [16.1 Introduction to Sorting Algorithms](#)
2. [16.2 Introduction to Search Algorithms](#)
3. [16.3 Analysis of Algorithms](#)
4. [16.4 Common Errors to Avoid](#)

16.1 Introduction to Sorting Algorithms

CONCEPT: Sorting algorithms are used to arrange data into some order.

Often, the data in an array must be sorted in some order. Customer lists, for instance, are commonly sorted in alphabetical order. Student grades might be sorted from highest to lowest. Product codes could be sorted so all the products of the same color are stored together. To sort the data in an array, the programmer must use an appropriate *sorting algorithm*. A sorting algorithm is a technique for scanning through an array and rearranging its contents in some specific order.

The data in an array can be sorted in either ascending or descending order. If an array is sorted in *ascending order*, it means the values in the array are stored from lowest to highest. If the values are sorted in *descending order*, they are stored from highest to lowest. This section will introduce four sorting algorithms that can be used to sort the data in an array: the *bubble sort*, the *selection sort*, the *insertion sort*, and the *Quicksort*.

The Bubble Sort



The Bubble Sort Algorithm

The bubble sort is an easy way to arrange data in ascending or descending order. In this section we will see how the bubble sort can be used to sort an array in ascending order. This particular algorithm is called the *bubble sort* because it makes several passes through the elements of the array and the larger values “bubble” toward the end of the array with each pass.

Suppose we have the array shown in [Figure 16-1](#). Let’s see how the bubble sort can be used in arranging the array’s elements in ascending order.

7	2	3	8	9	1
Element 0	Element 1	Element 2	Element 3	Element 4	Element 5

Figure 16-1 An array

The bubble sort starts by comparing the first two elements in the array. If element 0 is greater than element 1, they are exchanged. After the exchange, the array would appear as shown in [Figure 16-2](#).




Figure 16-2 Elements 0 and 1 are exchanged

This step is repeated with elements 1 and 2. If element 1 is greater than element 2, they are exchanged. The array would then appear as shown in [Figure 16-3](#).




Figure 16-3 Elements 1 and 2 are exchanged

Next, elements 2 and 3 are compared. In this array, these elements are already in the proper order (element 2 is less than element 3), so no exchange takes place. As the cycle continues, elements 3 and 4 are compared. Once again, no exchange is necessary because they are already in the proper order.

When elements 4 and 5 are compared, however, an exchange must take place because element 4 is greater than element 5. The array now appears as shown in [Figure 16-4](#).




Figure 16-4 Elements 4 and 5 are exchanged

At this point, the entire array has been scanned, and the largest value, 9, is in the correct position. There are other elements, however, that are not yet in their final positions. So, we make another pass through the array, comparing each element with its neighbor. In the next pass we can stop comparing after we reach the next-to-last element because the last element is already at the correct position.

The second pass starts by comparing elements 0 and 1. Because those two are in the proper order, no exchange takes place. Elements 1 and 2 are compared next, but once again, no exchange takes place. This continues until elements 3 and 4 are compared. Because element 3 is greater than element 4, they are exchanged. Element 4 is the last element that we will compare during this pass, so this pass stops. The array now appears as shown in [Figure 16-5](#).




Figure 16-5 Elements 3 and 4 are exchanged

At the end of the second pass, the last two elements in the array are in their correct positions. The third pass starts now, comparing each element with its neighbor. The third pass will not involve the last two elements, however, because they are already in their correct positions. When the third pass is finished, the last three elements will also have been correctly placed, as shown in [Figure 16-6](#).




Figure 16-6 The array after the third pass

Each time the algorithm makes a pass through the array, the portion of the array that is scanned is decreased in size by one element, and the largest value in the scanned portion of the array is moved to

its final position. When all of the passes have been made, the array will appear as shown in [Figure 16-7](#).

1	2	3	7	8	9
Element 0	Element 1	Element 2	Element 3	Element 4	Element 5

Figure 16-7 The array with all elements sorted

The following pseudocode shows the logic of the bubble sort.

```
For lastPosition = last subscript subscript in the array, decremented down to 0
  For index = 0 up through lastPosition - 1
    If array[index] > array[index + 1]
      Swap array[index] with array[index + 1]
    End If
  End For
End For
```

[Code Listing 16-1](#) shows the `IntBubbleSorter` class, which has a public static method, `bubbleSort`. You pass an `int` array to the method and it uses the bubble sort algorithm to sort the array's contents in ascending order.

Code Listing 16-1 (`IntBubbleSorter.java`)

```
1 /**
2  The IntBubbleSorter class provides a public static
3  method for performing a bubble sort on an int array.
4  */
5
6 public class IntBubbleSorter
7 {
8
9     /**
10     The bubbleSort method uses the bubble sort algorithm
11     to sort an int array.
12     @param array The array to sort.
13     */
14
15     public static void bubbleSort(int[] array)
16     {
17         int lastPos;    // Position of the last element to compare
18         int index;      // Index of an element to compare
19         int temp;       // Used to swap to elements
20
21         // The outer loop positions lastPos at the last element
```

```

22 // to compare during each pass through the array. Initially
23 // lastPos is the index of the last element in the array.
24 // During each iteration, it is decreased by one.
25 for (lastPos = array.length - 1; lastPos >= 0; lastPos--)
26 {
27     // The inner loop steps through the array, comparing
28     // each element with its neighbor. All of the elements
29     // from index 0 through lastPos are involved in the
30     // comparison. If two elements are out of order, they
31     // are swapped.
32     for (index = 0; index <= lastPos - 1; index++)
33     {
34         // Compare an element with its neighbor.
35         if (array[index] > array[index + 1])
36         {
37             // Swap the two elements.
38             temp = array[index];
39             array[index] = array[index + 1];
40             array[index + 1] = temp;
41         }
42     }
43 }
44 }
45 }

```

Lines 17, 18, and 19 declare three variables: `lastPos`, `index`, and `temp`. Here are the purposes of the variables:

- The `lastPos` variable will hold the index of the last element that is to be compared to its neighbor during a pass through the array.
- The `index` variable is used as an index into the array during each pass.
- The `temp` variable is used to hold the value of an element temporarily during a swap.

The method uses two `for` loops, one nested inside another. The outer loop begins in line 25 as follows:

```
for (lastPos = array.length - 1; lastPos >= 0; lastPos--)
```

This loop will iterate once for each element in the array. It causes the `lastPos` variable to take on all of the array's subscripts, from the highest subscript down to 0. After each iteration, `lastPos` is decremented by one.

The second loop, which is nested inside the first loop, begins in line 32 as follows:

```
for (index = 0; index <= lastPos - 1; index++)
```

This loop iterates once for each element in the portion of the array that is still unsorted. It starts `index` at 0 and increments it up through `lastPos - 1`. During each iteration, the comparison in line 35 is performed:

```
if (array[index] > array[index + 1])
```

This `if` statement compares the element at `array[index]` with its neighbor `array[index + 1]`. If the element's neighbor is less, then the two are swapped by the code in lines 38 through 40.

The program in [Code Listing 16-2](#) tests the method.

Code Listing 16-2 (BubbleSortTest.java)

```
1  /**
2   * This program tests the bubbleSort method in the
3   * IntBubbleSorter class.
4   */
5
6  public class BubbleSortTest
7  {
8      public static void main(String[] args)
9      {
10         // Create an int array with test values.
11         int[] values = { 5, 1, 3, 6, 4, 2 };
12
13         // Display the array's contents.
14         System.out.println("Original order: ");
15         for (int element : values)
16             System.out.print(element + " ");
17
18         // Sort the array.
19         IntBubbleSorter.bubbleSort(values);
20
21         // Display the array's contents.
22         System.out.println("\nSorted order: ");
23         for (int element : values)
24             System.out.print(element + " ");
25
26         System.out.println();
27     }
28 }
```

Program Output

```
Original order:
5 1 3 6 4 2
Sorted order:
1 2 3 4 5 6
```

Using the Bubble Sort to Sort Objects

The previous example demonstrated how the bubble sort algorithm can be used to sort an array of `int`

values. But, suppose we have an array of objects that we wish to sort. How could we modify the `bubbleSort` method shown previously so it can work with objects?

To sort an array of objects, we must have a consistent way to compare the objects. In the `IntBubbleSorter` class, line 35 of the `bubbleSort` method compares two elements of the array and determines whether one is greater than the other. Because that method is comparing primitive `int` values, it simply compares the two elements with the greater than (`>`) operator. This approach will not work with object references, however. Reference variables hold only the addresses of objects, and you cannot use operators like greater than (`>`) or less than (`<`) to compare them. (Even if you could use these operators with object references, it would not make sense to do so. When you compare two objects, you usually want to compare the data contained within those objects.)

The Java API has an interface named `Comparable`, which helps in a situation like this. The `Comparable` interface, which is in the `java.lang` package, specifies a method named `compareTo`, which is used to compare the calling object with another object that is passed as an argument. The method returns an `int` value that specifies whether the calling object is greater than (`>`), less than (`<`), or equal to (`=`) the object that was passed as an argument.

Many of the classes in the Java API implement this interface, including `String` and the wrapper classes. You might recall from our discussion of the `String` class in [Chapter 3](#) how the `compareTo` method works:

- If the calling object is less than the object passed as an argument, the method returns a negative number.
- If the calling object is equal to the object passed as an argument, the method returns 0.
- If the calling object is greater than the object passed as an argument, the method returns a positive number.

If we want to write a `bubbleSort` method that can accept an array of objects, a good solution is to require that the objects in the array implement the `Comparable` interface. That way we can use the `compareTo` method to compare the elements. [Code Listing 16-3](#) shows the `ObjectBubbleSorter` class, which has a `bubbleSort` method that sorts an array of `Comparable` objects.

Code Listing 16-3 (`ObjectBubbleSorter.java`)

```
1 /**
2  The ObjectBubbleSorter class provides a public static
3  method for performing a bubble sort on an array of
4  objects implementing the Comparable interface.
5  */
6
7 public class ObjectBubbleSorter
8 {
9
10     /**
```

```

11     The bubbleSort method uses the bubble sort algorithm
12     to sort an array of objects implementing Comparable
13     @param array The array to sort.
14     */
15
16     public static void bubbleSort(Comparable[] array)
17     {
18         int lastPos;        // Marks the last position to compare
19         int index;          // Index of an element to compare
20         Comparable temp;    // Used to swap to elements
21
22         // The outer loop positions lastPos at the last element
23         // to compare during each pass through the array. Initially
24         // lastPos is the index of the last element in the array.
25         // During each iteration, it is decreased by one.
26         for (lastPos = array.length - 1; lastPos >= 0; lastPos--)
27         {
28             // The inner loop steps through the array, comparing
29             // each element with its neighbor. All of the elements
30             // from index 0 through lastPos are involved in the
31             // comparison. If two elements are out of order, they
32             // are swapped.
33             for (index = 0; index <= lastPos - 1; index++)
34             {
35                 // Compare an element with its neighbor.
36                 if (array[index].compareTo(array[index + 1]) > 0)
37                 {
38                     // Swap the two elements.
39                     temp = array[index];
40                     array[index] = array[index + 1];
41                     array[index + 1] = temp;
42                 }
43             }
44         }
45     }
46 }

```

NOTE:

If you compile the class shown in [Code Listing 16-3](#) you will likely see a warning such as:

```

ObjectBubbleSorter.java uses unchecked or unsafe operations.
Note: Recompile with -Xlint:unchecked for details.

```

For now, ignore this warning. In [Chapter 17](#), when we discuss generic types, you will learn how to eliminate these warnings.

Notice in line 16 that the `bubbleSort` method's parameter is declared as `Comparable[] array`. When we call this method, the argument that we pass must be an array of objects that implement `Comparable`. The comparison that is made in line 36 reads as follows:

```

if (array[index].compareTo(array[index + 1]) > 0)

```

This statement uses the `array[index]` object's `compareTo` method to determine whether it is greater than the object at `array[index + 1]`. If it is, the two elements are swapped. (Also notice that the `temp` variable's data type has been changed. It is now a `Comparable` reference variable.)

Because the `String` class implements the `Comparable` interface, an easy way to test the method is to pass an array of `String` objects to it, as shown in [Code Listing 16-4](#).

Code Listing 16-4

(ObjectBubbleSortTest.java)

```
1  /**
2   * This program tests the bubbleSort method in the
3   * ObjectBubbleSorter class.
4   */
5
6  public class ObjectBubbleSortTest
7  {
8      public static void main(String[] args)
9      {
10         // Create a String array with test values.
11         // Note that String implements Comparable.
12         String[] values = { "David", "Abe", "Katherine",
13                             "Beth", "Jeff", "Daisy" };
14
15         // Display the array's contents.
16         System.out.println("Original order:");
17         for (String element : values)
18             System.out.print(element + " ");
19
20         // Sort the array.
21         ObjectBubbleSorter.bubbleSort(values);
22
23         // Display the array's contents.
24         System.out.println("\nSorted order:");
25         for (String element : values)
26             System.out.print(element + " ");
27
28         System.out.println();
29     }
30 }
```

Program Output

```
Original order:
David Abe Katherine Beth Jeff Daisy
Sorted order:
Abe Beth Daisy David Jeff Katherine
```

When you are designing a class of your own, and you want to be able to sort objects of that class

using the `bubbleSort` method shown in the `ObjectBubbleSorter` class ([Code Listing 16-3](#)), your class must implement the `Comparable` interface, and provide a `compareTo` method. In [Chapter 17](#), Generics, we will take a closer look at the `Comparable` interface and how to implement it.

The Selection Sort

The *selection sort* works like this: The smallest value in the array is located and moved to position 0. Then, the next smallest value is located and moved to position 1. This process continues until all of the elements have been placed in their proper order. Let's see how the selection sort works when arranging the elements of the array in [Figure 16-8](#).




Figure 16-8 Values in an array

The selection sort scans the array, starting at position 0, and locates the element with the smallest value. Then, this smallest element is swapped with the element at position 0. In this example, the 1 stored at position 5 is swapped with the 5 stored at position 0. After the exchange, the array appears as shown in [Figure 16-9](#).




Figure 16-9 Values in the array after the first swap

Then, the algorithm repeats the process, but because the element at position 0 is already the smallest value in the array, it can be left out of the procedure. This time, the algorithm begins the scan at position 1. In this example, the element at 2 is exchanged with the element at position 1. Then, the array appears as shown in [Figure 16-10](#).

These two elements were swapped.




Figure 16-10 Values in the array after the second swap

Once again the process is repeated, but this time the scan begins at position 2. The algorithm will find that position 5 contains the next smallest value. This element is swapped with the element at position 2, causing the array to appear as shown in [Figure 16-11](#).

These two elements were swapped.




Figure 16-11 Values in the array after the third swap

Next, the scanning begins at position 3. Its value is swapped with that at position 5, causing the array to appear as shown in [Figure 16-12](#).

These two elements were swapped.




Figure 16-12 Values in the array after the fourth swap

At this point, there are only two elements left in the unsorted portion of the array. The algorithm finds that the value at 5 is smaller than that at position 4, so the two are swapped. This puts the array in its final arrangement, as shown in [Figure 16-13](#).




Figure 16-13 Values in the array after the fifth swap

Here is the selection sort algorithm in pseudocode:

```
For startScan equals each subscript in array from 0 through the next-to-last sub.  
  Set index variable to startScan  
  Set minIndex variable to startScan  
  Set minValue variable to array[startScan]  
  For index is each subscript in array from (startScan + 1) through the last s  
    If array[index] is less than minValue  
      Set minValue to array[index]  
      Set minIndex to index  
    End If  
  Increment index  
End For  
Set array[minIndex] to array[startScan]  
Set array[startScan] to minValue  
End For
```

[Code Listing 16-5](#) shows the `IntSelectionSorter` class, which has a public static method, `selectionSort`. You pass an `int` array to the method and it uses the selection sort algorithm to sort the array's contents in ascending order. The program in [Code Listing 16-6](#) tests the method.

Code Listing 16-5 (`IntSelectionSorter.java`)

```

1  /**
2   The IntSelectionSorter class provides a public static
3   method for performing a selection sort on an int array.
4  */
5
6  public class IntSelectionSorter
7  {
8
9   /**
10    The selectionSort method performs a selection sort on an
11    int array. The array is sorted in ascending order.
12    @param array The array to sort.
13   */
14
15   public static void selectionSort(int[] array)
16   {
17       int startScan; // Starting position of the scan
18       int index;      // To hold a subscript value
19       int minIndex;   // Element with smallest value in the scan
20       int minValue;   // The smallest value found in the scan
21
22       // The outer loop iterates once for each element in the
23       // array. The startScan variable marks the position where
24       // the scan should begin.
25       for (startScan = 0; startScan < (array.length-1); startScan++)
26       {
27           // Assume the first element in the scannable area
28           // is the smallest value.
29           minIndex = startScan;
30           minValue = array[startScan];
31
32           // Scan the array, starting at the 2nd element in
33           // the scannable area. We are looking for the smallest
34           // value in the scannable area.
35           for(index = startScan + 1; index < array.length; index++)
36           {
37               if (array[index] < minValue)
38               {
39                   minValue = array[index];
40                   minIndex = index;
41               }
42           }
43
44           // Swap the element with the smallest value
45           // with the first element in the scannable area.
46           array[minIndex] = array[startScan];
47           array[startScan] = minValue;
48       }
49   }
50 }

```

Code Listing 16-6 (SelectionSortTest.java)

```

1  /**
2   This program tests the selectionSort method in the

```

```

3     IntSelectionSorter class.
4  */
5
6  public class SelectionSortTest
7  {
8      public static void main(String[] args)
9      {
10         // Create an int array with test values.
11         int[] values = { 5, 1, 3, 6, 4, 2 };
12
13         // Display the array's contents.
14         System.out.println("Original order:");
15         for (int element : values)
16             System.out.print(element + " ");
17
18         // Sort the array.
19         IntSelectionSorter.selectionSort(values);
20
21         // Display the array's contents.
22         System.out.println("\nSorted order:");
23         for (int element : values)
24             System.out.print(element + " ");
25
26         System.out.println();
27     }
28 }

```

Program Output

```

Original order:
5 1 3 6 4 2
Sorted order:
1 2 3 4 5 6

```

Using the Selection Sort to Sort Objects

Previously, we demonstrated how the `bubbleSort` method could be written to accept an array of `Comparable` objects, and then use the `compareTo` method in the bubble sort algorithm to compare the objects. We can take the same approach with the selection sort, as shown in [Code Listing 16-7](#). This class has a public static `selectionSort` method that accepts an array of `Comparable` objects.

Code Listing 16-

7 (ObjectSelectionSorter.java)

```

1  /**
2   * The ObjectSelectionSorter class provides a public static
3   * method for performing a selection sort on an array of

```

```

4  objects that implement the Comparable interface.
5  */
6
7  public class ObjectSelectionSorter
8  {
9
10     /**
11      The selectionSort method performs a selection sort on an
12      array of objects that implement the Comparable interface.
13      @param array The array to sort.
14     */
15
16     public static void selectionSort(Comparable[] array)
17     {
18         int startScan;           // Starting position of the scan
19         int index;               // To hold a subscript value
20         int minIndex;           // Element with smallest value in the scan
21         Comparable minValue;     // The smallest value found in the scan
22
23         // The outer loop iterates once for each element in the
24         // array. The startScan variable marks the position where
25         // the scan should begin.
26         for (startScan = 0; startScan < (array.length-1); startScan++)
27         {
28             // Assume the first element in the scannable area
29             // is the smallest value.
30             minIndex = startScan;
31             minValue = array[startScan];
32
33             // Scan the array, starting at the 2nd element in
34             // the scannable area. We are looking for the smallest
35             // value in the scannable area.
36             for(index = startScan + 1; index < array.length; index++)
37             {
38                 if (array[index].compareTo(minValue) < 0)
39                 {
40                     minValue = array[index];
41                     minIndex = index;
42                 }
43             }
44
45             // Swap the element with the smallest value
46             // with the first element in the scannable area.
47             array[minIndex] = array[startScan];
48             array[startScan] = minValue;
49         }
50     }
51 }

```

The program in [Code Listing 16-8](#) tests the `selectionSort` method with an array of `String` objects.

Code Listing 16-8

(ObjectSelectionSortTest.java)

```

1  /**
2   This program tests the selectionSort method in the
3   ObjectSelectionSorter class.
4  */
5
6  public class ObjectSelectionSortTest
7  {
8      public static void main(String[] args)
9      {
10         // Create a String array with test values.
11         // Note that String implements Comparable.
12         String[] values = { "David", "Abe", "Katherine",
13                             "Beth", "Jeff", "Daisy" };
14
15         // Display the array's contents.
16         System.out.println("Original order:");
17         for (String element : values)
18             System.out.print(element + " ");
19
20         // Sort the array.
21         ObjectSelectionSorter.selectionSort(values);
22
23         // Display the array's contents.
24         System.out.println("\nSorted order:");
25         for (String element : values)
26             System.out.print(element + " ");
27
28         System.out.println();
29     }
30 }

```

Program Output

```

Original order:
David Abe Katherine Beth Jeff Daisy
Sorted order:
Abe Beth Daisy David Jeff Katherine

```

The Insertion Sort

In the insertion sort, we begin by looking at the first two elements of the array. We compare these elements and, if necessary, we swap them so they are in the proper order. This becomes a sorted portion of the array.

Then, our objective is to incorporate the third element of the array into the part of the array already sorted. We do this by inserting it into proper position, relative to the first two elements. If we need to shift either of the first two elements to accommodate the third element, we do so. Once we have inserted the third element into the correct position (relative to the first two elements), the initial portion of the array consisting of the first three elements will be sorted.

This process continues with the fourth and subsequent elements, until all of the elements have been inserted into their proper positions. Let's look at an example. Suppose we start with the `int` array shown in [Figure 16-14](#). As shown in the figure, the values in the first and second elements are out of order, so they will be swapped.



Figure 16-14 An unsorted array

After the swap takes place, the first and second elements will be in sorted order. The next step is to move the value of the third element so it is in the correct position, relative to the first two elements. As shown in [Figure 16-15](#), the value in the third element must be positioned between the values in the first and second elements.




Figure 16-15 The third element must be moved

After the value in the third element is moved to its new position, the first three elements will be in sorted order. The next step is to move the value of the fourth element so it is in the correct position, relative to the first three elements. As shown in [Figure 16-16](#), the value in the fourth element must be positioned between the values in the second and third elements.




Figure 16-16 The fourth element must be moved

After the value in the fourth element is moved to its new position, the first four elements will be in sorted order. The next step is to move the value of the fifth element so it is in the correct position, relative to the first four elements. As shown in [Figure 16-17](#), the value in the fifth element must be positioned between the values in the first and second elements.




Figure 16-17 The fifth element must be moved

After the value in the fifth element is moved to its new position, the first five elements will be in sorted order. The next step is to move the value of the sixth element so it is in the correct position, relative to the first five elements. As shown in [Figure 16-18](#), the value in the sixth element must be moved to the beginning of the array.




Figure 16-18 The sixth element must be moved

The sixth element is the last element in the array. Once it is moved to its correct position, the entire array is sorted. This is shown in [Figure 16-19](#).




Figure 16-19 All of the elements are in the correct position

Here is the insertion sort algorithm in pseudocode:

```

For index equals each subscript in array from 1 through the last subscript
  Save the value in index in scan
  Save the value in array[index] in unsortedValue
  While scan is greater than 0 and array[scan - 1] is greater than unsortedValue
    Move array[scan - 1] to array[scan]
    Subtract 1 from scan
  End While
  Store the value in unsortedValue in array[scan]
End For

```

[Code Listing 16-9](#) shows the `IntInsertionSorter` class, which has a public static method, `insertionSort`. You pass an `int` array to the method and it uses the insertion sort algorithm to sort

the array's contents in ascending order. The program in [Code Listing 16-10](#) tests the method.

Code Listing 16-9 (IntInsertionSorter.java)

```
1 /**
2  The IntInsertionSorter class provides a public static
3  method for performing an insertion sort on an int array.
4  */
5
6 public class IntInsertionSorter
7 {
8
9     /**
10     The insertionSort method performs an insertion sort on
11     an int array. The array is sorted in ascending order.
12     @param array The array to sort.
13     */
14
15     public static void insertionSort(int[] array)
16     {
17         int unsortedValue; // The first unsorted value
18         int scan;          // Used to scan the array
19
20         // The outer loop steps the index variable through each
21         // subscript in the array, starting at 1. The portion of the
22         // array consisting of element 0 by itself is already sorted.
23         for (int index = 1; index < array.length; index++)
24         {
25             // The first element outside the sorted portion is
26             // array[index]. Store the value of this element
27             // in unsortedValue.
28             unsortedValue = array[index];
29
30             // Start scan at the subscript of the first element
31             // in the still unsorted part.
32             scan = index;
33
34             // Move the first element in the still unsorted part
35             // into its proper position within the sorted part.
36             while (scan > 0 && array[scan - 1] > unsortedValue)
37             {
38                 array[scan] = array[scan-1];
39                 scan--;
40             }
41
42             // Insert the unsorted value in its proper position
43             // within the sorted part.
44             array[scan] = unsortedValue;
45         }
46     }
47 }
```

Code Listing 16-10 (InsertionSortTest.java)

```
1 /**
2  This program tests the insertionSort method in the
3  IntInsertionSorter class.
4  */
5
6 public class InsertionSortTest
7 {
8     public static void main(String[] args)
9     {
10         // Create an int array with test values.
11         int[] values = { 5, 1, 3, 6, 4, 2 };
12
13         // Display the array's contents.
14         System.out.println("Original order:");
15         for (int element : values)
16             System.out.print(element + " ");
17
18         // Sort the array.
19         IntInsertionSorter.insertionSort(values);
20
21         // Display the array's contents.
22         System.out.println("\nSorted order:");
23         for (int element : values)
24             System.out.print(element + " ");
25
26         System.out.println();
27     }
28 }
```

Program Output

```
Original order:
5 1 3 6 4 2
Sorted order:
1 2 3 4 5 6
```

Using the Insertion Sort to Sort Objects

We have demonstrated how the `bubbleSort` and `selectionSort` methods can be written to accept an array of `Comparable` objects, and then use the `compareTo` method to compare the objects. A similar approach can be taken with the `insertionSort` method. This will be left as a programming exercise for you at the end of the chapter.

The Quicksort Algorithm

The Quicksort algorithm is a popular sorting routine developed in 1960 by C. A. R. Hoare. It sorts an array by dividing it into two sublists. Between the sublists is a selected value known as the *pivot*, as shown in [Figure 16-20](#).




Figure 16-20 Sublists and pivot

Notice in the figure that sublist 1 is positioned to the left of (before) the pivot, and sublist 2 is positioned to the right of (after) the pivot. Once a pivot value has been selected, the algorithm reorders the other values in the array until all the elements in the left sublist are less than the pivot, and all the elements in the right sublist are greater than or equal to the pivot. This does not mean that the elements in the sublists are sorted in order, just that the smaller values in the array are in the left sublist and the larger values are in the right sublist.

Once this is done, the algorithm recursively repeats the procedure on sublist 1, and then on sublist 2. During this process, the sublists are themselves divided into sublists with a pivot. All the elements less than the pivot are moved to the left sublist, and all the elements greater than the pivot are moved to the right sublist. The recursion stops when there is at most one element in a sublist. When all of the recursive calls have completed, the array is completely sorted.

To demonstrate the Quicksort algorithm we have provided a class, `IntQuickSorter.java`. This class has the following public static method:

```
14     public static void quickSort(int array[])
15     {
16         doQuickSort(array, 0, array.length - 1);
17     }
```

This method accepts an `int` array as its argument. Then it calls a private static method, `doQuickSort`. The reason we have this extra method call is that the Quicksort algorithm requires the starting subscript and the ending subscript of the area of the array that is to be sorted. When you first invoke the Quicksort algorithm, you provide the starting and ending subscripts of the entire array. As a convenience to anyone using the `IntQuickSorter` class, we have provided the `quickSort` method, which simply accepts an array. Then it calls the `doQuickSort` method, providing the starting and ending subscripts as arguments.

The pseudocode for the `doQuickSort` method is shown here:

```
doQuickSort:
    If Starting Index < Ending Index
```

```

    Partition the List around a Pivot
    doQuickSort Sublist 1
    doQuickSort Sublist 2
End If
End doQuickSort

```

Here is the Java code for the `doQuickSort` method:

```

27 private static void doQuickSort(int array[], int start, int end)
28 {
29     int pivotPoint;
30
31     if (start < end)
32     {
33         // Get the pivot point.
34         pivotPoint = partition(array, start, end);
35
36         // Sort the first sublist.
37         doQuickSort(array, start, pivotPoint - 1);
38
39         // Sort the second sublist.
40         doQuickSort(array, pivotPoint + 1, end);
41     }
42 }

```

The first argument is the array holding the list that is to be sorted. The second and third arguments are the starting and ending subscripts of the list.

The subscript of the pivot element is returned by the `partition` method in line 34. The `partition` method not only determines which element will be the pivot, but also rearranges the other values in the list. The method selects the element in the middle of the array as the pivot, and then scans the remainder of the array searching for values less than the pivot.

The code for the `partition` method is shown here:

```

56 private static int partition(int array[], int start, int end)
57 {
58     int pivotValue;           // To hold the pivot value
59     int endOfLeftList;        // Last element in the left sublist.
60     int mid;                  // To hold the mid-point subscript
61
62     // Find the subscript of the middle element.
63     // This will be our pivot value.
64     mid = (start + end) / 2;
65
66     // Swap the middle element with the first element.
67     // This moves the pivot value to the start of
68     // the list.
69     swap(array, start, mid);
70
71     // Save the pivot value for comparisons.
72     pivotValue = array[start];
73
74     // For now, the end of the left sublist is

```

```

75         // the first element.
76         endOfLeftList = start;
77
78         // Scan the entire list and move any values that
79         // are less than the pivot value to the left
80         // sublist.
81         for (int scan = start + 1; scan <= end; scan++)
82         {
83             if (array[scan] < pivotValue)
84             {
85                 endOfLeftList++;
86                 swap(array, endOfLeftList, scan);
87             }
88         }
89
90         // Move the pivot value to end of the
91         // left sublist.
92         swap(array, start, endOfLeftList);
93
94         // Return the subscript of the pivot value.
95         return endOfLeftList;
96     }

```

The `partition` method does not initially sort the values into their final order. Its job is only to move the values that are less than the pivot to the pivot's left, and move the values that are greater than the pivot to the pivot's right. As long as that condition is met, they may appear in any order. The ultimate sorting order of the entire array is achieved cumulatively, through the recursive calls to the `doQuickSort` method.

There are many different ways of partitioning the array. As previously stated, the technique shown in this `partition` method selects the middle value as the pivot. That value is then moved to the beginning of the array (by swapping it with the value stored there). This simplifies the next step, which is to scan the array.

A `for` loop scans the remainder of the array, and when an element is found whose value is less than the pivot, that value is moved to a location left of the pivot point. A third method, `swap`, is used to swap the values found in any two elements of the array. The method is shown here:

```

106     private static void swap(int[] array, int a, int b)
107     {
108         int temp;
109
110         temp = array[a];
111         array[a] = array[b];
112         array[b] = temp;
113     }

```

The program shown in [Code Listing 16-11](#) uses these methods to test the `Quicksort` algorithm.

Code Listing 16-11 (`QuickSortTest.java`)

```

1  /**
2   This program tests the quickSort method in the
3   IntQuickSorter class.
4  */
5
6  public class QuickSortTest
7  {
8      public static void main(String[] args)
9      {
10         // Create an int array with test values.
11         int[] values = { 5, 1, 3, 6, 4, 2 };
12
13         // Display the array's contents.
14         System.out.println("Original order:");
15         for (int element : values)
16             System.out.print(element + " ");
17
18         // Sort the array.
19         IntQuickSorter.quickSort(values);
20
21         // Display the array's contents.
22         System.out.println("\nSorted order:");
23         for (int element : values)
24             System.out.print(element + " ");
25
26         System.out.println();
27     }
28 }

```

Program Output

```

Original order:
5 1 3 6 4 2
Sorted order:
1 2 3 4 5 6

```

Using Quicksort to Sort Objects

As with the other sorting algorithms shown in this chapter, the Quicksort methods we have demonstrated can be modified to accept an array of `Comparable` objects, and then we can use the `compareTo` method to make comparisons. Modifying these Quicksort methods to accept an array of `Comparable` objects will be left as a programming exercise for you at the end of the chapter.

Checkpoint

1. 16.1 Which of the sorting algorithms that we discussed makes several passes through an array and causes the larger values to move gradually toward the end of the array with each pass?

2. 16.2 One of the sorting algorithms that we discussed works like this: It begins by putting the initial portion of the array consisting of the first two elements in sorted order. Then the third element is moved to its correct position, relative to the first two elements. At that point the first three elements are in sorted order. This process continues with the fourth and subsequent elements until the entire array is sorted. Which algorithm is this?
3. 16.3 One of the sorting algorithms that we discussed works like this: The smallest value in the array is located and moved to element 0. Then the next smallest value is located and moved to element 1. This process continues until all of the elements have been placed in their proper order. Which algorithm is this?
4. 16.4 One of the algorithms we discussed divides an array into two sublists, with a pivot value between them. It arranges the sublists so that all the values less than the pivot are stored in the left sublist and all the values greater than the pivot are stored in the right sublist. The algorithm recursively sorts the sublists in this same way. When all of the recursive calls have completed, the array is completely sorted. Which algorithm is this?

16.2 Introduction to Search Algorithms

CONCEPT: A search algorithm is a method of locating a specific item in a larger collection of data. This section discusses two algorithms for searching the contents of an array.

It's very common for programs not only to store and process data stored in arrays, but also to search arrays for specific items. This section will show you two methods of searching an array: the sequential search and the binary search. Each has its advantages and disadvantages.

The Sequential Search Algorithm

The *sequential search algorithm* uses a loop to step sequentially through an array, starting with the first element. It compares each element with the value being searched for and stops when the value is found or the end of the array is encountered. If the value being searched for is not in the array, the algorithm unsuccessfully searches to the end of the array.

Here is the pseudocode for a method that performs the sequential search:

```
Set found to false
Set position to -1
Set index to 0
While found is false and index < number of elements
    If list[index] is equal to search value
        found = true
        position = index
    End If
    Add 1 to index
End While
Return position
```

This pseudocode method returns the subscript of the element holding the value. If the value is not found, however, it returns -1 . The reason -1 is returned when the value is not found is that -1 is not a valid array subscript, and it cannot be mistaken as a position within the array.

The `IntSequentialSearcher` class in [Code Listing 16-12](#) has a public static method that performs a sequential search on an `int` array. The program in [Code Listing 16-13](#) tests the method.

Code Listing 16-12 (IntSequentialSearcher.java)

```
1 /**
2     The IntSequentialSearcher class provides a public static
```

```

3  method for performing a sequential search on an int array.
4  */
5
6  public class IntSequentialSearcher
7  {
8      /**
9       * The search method searches an array for a value.
10      * @param array The array to search.
11      * @param value The value to search for.
12      * @return The subscript of the value if found in the
13      *         array, otherwise -1.
14      */
15
16     public static int search(int[] array, int value)
17     {
18         int index;        // Loop control variable
19         int position;     // Position the value is found at
20         boolean found;    // Flag indicating search results
21
22         // Element 0 is the starting point of the search.
23         index = 0;
24
25         // Store the default values position and found.
26         position = -1;
27         found = false;
28
29         // Search the array.
30         while (!found && index < array.length)
31         {
32             if (array[index] == value)
33             {
34                 found = true;
35                 position = index;
36             }
37             index++;
38         }
39
40         // Return the found element's position,
41         // or -1 if not found.
42         return position;
43     }
44 }

```

Code Listing 16-13

(SequentialSearchTest.java)

```

1  /**
2   * This program sequentially searches an
3   * int array for a specified value.
4   */
5
6  public class SequentialSearchTest
7  {
8      public static void main(String[] args)

```

```

9      {
10         int[] tests = { 87, 75, 98, 100, 82 };
11         int results;
12
13         // Search the array for the value 100.
14         results = IntSequentialSearcher.search(tests, 100);
15
16         // Determine whether 100 was found and
17         // display an appropriate message.
18         if (results == -1)
19         {
20             System.out.println("You did not " +
21                               "earn 100 on any test.");
22         }
23         else
24         {
25             System.out.println("You earned 100 " +
26                               "on test " + (results + 1));
27         }
28     }
29 }

```

Program Output

```
You earned 100 on test 4
```

Average Case Analysis of the Sequential Search

The advantage of the sequential search is its simplicity. It is very easy to understand and implement. Sequential search is very flexible, and can be used on any array. But this flexibility comes at a cost: If the array being searched contains 20,000 elements, the algorithm may have to look at all 20,000 elements. This might happen, for example, when the value being searched for is not in the array, or when it is found in the very last array entry.

In an average case, an item is just as likely to be found near the beginning of the array as near the end. Typically, for an array of n items, the sequential search will locate an item in $n/2$ attempts. If an array has 50,000 elements, the sequential search will make a comparison with 25,000 of them in a typical case. This is assuming, of course, that the search item has equal probability of being found at any position in the array. ($n/2$ is the average number of comparisons. The maximum number of comparisons is always n .)

The Binary Search

The *binary search* is a clever algorithm that is much more efficient than the sequential search. Its only requirement is that the values in the array must be sorted in some order. Let us suppose the input array is sorted in ascending order. Instead of testing the array's first element, this algorithm starts with the element in the middle. If that element happens to contain the desired value, then the search is over. Otherwise, the value in the middle element is either greater than or less than the value being searched for. If it is greater, then the desired value (if it is in the list) will be found somewhere in the first half of the array. If it is less, then the desired value (again, if it is in the list) will be found somewhere in the last half of the array. In either case, half of the array's elements have been eliminated from further searching.

If the desired value isn't found in the middle element, the procedure is repeated for the half of the array that potentially contains the value. For instance, if the last half of the array is to be searched, the algorithm tests *its* middle element. If the desired value isn't found there, the search is narrowed to the quarter of the array that resides before or after that element. This process continues until the value being searched for is either found, or there are no

more elements to test. Here is the pseudocode for a method that performs a binary search on an array:

```
Set first to 0
Set last to the last subscript in the array
Set found to false
Set position to -1
While found is not true and first is less than or equal to last
    Set middle to the subscript half-way between array[first] and array[last]
    If array[middle] equals the desired value
        Set found to true
        Set position to middle
    Else If array[middle] is greater than the desired value
        Set last to middle - 1
    Else
        Set first to middle + 1
    End If
End While
Return position
```

This algorithm uses three variables to mark positions within the array: `first`, `last`, and `middle`. The `first` and `last` variables mark the boundaries of the portion of the array currently being searched. They are initialized with the subscripts of the array's `first` and `last` elements. The subscript of the element halfway between `first` and `last` is calculated and stored in the `middle` variable. If the element in the middle of the array does not contain the search value, the `first` or `last` variables are adjusted so that only the top or bottom half of the array is searched during the next iteration. This cuts the portion of the array being searched in half each time the loop fails to locate the search value.

The `IntBinarySearcher` class shown in [Code Listing 16-14](#) has a public static method that performs a binary search on an integer array. The first parameter, `array`, is searched for an occurrence of the number stored in `value`. If the number is found, its array subscript is returned. Otherwise, `-1` is returned indicating the value did not appear in the array. The program in [Code Listing 16-15](#) tests the method.

Code Listing 16-14 (IntBinarySearcher.java)

```
1 /**
2  The IntBinarySearcher class provides a public static
3  method for performing a binary search on an int array.
4  */
5
6 public class IntBinarySearcher
7 {
8     /**
9      The search method performs a binary search on an int
10     array. The array is searched for the number passed to
11     value. If the number is found, its array subscript is
12     returned. Otherwise, -1 is returned indicating the
13     value was not found in the array.
14     @param array The array to search.
15     @param value The value to search for.
16     */
17
18     public static int search(int[] array, int value)
19     {
20         int first;        // First array element
21         int last;         // Last array element
22         int middle;       // Mid point of search
23         int position;     // Position of search value
24         boolean found;    // Flag
25
26         // Set the initial values.
27         first = 0;
28         last = array.length - 1;
29         position = -1;
30         found = false;
31
32         // Search for the value.
33         while (!found && first <= last)
34         {
35             // Calculate midpoint
36             middle = (first + last) / 2;
37
38             // If value is found at mid point...
39             if (array[middle] == value)
40             {
41                 found = true;
42                 position = middle;
43             }
44             // else if value is in lower half...
45             else if (array[middle] > value)
46                 last = middle - 1;
47             // else if value is in upper half....
48             else
49                 first = middle + 1;
50         }
51
52         // Return the position of the item, or -1
53         // if it was not found.
```

```

54         return position;
55     }
56 }

```

Code Listing 16-15 (BinarySearchTest.java)

```

1  import java.util.Scanner;
2
3  /**
4   * This program demonstrates the search method in
5   * the IntBinarySearcher class.
6   */
7
8  public class BinarySearchTest
9  {
10     public static void main(String[] args)
11     {
12         int result, searchValue;
13         String input;
14
15         // An array of numbers to search.
16         int numbers[] = { 536, 289, 296, 429, 319, 142, 394,
17                          101, 388, 147, 417, 199, 207, 222,
18                          189, 310, 447, 521, 234, 600};
19
20         // Create a Scanner object for keyboard input.
21         Scanner keyboard = new Scanner(System.in);
22
23         // First we must sort the array in ascending order.
24         IntQuickSorter.quickSort(numbers);
25
26         do
27         {
28             // Get a value to search for.
29             System.out.print("Enter a value to search for: ");
30             searchValue = keyboard.nextInt();
31
32             // Search for the value
33             result = IntBinarySearcher.search(numbers, searchValue);
34
35             // Display the results.
36             if (result == -1)
37                 System.out.println(searchValue + " was not found.");
38             else
39             {
40                 System.out.println(searchValue + " was found at " +
41                                   "element " + result);
42             }
43
44             // Consume the remaining newline...
45             keyboard.nextLine();
46
47             // Does the user want to search again?
48             System.out.print("Do you want to search again? (Y or N): ");
49             input = keyboard.nextLine();

```

```
50         } while (input.charAt(0) == 'y' || input.charAt(0) == 'Y');
51     }
52 }
```

Program Output with Example Input Shown in Bold

```
Enter a value to search for: 289 [Enter]
289 was found at element 8
Do you want to search again? (Y or N): y [Enter]
Enter a value to search for: 101 [Enter]
101 was found at element 0
Do you want to search again? (Y or N): y [Enter]
Enter a value to search for: 999 [Enter]
999 was not found.
Do you want to search again? (Y or N): n [Enter]
```

Efficiency of the Binary Search

For sorted arrays, the binary search is much more efficient than the sequential search. Every time it makes a comparison and fails to find the desired item, it eliminates half of the remaining portion of the array that must be searched. For example, consider an array with 1,000 elements.

If the binary search fails to find an item on the first attempt, the number of elements that remains to be searched is 500. If the item is not found on the second attempt, the number of elements that remains to be searched is 250. This process continues until the binary search has either located the desired item or determined that it is not in the array. With 1,000 elements this takes no more than 10 comparisons. (Compare this to the sequential search, which would make an average of 500 comparisons!)

Powers of 2 are used to calculate the maximum number of comparisons the binary search will make on an array of any size. Simply find the smallest power of 2 that is greater than or equal to the number of elements in the array. For example, a maximum of 16 comparisons will be made on an array of 50,000 elements ($2^{16}=65,536$), and a maximum of 20 comparisons will be made on an array of 1,000,000 elements ($2^{20}=1,048,576$).

A Recursive Binary Search

The binary search algorithm that we previously examined used an iterative approach, meaning that it used a loop to step through the array. The binary search algorithm can also be implemented recursively. For example, the procedure can be expressed as follows:

```
If array[middle] equals the search value, then the value is found
Else if array[middle] is less than the search value, perform a binary
```

search on the upper half of the array

*Else if array[middle] is greater than the search value, perform
a binary search on the lower half of the array*

End If

When you compare the recursive algorithm to its iterative counterpart, it becomes evident that the recursive version is much more elegant and easier to understand. The recursive binary search algorithm is also a good example of repeatedly breaking a problem down into smaller pieces until it is solved. [Code Listing 16-16](#) shows the `RecursiveBinarySearcher` class. This class provides a public static method named `search`, which calls the recursive `binarySearch` method to find a value in an `int` array. The program in [Code Listing 16-17](#) tests the method.

Code Listing 16-16

(RecursiveBinarySearcher.java)

```
1 /**
2     The RecursiveBinarySearcher class provides a public static
3     method for performing a recursive binary search on an int array.
4 */
5
6 public class RecursiveBinarySearcher
7 {
8     /**
9         The search method calls the doBinarySearch method
10        to search for a value in an array.
11        @param array The array to search.
12        @param value The value to search for.
13    */
14
15    public static int search(int[] array, int value)
16    {
17        return binarySearch(array, 0, array.length - 1, value);
18    }
19
20    /**
21        The binarySearch method performs a recursive binary
22        search on an integer array.
23        @param array The array to search.
24        @param first The first element in the search range.
25        @param last The last element in the search range.
26        @param value The value to search for.
27        @return The subscript of the value if found,
28                otherwise -1.
29    */
30
31    private static int binarySearch(int[] array, int first,
32                                   int last, int value)
33    {
34        int middle;        // Mid point of search
35
36        // Test for the base case where the
37        // value is not found.
```



```

38     if (first > last)
39         return -1;
40
41     // Calculate the middle position.
42     middle = (first + last) / 2;
43
44     // Search for the value.
45     if (array[middle] == value)
46         return middle;
47     else if (array[middle] < value)
48         return binarySearch(array, middle + 1,
49                             last, value);
50     else
51         return binarySearch(array, first,
52                             middle - 1, value);
53 }
54
55 }

```

Code Listing 16-17

(RecursiveBinarySearchTest.java)

```

1 import java.util.Scanner;
2
3 /**
4  * This program tests the search method in the
5  * RecursiveBinarySearcher class.
6  */
7
8 public class RecursiveBinarySearchTest
9 {
10     public static void main(String [] args)
11     {
12         int result, searchValue;
13         String input;
14
15         // An array of numbers to search.
16         int numbers[] = { 536, 289, 296, 429, 319, 142, 394,
17                         101, 388, 147, 417, 199, 207, 222,
18                         189, 310, 447, 521, 234, 600};
19
20         // Create a Scanner object to read keyboard input.
21         Scanner keyboard = new Scanner(System.in);
22
23         // First we must sort the array in ascending order.
24         IntQuickSorter.quickSort(numbers);
25
26         do
27         {
28             // Get a value to search for.
29             System.out.print("Enter a value to search for: ");
30             searchValue = keyboard.nextInt();
31
32             // Search for the value

```

```

33         result =
34             RecursiveBinarySearcher.search(numbers, searchValue);
35
36         // Display the results.
37         if (result == -1)
38             System.out.println(searchValue + " was not found.");
39         else
40         {
41             System.out.println(searchValue + " was found at " +
42                 "element " + result);
43         }
44
45         // Consume the remaining newline.
46         keyboard.nextLine();
47
48         // Does the user want to search again?
49         System.out.print("Do you want to search again? (Y or N): ");
50         input = keyboard.nextLine();
51     } while (input.charAt(0) == 'y' || input.charAt(0) == 'Y');
52 }
53 }

```

Program Output with Example Input Shown in Bold

```

Enter a value to search for: 289 [Enter]
289 was found at element 8
Do you want to search again? (Y or N): y [Enter]
Enter a value to search for: 101 [Enter]
101 was found at element 0
Do you want to search again? (Y or N): y [Enter]
Enter a value to search for: 999 [Enter]
999 was not found.
Do you want to search again? (Y or N): n [Enter]

```

Checkpoint

1. 16.5 Describe the difference between the sequential search and the binary search.
2. 16.6 On average, with an array of 20,000 elements, how many comparisons will the sequential search perform? (Assume the items being searched have equal probability of being found at any of the positions in the array.)
3. 16.7 With an array of 20,000 elements, what is the maximum number of comparisons the binary search will perform?
4. 16.8 If a sequential search is performed on an array, and it is known that some items are

searched for more frequently than others, how can the contents of the array be reordered to improve the average performance of the search?

16.3 Analysis of Algorithms

CONCEPT: We can estimate the efficiency of an algorithm by counting the number of steps that the algorithm requires to solve a problem.

An algorithm is a mechanical step-by-step procedure for solving a problem and is the basic strategy used in designing a program. There is often more than one algorithm that can be used to solve a given problem. For example, we saw in the previous section that the problem of searching a sorted array can be solved by two different methods: *sequential search* and *binary search*.

How can we decide which of two algorithms for solving a problem is better? To answer this question, we need to establish criteria for judging the “goodness” or efficiency of an algorithm. The two criteria that are most often used are space and time. The *space* criterion refers to the amount of memory the algorithm requires to solve the problem, while the *time* criterion refers to the length of execution time. In this chapter, we will use the time criterion to evaluate the efficiency of algorithms.

One possibility for comparing two algorithms is to code them and then time the execution of the resulting Java programs. This experimental approach can yield useful information, but it has the following shortcomings:

- It measures the efficiency of programs rather than algorithms.
- The results obtained depend on the programming language used to code the algorithms, and on the quality of the compiler used to generate machine code. The programs may run faster or slower if they are coded in a different language, or compiled by a different compiler.
- The results obtained depend on how the operating system executes programs, and on the nature of the hardware on which the programs are executing. The execution times may be different if we run the programs on a different computer and a different operating system.
- The results obtained apply only to those inputs that were part of the execution runs and may not be representative of the performance of the algorithms on a different set of inputs.

A better approach is to count the number of basic steps an algorithm requires to process an input of a given size. To make sense of this approach, we need more precise definitions of the concepts of a computational problem, problem input, input size, and basic step.

Computational Problems and Basic Steps

A *computational problem* is a problem to be solved using an algorithm. Such a problem is a collection of *instances*, with each instance being specified by input data given in some prescribed format. For example, if the problem P is to sort an array of integers, then an instance of P is just a

specific integer array. The *size* of an instance refers to the amount of memory needed to hold the input data. The input size is usually given as a number that allows us to infer the total number of bits occupied by the input data. If the number of bits occupied by an entry of the array is fixed, say at 64 bits, then the length of the array is a good measure of input size. In contrast, the length of the array is not a good measure of input size if there is no a priori bound on the size of array elements.

A step executed by an algorithm is a *basic step* (also called a *basic operation*) if the algorithm can execute the step in time bounded by a constant regardless of the size of the input. In sorting an array of integers, the step

- *Swap the elements in positions k and $k+1$*

is basic because the time required to swap two array elements remains constant even if the length of the array increases. In contrast, a step such as

- *Find the largest element of the array*

is not basic because the time required to complete the step depends on the length of the array. Intuitively, a basic step is one that can conceivably be built into the hardware of some physical computers.

The definition of a basic step does not specify the size of the constant that bounds the time required to execute the step. Ignoring the exact value of these constants reflects the reality that the same operation may be executed with different speeds on different hardware, and that an operation that can be executed with one hardware instruction on one computer may require several hardware instructions on another computer. A consequence of this definition is that we can count any constant number of basic steps as one basic step. In particular, an algorithm that executes $500n$ basic steps can accurately be described as executing n basic steps.

It is important to realize that ordinary arithmetic and logic operations such as addition and comparison are not basic unless a constant bound is put on the size of the numbers being added or compared. The size of the bound does not matter as long as the bound is constant: it may be 32, 64, 128, 1024 bits, or even larger, and these operations will still be basic. In the following discussion, we assume that numbers that occur in our algorithms as inputs, outputs, or are computed as intermediate results are all bounded in size. This allows us to consider operations on them as basic.

It only makes sense to describe an algorithm after we have described the problem the algorithm is supposed to solve. A computational problem is described by giving, for each positive integer n , the format of a typical input of size n followed by a description of output to be produced by an algorithm that solves the problem. For example, the problem of summing an integer array can be described as follows:

- INPUT: An integer array $a[]$ of size n
- SIZE OF INPUT: The number n of array entries

- OUTPUT: An integer *sum* representing the sum total of the array entries

In this context, an algorithm for solving the problem can be given as follows:

ALGORITHM:

```

1:  sum = 0
2:  k = 0      //array index
3:  While k < n do
4:      sum = sum + a[k]
5:      k = k + 1
6:  End While

```

Complexity of Algorithms

We can determine the number of basic steps executed by this algorithm on an input of size n by reasoning as follows. The algorithm consists of two statements (Statements 1 and 2) and a loop (Statements 3–6). The two statements before the loop can be counted as one basic operation. Clearly, the amount of time the algorithm takes to execute a *single* iteration of the loop does not depend on n , and so is constant. (This can also be seen by noting that each statement inside the loop is a basic operation.) This allows us to count each loop iteration as a single basic operation. Because the algorithm executes the loop n times, it executes $n+1$ basic operations altogether. For large problem sizes, the number of basic operations performed is approximately n . Because of the hidden constant factors, we say that the algorithm requires time proportional to n to process an input of size n .

Notice that in our analysis of the array-summing algorithm, we could have gotten the same result by counting the number of loop iterations performed by the algorithm. Moreover, notice that the crucial operation in summing an array is the addition of array elements (Statement 4), and that there are as many additions of array elements as there are loop iterations. Thus, we could have gotten the same result by just counting additions of array elements. It turns out that for most algorithms, it suffices to identify and count only one or two basic operations that are in some way crucial to the problem being solved. For example, in many array searching and sorting algorithms, it suffices to count the number of comparisons between array elements.

The array-summing algorithm just considered is particularly simple to analyze because it performs the same amount of work for all inputs of a given size. Consider the problem of searching an array of integers for a given integer X , solved here by the familiar sequential search algorithm:

- INPUT: An integer array $a[]$ of size n , and an integer X
- SIZE OF INPUT: The number n of array entries
- OUTPUT: An integer k in the range $0 \leq k \leq n-1$ such that $a[k]=X$, or $k=n$

ALGORITHM:

```

k = 0
While k < n and a[k] != X do
    k = k + 1;

```

End While

This algorithm starts at one end and searches sequentially through the array. The algorithm stops as soon as it encounters X , but will search the entire array if X is not in the array. The algorithm may stop after making only one comparison (X is found in the first entry examined), or it may stop after making n comparisons (X is found in the last place examined, or is not in the array). In fact, the algorithm may perform m comparisons for any value of m between 1 and n . In cases where the algorithm may perform different amounts of work for different inputs of the same size, it is common to measure the efficiency of the algorithm by the work done on an input of size n that requires the most work. This is called measuring the algorithm by its *worst-case complexity function*.

Worst Case Complexity of Algorithms

The worst-case complexity function $f(n)$ of an algorithm is the number of steps it performs on an input of size n that requires the most work. It gives an indication of the longest time an algorithm will ever take to solve an instance of size n , and is a good measure of efficiency to use when we are looking for a performance guarantee.

Let us determine the worst-case complexity of binary search. This algorithm is used to locate an item X in an array sorted in ascending order. The worst case occurs when X is not found in the array. We will show that in this case, the algorithm performs $L+1$ steps, where L is the number of loop iterations.

As shown earlier in this chapter, this algorithm consists of some initialization of some variables followed by a loop. The initialization requires constant time, and can therefore be considered to be one basic operation. Likewise, each iteration of the loop is a basic step (increasing the number of entries in the array does not increase the amount of time required by a single iteration of the loop). This shows that the number of steps required by binary search is $L+1$. Now L is approximately equal to the integer part of $\log n$, the logarithm of n to the base 2. To see this, notice that the size of the array to be searched is initially n , and each iteration reduces the size of the remaining portion of the array by one half. Because each loop iteration performs at most two comparisons, binary search performs a total of $2 \log n$ comparisons. By ignoring the constant factor, we can summarize our findings as follows:

- *In the worst case, binary search requires time proportional to $\log n$.*

For our last example of determining the complexity of an algorithm, consider the following outline of a variant of the selection sort algorithm, also studied earlier in this chapter:

- INPUT: An array $a[]$ of n integers
- SIZE OF INPUT: The number n of array entries
- OUTPUT: The array $a[]$ rearranged so that $a[0] \leq a[1] \leq \dots \leq a[n-1]$

ALGORITHM:

```
1:   For (k = n-1; k >= 1; k--)
2:       // a[0..k] is what remains to be sorted
3:       determine position p of largest entry in a[0..k]
4:       swap a[p] with a[k]
5:   End For
```

Let us determine the number of comparisons (involving array entries) executed by this algorithm in sorting an array of n entries. All such comparisons occur in step 3, which is clearly not a basic step (it requires time proportional to k , which can be as large as n). Let us express step 3 using operations that are clearly basic:

- INPUT: array $a[0..k]$ of $k+1$ entries
- SIZE OF INPUT: number $k+1$ of array entries

```
6:   p = 0    //position of largest
7:   For (m = 1; m <= k; m++)
8:       If a[m] > a[p] Then
9:           p = m
10:      End if
11:   End For
```

It is clear that this algorithm requires k comparisons between array entries.

Now returning to the main sorting algorithm, we observe that there will be $n-1$ iterations of the loop that starts at line 1 and ends at line 5, with one such iteration for each value of k in the range $n-1$ to 1. Moreover, on the k -th such iteration, the step on line 3 will perform k comparisons:

$k = n-1$: step 3 performs $n-1$ comparisons $k = n-2$: step 3 performs $n-2$ comparisons ... $k = 1$: step 3 performs 1 comparison

Now we see that the total number of comparisons performed by this simple sorting algorithm is given by the expression

$$1 + 2 + 3 + \dots + n-1 = (n-1)n/2$$

For large n , this expression is very close to $n^2/2$. We have proved the following:

- In the worst case, selection sort requires time proportional to n^2 .

Average Case Complexity

The worst-case complexity will not give a good indication of how an algorithm will perform in practical situations in which inputs that yield worst case performance are rare. An alternative method, the *average case complexity function*, can be used if the relative frequencies with which the different inputs are likely to occur in practice are known. The average case complexity function uses such

frequencies to form a weighted average of the number of steps performed on each input. Although it yields a good measure of the expected performance of an algorithm, it is difficult to use in practice because reliable estimates of input frequencies are usually not available.

Asymptotic Complexity and the Big O Notation

We can compare two algorithms F and G for solving a problem by comparing their complexity functions. More specifically, if $f(n)$ and $g(n)$ are the complexity functions for the two algorithms, we can compare the algorithms against each other by looking at what happens to the ratio $f(n)/g(n)$ when n gets large. This is easiest to understand if this ratio tends to some limit. Let us consider some specific examples. Throughout, we assume that $f(n) \geq 1$ and $g(n) \geq 1$ for all $(n) \geq 1$.

- $f(n)=3n^2+5n$ and $g(n)=n^2$. In this case

$$f(n)/g(n) = (3n^2 + 5n)/n^2 = 3 + 5/n \rightarrow 3 \quad \text{as } n \rightarrow \infty$$

That is, the value of $f(n)/g(n)$ gets closer and closer to 3 as n gets large. What this means is that for very large input sizes F performs three times as many basic operations as G . Because the two algorithms differ in performance by a constant factor, we consider them to be equivalent in efficiency.

- $f(n)=3n^2+5n$ and $g(n)=n$. In this case

$$f(n)/g(n) = (3n^2 + 5n)/n = 3n + 5 \rightarrow \infty \quad \text{as } n \rightarrow \infty$$

Here, the ratio $f(n)/g(n)$ gets larger and larger as n gets large. This means F does a lot more work than G on large input sizes. This makes G the better algorithm for large inputs.

- $f(n)=3n^2+5n$ and $g(n)=n^3$. In this case

$$f(n)/g(n) = (3n^2 + 5n)/n^3 = 3/n + 5/n^2 \rightarrow 0 \quad \text{as } n \rightarrow \infty$$

This means that for large inputs the algorithm G is doing a lot more work than F , making F the more efficient algorithm.

In general, we can compare two complexity functions $f(n)$ and $g(n)$ by looking at what happens to $f(n)/g(n)$ as n gets large. Although thinking in terms of a limit of this ratio is helpful in comparing the two algorithms, we cannot assume that such a limit will always exist. It turns out that a limit does not have to exist for us to gain useful information from this ratio: we can usefully compare the two complexity functions if we can find a positive constant K such that

$$f(n)/g(n) \leq K \quad \text{for all } n \geq 1.$$

(16.1)

If this can be done, it means that the algorithm F is no worse than G for large problems, or in other words, F is at least as good as G . In this case, we say that $f(n)$ is in $O(g(n))$, pronounced “ f is in Big O of g .” The condition that defines $f(n)$ is in $O(g(n))$ is often written in the equivalent form shown in [Equation \(16.2\)](#).

$$f(n) \leq K g(n) \text{ whenever } n \geq 1.$$

(16.2)

Showing that $f(n)$ is in $O(g(n))$ is usually straightforward: you look at the ratio $f(n)/g(n)$ and try to find a positive constant K that satisfies [Equation \(16.1\)](#) for all $n \geq 1$. For example, to show that $3n^2 + 5n$ is in $O(n^2)$, look at the ratio

$$\frac{3n^2 + 5n}{n^2} = 3 + \frac{5}{n}$$

and notice that $5/n$ will be at most 5 for all $n \geq 1$. You can then use $K=8$ in [Equation \(16.1\)](#).

To show that $f(n)$ is *not* in $O(g(n))$, you have to show that there is no way to find a positive K that will satisfy [Equation \(16.1\)](#) for all $n \geq 1$. For example, the function $3n^2 + 5n$ is not in $O(n)$ because there is no constant K that satisfies

$$\frac{3n^2 + 5n}{n} = 3n + 5 \leq K \text{ for all } n \geq 1.$$

Although defined for functions, the “Big O” notation and terminology is also used to characterize algorithms and computational problems. Thus, we say that an algorithm F is in $O(g(n))$ for some function $g(n)$ if the complexity function $f(n)$ of F is in Big O of $g(n)$. Accordingly, sequential search of an array is in $O(n)$ whereas binary search is in $O(\log n)$.

Similarly, a computational problem is said to be in $O(g(n))$ if there exists an algorithm for the problem whose worst case complexity function is in $O(g(n))$. Thus, the problem of sorting an array is in $O(n^2)$, whereas the problem of searching a sorted array is in $O(\log n)$.

If $g(n)$ is a function, $O(g(n))$ can be regarded as a family of functions that grow no faster than $g(n)$. These families are called complexity classes, and a few of them are important enough to merit specific names. We list them here in order of their rate of growth:

1. $O(1)$: A function $f(n)$ is in this class if there is a constant $K > 0$ such that $f(n) \leq K$ for all $n \geq 1$. An algorithm whose worst case complexity function is in this class is said to run in *constant time*.
2. $O(\log n)$: Algorithms in this class run in *logarithmic time*. Because $\log n$ grows much slower than n , a huge increase in the size of the problem results in a small increase in the running time of the algorithm. This complexity is characteristic of search problems that eliminate half of the search space with each operation. The *binary search* algorithm is in this class.
3. $O(n)$: Algorithms in this class run in *linear time*. Any increase in the size of the problem results in a proportionate increase in the running time of the algorithm. This complexity is characteristic of algorithms like *sequential search* that make a single pass, or a constant number of passes,

over their input.

4. $O(n \log n)$: This class is called “ $n \log n$ ” *time*. An increase in the size of the problem results in a slightly greater increase in the running time of the algorithm. The average case complexity of *Quicksort* lies in this class. Two well-known sorting algorithms, *Mergesort* and *Heapsort*, have worst case complexity functions that also lie in this class. We will discuss the *Heapsort* algorithm in [Chapter 22](#).
5. $O(n^2)$: This class is called *quadratic time*. This performance is characteristic of algorithms that make multiple passes over the input data using two nested loops. An increase in the size of the problem causes a much greater increase in the running time of the algorithm. The worst case complexity functions of *bubble sort*, *selection sort*, *insertion sort*, and *Quicksort* all lie in this class.
6. The union of all classes $O(n^1)$, $O(n^2)$, $O(n^3)$, ..., forms a class called *polynomial time*.



In the Spotlight:

The Importance of Choosing an Efficient Algorithm

The algorithm selected for a particular task can have a significant impact on the running time of a program. As an example, let us consider the task of computing terms of the Fibonacci sequence. You first encountered this sequence in [Chapter 15](#), and as you may recall, terms can be computed by this method:

```
public static long fib(long n)
{
    if (n <= 1)
        return 1;
    else
        return fib(n-2) + fib(n-1);
}
```

This implementation has the advantages of being short and easy to understand, and for that reason, it is often used to illustrate recursion. It is, however, very inefficient. We can get an idea of how inefficient it is by embedding it in a program that determines the time required to execute calls to the method. The following program allows a user to specify six consecutive values to be passed as arguments to the method and prints the time needed by the computer to execute each call.

Code Listing 16-18

(RecursiveFibonacciTimer.java)

```

1 import java.util.Scanner;
2
3 /**
4  This program times calls to the recursive Fibonacci method
5  for 6 consecutive calls and displays the results.
6  */
7
8 public class RecursiveFibonacciTimer
9 {
10     public static void main(String[] args)
11     {
12         // Get the starting argument
13         System.out.print("Enter a positive integer: ");
14         Scanner sc = new Scanner(System.in);
15         int number = sc.nextInt();
16
17         // Variables used to determine time for a function call
18         long currentTime = System.currentTimeMillis();
19         long previousTime;
20         long elapsedTime = 0;
21
22         for (int k = 0; k <= 5; k++)
23         {
24             // Record time before calling the method
25             previousTime = currentTime;
26             System.out.print("The Fibonacci term at position ");
27             System.out.print((number + k) + " is " );
28
29             // Compute and print fib term for the next argument
30             System.out.println(fib(number + k));
31
32             // Record time after calling the method
33             currentTime = System.currentTimeMillis();
34
35             // Compute and print elapsed time in seconds
36             elapsedTime = (currentTime - previousTime) / 1000;
37             System.out.println("Computed in " + elapsedTime + " seconds.");
38         }
39     }
40
41     /**
42      Computes a term of the Fibonacci sequence
43      @param n
44      @return nth term of the sequence
45     */
46     public static long fib(long n)
47     {
48         if (n <= 1)
49             return 1;
50         else
51             return fib(n-2) + fib(n-1);
52     }
53 }

```

Program Output with Example Input Shown in

Bold

```
Enter a positive integer: 45 [Enter]
The Fibonacci term at position 45 is 1836311903
Computed in 17 seconds.
The Fibonacci term at position 46 is 2971215073
Computed in 27 seconds.
The Fibonacci term at position 47 is 4807526976
Computed in 43 seconds.
The Fibonacci term at position 48 is 7778742049
Computed in 67 seconds.
The Fibonacci term at position 49 is 12586269025
Computed in 110 seconds.
The Fibonacci term at position 50 is 20365011074
Computed in 179 seconds.
```

In line 18, the program uses the `System.currentTimeMillis()` method to get the current time in milliseconds since midnight January 1, 1970. The value obtained is stored in the `currentTime` variable. At the beginning of each loop iteration, the current time is assigned to the `previousTime` variable in line 25. Then, after the call to the `fib` method, the current time is obtained again in line 33 and assigned to the `currentTime` variable. The time difference, in milliseconds, is calculated in line 36 by subtracting `previousTime` from `currentTime`. This difference is divided by 1000 to yield the time elapsed during the method call in seconds.

You can tell that the time required to compute a term of the sequence increases significantly for each successive term. For example, the computation of the term at position 45 requires a quarter of a minute, and five terms later, the 50th term requires nearly three minutes. And it gets worse from there. Clearly, this algorithm, although simple to write, is so inefficient as to be unusable for problems of moderately large sizes.

Checkpoint

1. 16.9 What is a basic operation?
2. 16.10 What is the worst case complexity function of an algorithm?
3. 16.11 One algorithm needs $100n$ basic operations to process an input of size n , and another algorithm needs $25n$ basic operations to process the same input. Which of the two algorithms is more efficient?
4. 16.12 What does it mean to say that $f(n)$ is not in $O(g(n))$??
5. 16.13 What does it mean to say that $f(n)$ is in $O(g(n))$??
6. 16.14 Show that $100n^3 + 50n^2 + 75$ is in $O(20n^3)$ by finding a positive K that satisfies

[Equation \(16.1\)](#) in this section.

7. 16.15 Let $a[]$ and $b[]$ be two integer arrays of size n . Following the examples of this section, give a formal description of the problem of determining if every element of $a[]$ is also an element of $b[]$. The output of the algorithm should be one of the words “true” or “false.”
8. 16.16 Give an algorithm for solving the problem you defined in Checkpoint 16.15, and determine its worst case complexity.
9. 16.17 Show that every function in $O(20)$ also is in $O(1)$.
10. 16.18 Assuming $g(n) \geq 1$ for all $n \geq 1$ show that every function in $O(g(n) + 100)$ is also in $O(g(n))$.

16.4 Common Errors to Avoid

- Using an inefficient sort or search algorithm on a large array. Simple algorithms like the bubble sort and the sequential search are inefficient because they access array elements so many times. More efficient algorithms like Quicksort and the binary search are preferable, especially on large arrays.
- Forgetting to sort the data in an array before using the binary search algorithm. The binary search algorithm requires that its data already be sorted.
- Using timing to determine the efficiency of an algorithm. The results obtained from timing an application depend on many variables, including the programming language used and on the quality of the compiler used to generate machine code. A better approach is to count the number of basic steps an algorithm requires to process an input of a given size.

Review Questions and Exercises

Multiple Choice and True/False

1. This search algorithm steps sequentially through an array, comparing each item with the search value.
 1. sequential search
 2. binary search
 3. natural order search
 4. selection search
2. This search algorithm repeatedly divides the portion of an array being searched in half.
 1. sequential search
 2. binary search
 3. natural order search
 4. selection search
3. This search algorithm will search half the array on average.
 1. sequential search
 2. binary search
 3. natural order search
 4. selection search
4. This search algorithm requires that the array's contents be sorted.
 1. sequential search
 2. binary search
 3. natural order search
 4. selection search

5. If an array is sorted in this order, the values are stored from lowest to highest.
 1. asymptotic
 2. logarithmic
 3. ascending
 4. descending
6. If an array is sorted in this order, the values are stored from highest to lowest.
 1. asymptotic
 2. logarithmic
 3. ascending
 4. descending
7. This sorting algorithm makes several passes through an array and causes the larger values to gradually move toward the end of the array with each pass.
 1. bubble sort
 2. selection sort
 3. insertion sort
 4. Quicksort
8. This sorting algorithm recursively divides an array into sublists.
 1. bubble sort
 2. selection sort
 3. insertion sort
 4. Quicksort
9. In this sorting algorithm, the smallest value in the array is located and moved to position 0. Then the next smallest value is located and moved to position 1. This process continues until all of the elements have been placed in their proper order.
 1. bubble sort
 2. selection sort

3. insertion sort
 4. Quicksort
10. This sorting algorithm begins by sorting the initial portion of the array consisting of two elements. Then, the third element is moved to its correct position, relative to the first two elements. At that point, the first three elements are in sorted order. This process continues with the fourth and subsequent elements until the entire array is sorted.
1. bubble sort
 2. selection sort
 3. insertion sort
 4. Quicksort
11. **True or False:** If data is sorted in ascending order, it means it is ordered from lowest value to highest value.
12. **True or False:** If data is sorted in descending order, it means it is ordered from lowest value to highest value.
13. **True or False:** The *average* number of comparisons performed by the sequential search on an array of n elements is $n/2$.
14. **True or False:** The *maximum* number of comparisons performed by the sequential search on an array of n elements is $n/2$.

Find the Error

1. Assume this code is using the IntBinarySearcher class presented in this chapter:

```
int[] numbers = { 8, 0, 9, 4, 3, 7, 2 };  
// Search for 7.  
int location = IntBinarySearcher.search(numbers, 7);  
// Display the result.  
if (location != -1)  
    System.out.println("7 was found in the array.");  
else  
    System.out.println("7 was NOT found in the array.");
```

Algorithm Workbench

1. Complete the following table calculating the average and maximum number of comparisons the

sequential search will perform, and the maximum number of comparisons the binary search will perform.

Array Size: 50 elements	Array Size: 500 elements	Array Size: 10,000 elements	Array Size: 100,000 elements	Array Size: 10,000,000 elements

Sequential Search
(Average
Comparisons)

Sequential Search
(Maximum
Comparisons)

Binary Search
(Maximum
Comparisons)

Short Answer

1. If a sequential search method is searching for a value that is stored in the last element of a 10,000-element array, how many elements will the search code have to examine to locate the value?
2. In an average case involving an array of n elements, how many array elements does a sequential search algorithm have to access to locate a specific value?
3. A binary search method is searching for a value that is stored in the middle element of an array. How many array elements have to be examined before finding the value?
4. What is the maximum number of comparisons that a binary search method will make when searching for a value in a 1,000-element array?
5. Why is the bubble sort inefficient for large arrays?
6. Under what conditions might the insertion sort be more efficient than the bubble sort?
7. List the steps that the selection sort algorithm would make in sorting the following values: 4, 1,

3, 2.

8. List the steps that the insertion sort algorithm would make in sorting the following values: 4, 1, 3, 2.

Programming Challenges

1. Sorting Objects with the Insertion Sort Algorithm

The `insertionSort` method in the `IntInsertionSorter` class presented in this chapter sorts an array of `int` values. Create an `ObjectInsertionSorter` class that can sort `Comparable` objects. Demonstrate the class in a program that sorts an array of `String` objects.

2. Sorting Objects with the Quicksort Algorithm

The `IntQuickSorter` class presented in this chapter sorts an array of `int` values. Create an `ObjectQuickSorter` class that can sort `Comparable` objects. Demonstrate the class in a program that sorts an array of `String` objects.

3. Searching for Objects with the Binary Search Algorithm

The `IntBinarySearcher` class presented in this chapter searches an `int` array for a specific value. Create an `ObjectBinarySearcher` class that can search an array of `Comparable` objects. Demonstrate the class in a program that searches for a `String` in an array of `String` objects.



Searching for Objects with the Binary Search Algorithm

4. Charge Account Validation

Create a class with a method that accepts a charge account number as its argument. The method should determine whether the number is valid by comparing it to the following list of valid charge account numbers:

5658845 4520125 7895122 8777541 8451277 1302850

8080152 4562555 5552012 5050552 7825877 1250255

1005231 6545231 3852085 7576651 7881200 4581002

These numbers should be stored in an array. Use a sequential search to locate the number passed as an argument. If the number is in the array, the method should return `true`, indicating the number is valid. If the number is not in the array, the method should return `false`, indicating the number is invalid.

Write a program that tests the class by asking the user to enter a charge account number. The program should display a message indicating whether the number is valid or invalid.

5. Charge Account Validation Modification

Modify the program you wrote for Programming Challenge 4 (Charge Account Validation) so it performs a binary search to locate valid account numbers. Use your choice of sorting algorithm to sort the array before the binary search is performed.

6. Name Search

Download the [Chapter 16](http://www.pearson.com/gaddis) source code from www.pearson.com/gaddis, and, you will find a text file by the name of *names.txt*. This file contains a list of the 65 most popular female first names in the United States according to the 1990 census. Write a program that reads these names into a `String` array. Use the Quicksort algorithm to sort the names in ascending order. Then allow the user to search for a name in the array. Use the binary search algorithm to perform the search.

7. Search Benchmarks

Write an application that has an array of at least 20 integers. It should call a method that uses the sequential search algorithm to locate one of the values. The method should keep a count of the number of comparisons it makes until it finds the value. Then the program should call another method that uses the binary search algorithm to locate the same value. It should also keep count of the number of comparisons it makes. Display these values on the screen.

8. Sorting Benchmarks

Modify the methods presented in this chapter that perform the bubble sort, selection sort, insertion sort, and Quicksort algorithms on an `int` array, such that each method keeps a count of the number of swaps it makes.

Then, write an application that uses four identical arrays of at least 20 integers. It should call each method on a different array, and display the number of exchanges made by each algorithm.

9. Efficient Computation of Fibonacci Numbers

Modify the application in [Code Listing 16-18](#) by replacing the recursive method for computing Fibonacci terms with an iterative version and try to make it as efficient as possible. Modify the application so it outputs the elapsed time in both seconds and milliseconds. Finally, run the application and generate output showing the time required to compute the terms of the sequence in positions 45, 46, 47, 48, 49, and 50. Comment on the results.

Chapter 17 Generics

TOPICS

1. [17.1 Introduction to Generics](#)
2. [17.2 Writing a Generic Class](#)
3. [17.3 Passing Objects of a Generic Class to a Method](#)
4. [17.4 Writing Generic Methods](#)
5. [17.5 Constraining a Type Parameter in a Generic Class](#)
6. [17.6 Inheritance and Generic Classes](#)
7. [17.7 Defining Multiple Type Parameters](#)
8. [17.8 Generics and Interfaces](#)
9. [17.9 Erasure](#)
10. [17.10 Restrictions on the Use of Generic Types](#)
11. [17.11 Common Errors to Avoid](#)

17.1 Introduction to Generics

CONCEPT: A generic class or method is one that permits you to specify the allowable types of objects that the class or method may work with. If you attempt to use the class or method with an object that is not of an allowable type, an error occurs at compile time.

In [Chapter 7](#), you were introduced to the `ArrayList` class, which is in the `java.util` package. Once you create an `ArrayList` object, you can use it as a container to hold other objects. An `ArrayList` object automatically expands in size to accommodate the number of objects that are stored in it, which is an advantage over traditional arrays. The `ArrayList` class can be used in such a way that a mixture of objects, of any type, can be simultaneously stored in it. For example, look at [Code Listing 17-1](#). This program stores a `String` object, a `Double` object, and a `Rectangle` object in the same `ArrayList`.

Code Listing 17-1 (`UnsafeArrayList.java`)

```
1  import java.util.ArrayList;
2
3  /**
4   * This program demonstrates that an ArrayList can accept
5   * a mixture of object types as its elements.
6   */
7
8  public class UnsafeArrayList
9  {
10     public static void main(String[] args)
11     {
12         // Create an ArrayList object.
13         ArrayList list = new ArrayList();
14
15         // Store a variety of objects in the list.
16         list.add("Java is fun!");           // Add a String
17         list.add(new Double(2.5));          // Add a Double
18         list.add(new Rectangle(10, 12));    // Add a Rectangle
19
20         // Retrieve a reference to each object in the list.
21         // Note that the reference returned from the get
22         // method must be cast to the correct type.
23         String s = (String)list.get(0);
24         Double d = (Double)list.get(1);
25         Rectangle r = (Rectangle)list.get(2);
26     }
27 }
```

In lines 16, 17, and 18 the `ArrayList` class's `add` method is called to add a `String`, a `Double`, and a `Rectangle` to the `ArrayList`. In lines 23, 24, and 25 the `get` method is called to get a reference to each object stored in the `ArrayList`. The statement in line 23 gets a reference to element 0 (the

String), line 24 gets a reference to element 1 (the Double), and line 25 gets a reference to element 2 (the Rectangle).

The `get` method returns an `Object` reference, so it is necessary to cast the method's return value to the desired type. Line 23 casts the method's return type to a `String` reference, line 24 casts it to a `Double` reference, and line 25 casts it to a `Rectangle` reference. If we do not cast the `get` method's return value, an error will occur at compile time. This is because the compiler does not know the exact type of the object being returned; only that it is an `Object`.

Even though the cast operator prevents an error from occurring at compile time, it does not prevent an error at runtime. We can still have problems if we inadvertently cast the `get` method's return value to the wrong type. For example, look at the following code:

```
12 // Create an ArrayList object.
13 ArrayList list = new ArrayList();
14
15 // Store a variety of objects in the list.
16 list.add("Java is fun!"); // Store a String at element 0
17 list.add(new Double(2.5)); // Store a Double at element 1
18 list.add(new Rectangle(10, 12)); // Store a Rectangle at element 2
19
20 // Retrieve a reference to each object in the list.
21 Double d = (Double)list.get(0); // Error! Element 0 is a String
22 String s = (String)list.get(1);
23 Rectangle r = (Rectangle)list.get(2);
```

This code starts out just like the `main` method in [Code Listing 17-1](#). An `ArrayList` is created, and then a `String`, a `Double`, and a `Rectangle` are stored in it. An error occurs in line 21, however. This statement gets a reference to the object at element 0. The problem is that it casts the return value of the `get` method to a `Double`, when the object that is actually returned is a `String`. The error that results from this statement does not happen at compile time, however. It happens at runtime. When the code runs, a `ClassCastException` will be thrown.

Fortunately, the Java language provides a feature known as generic classes and methods, commonly referred to as generics. *Generics* allow the compiler to perform type-checking and report type incompatibilities. Rather than waiting for your finished application to throw an exception, it is far better to discover at compile time the places in your code where types have been mismatched.

A generic class or method is one that permits you to specify the allowable types of objects that the class or method may work with. If you attempt to use the class or method with an object that is not of an allowable type, an error occurs at compile time. Recall from [Chapter 7](#) that the `ArrayList` class can also be used as a generic class. It allows you to specify the type of object that you want to store in an `ArrayList`, and the Java compiler will make sure that only objects of the specified type are stored in it. For example, we can create an `ArrayList` that will hold only `String` objects with a statement such as the following:

```
ArrayList<String> myList = new ArrayList<String>();
```

This statement creates an `ArrayList` object referenced by the variable `myList`. Because we have

written the notation `<String>` immediately after the word `ArrayList`, we will only be allowed to store `String` objects in the `ArrayList`. If we try to store any other type of object in the `ArrayList`, an error will occur at compile time.



The notation `ArrayList<String>` is pronounced “`ArrayList of String`.”

Similarly, if we wish to create an `ArrayList` object that can hold only `Student` objects, we would use a statement such as the following:

```
ArrayList<Student> myStudents = new ArrayList<Student>();
```

If we write a statement that stores an object of any type other than `Student` in `myStudents`, the compiler will generate an error. In these two examples, we can provide the notation `<String>` or `<Student>` because the `ArrayList` class can be used as a generic class.

Because the compiler knows the type of every element, we do not have to use a cast operator when retrieving an element from a generic `ArrayList`. For example, the following code creates an `ArrayList` that can hold only `String` objects, stores a string in it, and then retrieves the string. Notice that the last statement does not use a cast operator.

```
ArrayList<String> myList = new ArrayList<String>();  
myList.add("Java is fun!");  
String str = myList.get(0);
```

Using the Diamond Operator for Type Inference

Java allows you to use the diamond operator (`<>`) to simplify the instantiation of a generic class. Previously, you saw the following statement that creates an `ArrayList` object:

```
ArrayList<String> list = new ArrayList<String>();
```

Notice that the data type `String` is written between the angled brackets in two locations: first in the part that declares the reference variable (on the left side of the `=` operator), and then again in the part that calls the `ArrayList` constructor (on the right side of the `=` operator). You can use the diamond operator to simplify this statement, as follows:

```
ArrayList<String> list = new ArrayList<>();
```

By using the diamond operator, you are not required to write the data type in the part of the statement that calls the `ArrayList` constructor. The diamond operator causes the compiler to infer the required

data type from the reference variable declaration. Here is another example:

```
ArrayList<Student> list = new ArrayList<>();
```

Assuming that we have a class named `Student`, this statement creates an `ArrayList` that can hold `Student` objects.

Checkpoint

1. 17.1 Why is it risky to store objects of different types in an `ArrayList`?
2. 17.2 When `ArrayList` is used as a non-generic class, why must you use the cast operator with the `get` method?
3. 17.3 Suppose we use the following statement to instantiate the `ArrayList` class. What type of objects will we be allowed to store in the `ArrayList`?

```
ArrayList<Rectangle> myList = new ArrayList<>();
```

4. 17.4 Assume we have used the statement shown in [Checkpoint 17.3](#) to instantiate the `ArrayList` class. If we attempt to store an object that is not of the allowable type in the `ArrayList`, will we see an error message at compile time or at runtime?

17.2 Writing a Generic Class

CONCEPT: In Java you can write your own generic classes that accept data type arguments.

Not only does Java provide several generic classes for you to use, such as `ArrayList`, but it also allows you to write your own generic classes. [Code Listing 17-2](#) shows the `Point` class, which is a simple example of a generic class. (This version of the class is stored in the source code folder [Chapter 17\Point Class Version 1.](#)) An instance of the class holds two values: an *X* coordinate and a *Y* coordinate.

Code Listing 17-2 (Point.java)

```
1  /**
2     The Point class holds X and Y coordinates. The data type
3     of the coordinates is generic.
4  */
5
6  public class Point<T>
7  {
8      private T xCoordinate; // The X coordinate
9      private T yCoordinate; // The Y coordinate
10
11     /**
12         Constructor
13         @param x The X coordinate.
14         @param y The Y coordinate.
15     */
16
17     public Point(T x, T y)
18     {
19         xCoordinate = x;
20         yCoordinate = y;
21     }
22
23     /**
24         The setX method sets the X coordinate.
25         @param x The value for the X coordinate.
26     */
27
28     public void setX(T x)
29     {
30         xCoordinate = x;
31     }
32
33     /**
34         The setY method sets the Y coordinate.
35         @param y The value for the Y coordinate.
36     */
37
```

```

38     public void setY(T y)
39     {
40         yCoordinate = y;
41     }
42
43     /**
44         The getX method returns the X coordinate.
45         @return The value of the X coordinate.
46     */
47
48     public T getX()
49     {
50         return xCoordinate;
51     }
52
53     /**
54         The getY method returns the Y coordinate.
55         @return The value of the Y coordinate.
56     */
57
58     public T getY()
59     {
60         return yCoordinate;
61     }
62 }

```

Let's take a closer look at the class header in line 6:

```
public class Point<T>
```

After the name of the class, `T` appears inside angled brackets. In this class declaration, `T` is a type parameter. A *type parameter* is an identifier that represents an actual type. When you instantiate the class, you use an actual type as an argument for the type parameter. You do this by writing the type name inside angled brackets, immediately following the name of the class. For example, consider the following code:

```

Integer intX = new Integer(10);
Integer intY = new Integer(20);
Point<Integer> myPoint = new Point<>(intX, intY);


```

Recall that `Integer` is a wrapper class. In this statement we are passing `Integer` as a type argument to the `T` parameter. This is illustrated in [Figure 17-1](#).

```

Point<Integer> myPoint = new Point<>(intX, intY);

```



```

public class Point<T>

```

Inside the `Point` class,
`T` becomes `Integer`.

Figure 17-1 Passing a type argument to a

type parameter

[Figure 17-1 Full Alternative Text](#)

For this instance of the `Point` class, the type parameter `T` will become `Integer`. There are several places in the `Point` class where `T` is used, as shown in [Figure 17-2](#). In each place, `T` will become the type that was passed to it as an argument.




Figure 17-2 The type parameter `T` in the `Point` class

[Figure 17-2 Full Alternative Text](#)

For example, look at the following declarations of the `xCoordinate` and `yCoordinate` fields in lines 8 and 9:

```
private T xCoordinate;           // The X coordinate
private T yCoordinate;           // The Y coordinate
```

If we pass `Integer` as a type argument to `T`, these declarations will become as follows:

```
private Integer xCoordinate;    // The X coordinate
private Integer yCoordinate;    // The Y coordinate
```

The `T` type parameter is also used in the following method headers:

```
public Point(T x, T y)
public void setX(T x)
public void setY(T y)
public T getX()
public T getY()
```

If we pass `Integer` as a type parameter to `T`, these method headers will become as follows:

```
public Point(Integer x, Integer y)
public void setX(Integer x)
public void setY(Integer y)
public Integer getX()
public Integer getY()
```

The program in [Code Listing 17-3](#) tests the `Point` class by creating two instances of it. In the first instance, the `Double` type is passed to the type parameter `T`, and in the second instance the `Integer` type is passed to `T`. (This program is also stored in the source code folder [Chapter 17\Point Class Version 1.](#))

Code Listing 17-3 (TestPoint.java)

```
1 /**
2     This program demonstrates the Point class.
3 */
4
5 public class TestPoint
6 {
7     public static void main(String[] args)
8     {
9         // Create two Double objects to use as coordinates.
10        Double dblX = new Double(1.5);
11        Double dblY = new Double(2.5);
12
13        // Create a Point object that can hold Doubles.
14        Point<Double> dPoint = new Point<>(dblX, dblY);
15
16        // Create two Integer objects to use as coordinates.
17        Integer intX = new Integer(10);
18        Integer intY = new Integer(20);
19
20        // Create a Point object that can hold Integers.
21        Point<Integer> iPoint = new Point<>(intX, intY);
22
23        // Display the Double values stored in dPoint.
24        System.out.println("Here are the values in dPoint.");
```



```

25     System.out.println("X Coordinate: " + dPoint.getX());
26     System.out.println("Y Coordinate: " + dPoint.getY());
27     System.out.println();
28
29     // Display the Integer values stored in iPoint.
30     System.out.println("Here are the values in iPoint.");
31     System.out.println("X Coordinate: " + iPoint.getX());
32     System.out.println("Y Coordinate: " + iPoint.getY());
33 }
34 }

```

Program Output

```

Here are the values in dPoint.
X Coordinate: 1.5
Y Coordinate: 2.5
Here are the values in iPoint.
X Coordinate: 10
Y Coordinate: 20

```

When the following statement is executed in line 14, the `Double` type is passed as an argument to the type parameter `T`:

```
Point<Double> dPoint = new Point<>(dblX, dblY);
```

For this particular instance of the `Point` class, the `T` type parameter is replaced with `Double`. The method header for this particular object's constructor becomes the following:

```
public Point(Double x, Double y)
```

Notice that in line 14 references to two `Double` objects, `dblX` and `dblY`, are passed as arguments to the constructor. If we had passed anything other than `Double` references, the compiler would have generated an error.

Another instance of the class is created with the following statement in line 21:

```
Point<Integer> iPoint = new Point<>(intX, intY);
```

For this particular instance of the `Point` class, the `T` type parameter is replaced with `Integer`. The method header for this particular object's constructor becomes the following:

```
public Point(Integer x, Integer y)
```

In this statement we have passed two `Integer` references, `intX` and `intY`, as arguments to the constructor. If we had passed anything other than `Integer` references, an error would have occurred at compile time.

Only Reference Types Can Be Passed to Type

Parameters

When you create an instance of a generic class, you can only pass a reference type to the class's type parameter. This is the reason that in [Code Listing 17-3](#) we passed the `Double` and `Integer` types to the `Point` class instead of the primitive `double` and `int` types. The only way to use a generic type to hold a primitive value is to wrap the primitive value up in a wrapper class, as we demonstrated in [Code Listing 17-3](#).

Autoboxing and Unboxing

Although we can only pass reference types to a generic class's type parameter, we can use Java's autoboxing and unboxing feature to simplify our code if we need to store primitive values in a generic type. We briefly introduced autoboxing and unboxing in [Chapter 9](#). Recall that *autoboxing* is Java's process of automatically “boxing up” a primitive value inside an object of an appropriate wrapper class. For example, the statement in line 14 of [Code Listing 17-3](#) could have been written as follows:

```
Point<Double> dPoint = new Point<>(1.5, 2.5);
```

Because the `Point` constructor accepts two `Double` arguments, the literal values 1.5 and 2.5 are autoboxed into two objects of the `Double` class. References to those objects are then passed as arguments to the constructor.

Unboxing is the opposite of boxing. It is the process of automatically converting a wrapper class object to a primitive type, when needed. For example, look at the following code:

```
// Use autoboxing to pass doubles to the constructor.
Point<Double> dPoint = new Point<>(1.5, 2.5);
// Use unboxing to retrieve the X and Y coordinates
// and assign them to double variables.
double x = dPoint.getX();
double y = dPoint.getY();
```

In the last two statements, the `getX` and `getY` methods are called, which will return references to `Double` objects. Because the methods' return values are being assigned to `double` variables, however, Java will unbox the `Double` objects and store their values in the `double` variables. After this code executes the variables `x` and `y` will hold the values 1.5 and 2.5.

Instantiating a Generic Class without Specifying a Type Argument

Although it is not recommended, it is possible to create an instance of a generic class without providing a type argument. When you do this, you are using the class as a *raw type*. For example, look

at the program in [Code Listing 17-4](#). (This program is stored in the source code folder [Chapter 17\Point Class Version 1](#).)

Code Listing 17-4 (TestPoint2.java)

```
1  /**
2   * This demonstrates the Point class without a type argument.
3   */
4
5  public class TestPoint2
6  {
7      public static void main(String[] args)
8      {
9          // Create an Integer and a Double
10         // to use as coordinates.
11         Integer x = new Integer(1);
12         Double y = new Double(7.5);
13
14         // Create a Point object with an
15         // Integer x coordinate and a Double
16         // y coordinate.
17         Point myPoint = new Point(x, y);
18
19         // Display the object's X and Y
20         // coordinate values.
21         System.out.println("X Coordinate: " +
22                             myPoint.getX());
23         System.out.println("Y Coordinate: " +
24                             myPoint.getY());
25     }
26 }
```

Program Output

```
X Coordinate: 1
Y Coordinate: 7.5
```

An instance of the `Point` class is created with the following statement in line 17:

```
Point myPoint = new Point(x, y);
```

Notice that there is no type argument provided in angled brackets. When a type argument is not provided for an instance of a generic class, Java will provide `Object` as the type argument by default. This means that for the `myPoint` object, the declarations for the `xCoordinate` and `yCoordinate` fields will become the following:

```
private Object xCoordinate;    // The X coordinate
private Object yCoordinate;    // The Y coordinate
```

Likewise, each occurrence of `T` in the class methods will become `Object`.

NOTE:

If you compile the `TestPoint2.java` program shown in [Code Listing 17-4](#) you will see a warning such as the following:

```
TestPoint2.java uses unchecked or unsafe operations.  
Note: Recompile with -Xlint:unchecked for details.
```

The Java compiler issues this warning when we compile code that instantiates a generic class without specifying a type argument.

It is not recommended that you instantiate a generic class without providing a type argument, because you give up the benefits of type-safety that you would have otherwise. For example, in line 17 of [Code Listing 17-4](#), an `Integer` object is passed to the `Point` class constructor as the *X* coordinate, and a `Double` object is passed as the *Y* coordinate. This is legal because the constructor's parameter variables are of the `Object` type, and we can polymorphically pass references of any type to them.

It also means that the class's `xCoordinate` and `yCoordinate` fields can reference objects of any type. This opens up the door for runtime errors if we attempt to perform an unsupported operation on one of the class's fields.

Another reason that we should avoid instantiating a generic class without providing a type argument is that we have to use a cast operator to assign an `Object` reference to a variable of any other type. The following code illustrates this:

```
// Create an Integer and a Double.  
Integer x = new Integer(1);  
Double y = new Double(7.5);  
  
// Create a Point object with no type argument.  
Point myPoint = new Point(x, y);  
  
// Retrieve the X and Y coordinates.  
Integer xx = (Integer) myPoint.getX();  
Double yy = (Double) myPoint.getY();
```

The last two statements shown here use the `myPoint` object to call the `getX` and `getY` methods. Because these methods return an `Object` reference, we must cast the reference to the proper type in order to assign it to a variable. This is not necessary when we instantiate the `Point` class with a type argument.

Commonly Used Type Parameter Names

The generic `Point` class that we examined defines a type parameter named `T`. There's nothing magic about the name `T`. When you are defining a generic type parameter, you can use any legal identifier you wish. By convention, however, type parameters are usually single letters written in uppercase. This helps distinguish the name of a type parameter from the name of a class. Also, there are some names that are commonly used, depending on the purpose of the type parameter. [Table 17-1](#) lists the commonly used parameter names and briefly describes their meaning.

Table 17-1 Common Type Parameter Names

Name	Usual Meaning
<code>T</code>	<code>T</code> is commonly used for a general type.
<code>S</code>	<code>S</code> is commonly used for a general type when <code>T</code> has already been used.
<code>E</code>	<code>E</code> is commonly used to represent the type of an element, such as an element in an <code>ArrayList</code> or other collection.
<code>K</code>	<code>K</code> is commonly used to represent the type of a key in a class that keeps key/value pairs.
<code>V</code>	<code>V</code> is commonly used to represent the type of a value in a class that keeps key/value pairs.

As we progress through this chapter you will see several of these type parameter names being used.



Checkpoint

1. 17.5 What is a type parameter?
2. 17.6 When you create an instance of a generic class, can you pass reference types, primitive types, or either to the class's type parameter?

3. 17.7 Is it possible to create an instance of a generic class without providing a type argument? If so, is this recommended?
4. 17.8 Is it required that type parameters be single letters written in uppercase?

17.3 Passing Objects of a Generic Class to a Method

CONCEPT: When you create instances of a generic class, you must be able to pass those instances as arguments to methods. Method parameter variables can be declared to be of a generic type.

Suppose you want to write a method that accepts an instance of a generic class as an argument. For example, the following `printPoint` method accepts an instance of `Point<Integer>` as an argument and displays the *X* and *Y* coordinate values in that object.

```
public static void printPoint(Point<Integer> point)
{
    System.out.println("X Coordinate: " + point.getX());
    System.out.println("Y Coordinate: " + point.getY());
}
```

In this method, the parameter `point` is a reference to a `Point<Integer>` object. We can call the method and pass a reference to any `Point<Integer>` object, as shown here:

```
Point<Integer> iPoint = new Point<>(7, 12);
printPoint(iPoint);
```

A problem arises, however, if we want to pass an instance of `Point<Double>` to the method. The method's parameter is a reference to `Point<Integer>`, so only `Point<Integer>` objects can be passed to it.

If the `printPoint` method is intended to be a general purpose method that can display the *X* and *Y* coordinates of any `Point` object, this limitation will be too restrictive. It seems reasonable that any numeric wrapper class could be used as a valid type argument when creating an instance of the `Point` class. For example, we should expect that instances of `Point<Integer>`, `Point<Double>`, `Point<Long>`, `Point<Float>`, and `Point<Short>` are all valid. In that case, we would want the `printPoint` method to be able to accept instances of any of these types as an argument.

At this point you know something about inheritance and polymorphism, and you might be thinking along these lines: "I just checked the Java API documentation and found that all of the numeric wrapper classes extend the `Number` class, which is in the `java.lang` package. Why not use `Number` as the type argument in the method's parameter? That way we can use polymorphism to accept any `Point` object whose type argument extends the `Number` class."

In other words, why not write the `printPoint` method like this?

```
// Will this do what we want?
public static void printPoint(Point<Number> point)
{
    System.out.println("X Coordinate: " + point.getX());
}
```

```
System.out.println("Y Coordinate: " + point.getY());
```

```
}
```

Unfortunately, this will not give us the results we want. The Java compiler will allow us to pass *only* `Point<Number>` objects to this method. If we try to pass a `Point<Integer>` or `Point<Double>` object, we will get an error at compile time. One way to solve our problem is to use the `?` type wildcard with the parameter, as shown here:

```
public static void printPoint(Point<?> point)
{
    System.out.println("X Coordinate: " + point.getX());
    System.out.println("Y Coordinate: " + point.getY());
}
```

Notice that the parameter's type is `Point<?>`. The `?` character that appears inside the angled brackets is a *type wildcard*, and it indicates that any type argument can be used in its place. Now we can pass any `Point` object as an argument to the method, regardless of its type argument. The following statements will compile and execute without errors:

```
Point<Integer> iPoint = new Point<>(1, 2);
Point<Double> dPoint = new Point<>(1.5, 2.5);
printPoint(iPoint);
printPoint(dPoint);
```

Constraining a Type Parameter

Recall that the header for our original version of the `printPoint` method header looked like the following:

```
public static void printPoint(Point<Integer> point)
```

Because the parameter's type, `Point<Integer>`, was too restrictive, we changed the method header to the following:

```
public static void printPoint(Point<?> point)
```

The method now allows us to pass any `Point` object as an argument. However, this leads to a different problem: The parameter's new type, `Point<?>`, might not be restrictive enough. Suppose we only want to accept `Point` objects whose type argument is a numeric wrapper class. In other words, we only want to accept `Point` objects whose type argument is a subclass of `Number`. This can be accomplished by modifying the method once more, as shown here:

```
public static void printPoint(Point<? extends Number> point)
{
    System.out.println("X Coordinate: " + point.getX());
    System.out.println("Y Coordinate: " + point.getY());
}
```

In this version of the method, the parameter's type is `Point<? extends Number>`. This means that

the `Point` object's type argument may be `Number`, or any type that extends `Number`. The following statements will compile and execute without errors:

```
Point<Integer> iPoint = new Point<>(1, 2);
Point<Double> dPoint = new Point<>(1.5, 2.5);
printPoint(iPoint);
printPoint(dPoint);
```

However, the following code will not compile because the `String` class is not a subclass of `Number`:

```
Point<String> sPoint = new Point<>("1", "2");
printPoint(sPoint);    // Error!
```

[Code Listing 17-5](#) shows the `TestPoint3.java` program, which demonstrates this version of the `printPoint` method. (This program is stored in the source code folder [Chapter 17\Point Class Version 1.](#))

Code Listing 17-5 (TestPoint3.java)

```
1  /**
2      This program demonstrates the printPoint method.
3  */
4
5  public class TestPoint3
6  {
7      public static void main(String[] args)
8      {
9          // Create various Point objects.
10         Point<Integer> iPoint = new Point<>(1, 2);
11         Point<Double> dPoint = new Point<>(1.5, 2.5);
12         Point<Long> lPoint = new Point<>(10L, 20L);
13         Point<Float> fPoint = new Point<>(7.9f, 9.9f);
14
15         // Display each object's coordinates.
16         System.out.println("iPoint:");
17         printPoint(iPoint);
18
19         System.out.println("\ndPoint:");
20         printPoint(dPoint);
21
22         System.out.println("\nlPoint:");
23         printPoint(lPoint);
24
25         System.out.println("\nfPoint:");
26         printPoint(fPoint);
27     }
28
29     /**
30         The printPoint method prints a Point object's
31         X and Y coordinates.
32         @param point A Point object.
33     */
34 }
```

```

35     public static void printPoint(Point<? extends Number> point)
36     {
37         System.out.println("X Coordinate: " + point.getX());
38         System.out.println("Y Coordinate: " + point.getY());
39     }
40 }

```

Program Output

```

iPoint:
Coordinate: 1
Coordinate: 2

dPoint:
Coordinate: 1.5
Coordinate: 2.5

lPoint:
Coordinate: 10
Coordinate: 20

fPoint:
Coordinate: 7.9
Coordinate: 9.9

```



NOTE:

In line 12 the values being passed as arguments to the `Point` class constructor are `10L` and `20L`. The “L” suffix tells the compiler to treat these literal values as long integers. Without the “L” suffix, these literals would be treated as `int` values, and an error would occur. Similarly, in line 13 the values `7.9f` and `9.9f` are being passed as arguments to the constructor. The “f” suffix tells the compiler to treat these literals as `float` values. Without the “f” suffix, these literals would be treated as `double` values and an error would occur.

Defining a Type Parameter in a Method Header

In the previous example you saw how the notation `<? extends Number>` means “any type that is `Number` or a subclass of `Number`.” As an alternative to using the wildcard in this notation, you can actually define a type parameter in a method header. For example, the `printPoint` method could have been written as follows:

```

public static <T extends Number> void printPoint(Point<T> point)
{

```

```

System.out.println("X Coordinate: " + point.getX());
System.out.println("Y Coordinate: " + point.getY());
}

```

Notice that between the words `static` and `void`, the notation `<T extends Number>` appears. This defines a type parameter named `T`, and specifies that `T` can accept any type that is `Number` or a subclass of `Number`. Then notice that the type of the method's parameter is `Point<T>`. This means that the method can accept any `Point` object whose type argument is `Number` or a subclass of `Number`. This is illustrated in [Figure 17-3](#).




Figure 17-3 Type parameter defined in a method header

[Figure 17-3 Full Alternative Text](#)

Using this alternative syntax can simplify methods that accept multiple arguments of generic types. For example, look at the following method header:

```

public static void doSomething(Point<? extends Number> arg1,
                               Point<? extends Number> arg2,
                               Point<? extends Number> arg3)

```

This method accepts three arguments, each of which is a `Point<? extends Number>` object. We had to write `Point<? extends Number>` as the type for each parameter variable. We can simplify the notation by defining a type parameter, as shown here:

```

public static <T extends Number> void doSomething(Point<T> arg1,
                                                  Point<T> arg2,
                                                  Point<T> arg3)

```

The `extends` Key Word Constrains a Type to

an Upper Bound

When you use the `extends` key word with a generic type parameter, as we have shown in this section, it is said that you are constraining the type parameter to an *upper bound*. For example, in the notation `<T extends Number>` we are constraining `T` to the upper bound `Number`. This means that `T` can be any type that is below `Number` in the class hierarchy (or `Number` itself), but it cannot be a type that is above `Number` in the class hierarchy.

The `super` Key Word Constrains a Type to a Lower Bound

In addition to the `extends` key word, you can use the `super` key word to restrict a type parameter. Here is an example of how it might be used in a method header:

```
public static void doSomething(Point<? super Integer> arg)
```

In this method, the `arg` parameter's type is `Point<? super Integer>`. This means that the `Point` object's type argument may be `Integer`, or any superclass of `Integer`. Because the type can be any class that appears above `Integer` in the class hierarchy, it is said that we are constraining the type to a *lower bound*. The `Integer` class has only two superclasses: `Number` and `Object`. So, this method will only accept as arguments objects that are `Point<Integer>`, `Point<Number>`, or `Point<Object>`.



You can use either `extends` to constrain a type to an upper bound, or `super` to constrain a type to a lower bound, but not both.

Checkpoint

1. 17.9 In generic type notation, what does the `?` character (the type wildcard) indicate?
2. 17.10 In generic type notation, what does `<? extends Number>` mean?
3. 17.11 In generic type notation, what does `<? super Integer>` mean?
4. 17.12 In generic type notation, does the `extends` key word establish an upper boundary or a lower boundary?

5. 17.13 In generic type notation, does the `super` key word establish an upper boundary or a lower boundary?
6. 17.14 In generic type notation, can you use both `extends` and `super` to constrain a type to a set of boundaries?

17.4 Writing Generic Methods

CONCEPT: In addition to classes, methods can also be generic.

Previously, you saw how a method can accept an object of a generic class as an argument. Methods themselves can also be generic. This means that they can have their own type parameters, and can use those type parameters to represent the types of arguments, the types of local variables, and their return type. [Code Listing 17-6](#) shows a class with a generic method.

Code Listing 17-6 (GenericMethodDemo.java)

```
1  /**
2   * This program demonstrates a simple generic method.
3   */
4
5  public class GenericMethodDemo
6  {
7      public static void main(String[] args)
8      {
9          String[] names = { "Alfonso", "Beatrice", "Celine" };
10         displayArray(names);
11     }
12
13     /**
14      * The displayArray method displays each element
15      * in an array.
16      * @param array The array to display.
17      */
18
19     public static <E> void displayArray(E[] array)
20     {
21         for (E element : array)
22             System.out.println(element);
23     }
24 }
```

Program Output

```
Alfonso
Beatrice
Celine
```

Line 19 shows the header for the `displayArray` method. Notice that the method defines a type parameter named `E`. The method has a parameter variable, `array`, which is a reference to an array of `E` objects. The `E` type parameter is also used in line 21, in the enhanced `for` loop. Notice in line 10 that the `displayArray` method is called like a regular method. Even though it is a generic method, no

type argument is passed to it. When you call a generic method, the compiler determines which type to use from the context in which you are using the method. For example, in line 10 we are passing `names` as an argument to the method. Because `names` is a `String` array, the compiler uses the `String` type for `E` in the method.

You can constrain a type parameter in a generic method. For example, suppose the `displayArray` method had been written as follows:

```
public static <E extends Number> void displayArray(E[] array)
{
    for (E element : array)
        System.out.println(element);
}
```

In this version of the method, we have constrained the type parameter `E` to any type that is `Number` or a subclass of `Number`. When calling this version of the method, we can only pass an array of `Number` objects, or of objects of a subclass of `Number`.

Checkpoint

1. 17.15 When an instance of a generic class is created, a type argument is explicitly passed to the class's type parameter. When a generic method is called, how is a type argument provided for the method's type parameter?

17.5 Constraining a Type Parameter in a Generic Class

CONCEPT: You can constrain the type parameter in a generic class so that only certain type arguments can be used to create an instance of the class.

Previously you saw how we used generic notation in the `printPoint` method header so the method can accept an instance of the generic `Point` class as an argument. You also saw how we used generic notation to constrain the method so that it will only accept `Point` objects whose type argument is `Number` or a subclass of `Number`.

Suppose we wish to constrain the `Point` class itself so that only certain type arguments can be used to create an instance of the class. For example, suppose we want to allow instances of the `Point` class to be created using only the numeric wrapper classes as type arguments. That would prevent a programmer from using classes such as `String` as the type argument when instantiating the `Point` class. We can impose such a restriction by using the `extends` key word to constrain the class's type parameter.

[Code Listing 17-7](#) shows a modified version of the `Point` class, which uses the `extends` key word in its generic type notation. (This version of the class is stored in the source code folder [Chapter 17\Point Class Version 2.](#))

Code Listing 17-7 (Point.java)

```
1  /**
2   * The Point class holds X and Y coordinates. The data type
3   * of the coordinates is generic, but restricted to
4   * Number or a subclass of Number.
5   */
6
7  public class Point<T extends Number>
8  {
9      private T xCoordinate; // The X coordinate
10     private T yCoordinate; // The Y coordinate
11
12     /**
13      * Constructor
14      * @param x The X coordinate.
15      * @param y The Y coordinate.
16      */
17
18     public Point(T x, T y)
19     {
20         xCoordinate = x;
21         yCoordinate = y;
```



```

22     }
23
24     /**
25      * The setX method sets the X coordinate.
26      * @param x The value for the X coordinate.
27      */
28
29     public void setX(T x)
30     {
31         xCoordinate = x;
32     }
33
34     /**
35      * The setY method sets the Y coordinate.
36      * @param y The value for the Y coordinate.
37      */
38
39     public void setY(T y)
40     {
41         yCoordinate = y;
42     }
43
44     /**
45      * The getX method returns the X coordinate.
46      * @return The value of the X coordinate.
47      */
48
49     public T getX()
50     {
51         return xCoordinate;
52     }
53
54     /**
55      * The getY method returns the Y coordinate.
56      * @return The value of the Y coordinate.
57      */
58
59     public T getY()
60     {
61         return yCoordinate;
62     }
63 }

```

The only difference between this `Point` class and the one we studied earlier is the following type parameter notation in line 7:

```
public class Point<T extends Number>
```

The notation `<T extends Number>` defines a type parameter `T`, which is constrained with `Number` as its upper bound. In other words, only `Number` or a subclass of `Number` may be passed as a type argument to this parameter. The following statements will compile without error because each of the type arguments being used is a subclass of `Number`:

```

Point<Integer> iPoint = new Point<>(1, 2);
Point<Double> dPoint = new Point<>(1.5, 2.5);

```

```
Point<Long> lPoint = new Point<>(10L, 20L);
Point<Float> fPoint = new Point<>(7.9f, 9.9f);
```

These statements, however, will cause an error at compile time because the type arguments `String` and `Character` are not subclasses of `Number`:

```
// These statements will cause an error.
Point<String> iPoint = new Point<>("1", "2");
Point<Character> dPoint = new Point<>('1', '2');
```

You can also use the `super` key word in a generic class, to constrain a type parameter to a lower bound. For example, suppose we had constrained the `Point` class's type parameter as follows:

```
public class Point<T super Double>
```

This specifies that the type argument must be `Double`, or any superclass of `Double`. Otherwise, an error will occur at compile time.

Checkpoint

1. 17.16 Why would you want to constrain the type parameter in a generic class?
2. 17.17 Can you constrain a generic class's type parameter with an upper bound? With a lower bound?

17.6 Inheritance and Generic Classes

CONCEPT: A generic class can be a superclass, and it can extend other classes.



Inheritance and Generic classes

Inheritance can be used with generic classes. For example, [Code Listing 17-8](#) shows the `Point3D` class, which extends the `Point` class. This class inherits the `Point` class's `xCoordinate` and `yCoordinate` fields, and has its own field for a `Z` coordinate, named `zCoordinate`. (This version of the class is stored in the source code folder [Chapter 17\Point Class Version 2.](#))

Code Listing 17-8 (Point3D.java)

```
1  /**
2   The Point3D class holds a Z coordinate. The data type
3   of the coordinate is generic.
4  */
5
6  public class Point3D<T extends Number> extends Point<T>
7  {
8      private T zCoordinate; // The z coordinate
9
10     /**
11      Constructor
12      @param x The X coordinate.
13      @param y The Y coordinate.
14      @param z The Z coordinate.
15     */
16
17     public Point3D(T x, T y, T z)
18     {
19         // Call the Point class constructor.
20         super(x, y);
21
22         // Assign the Z coordinate.
23         zCoordinate = z;
24     }
25
26     /**
27      The setZ method sets the Z coordinate.
28      @param z The value for the Z coordinate.
29     */
30
31     public void setZ(T z)
32     {
```

```

33         zCoordinate = z;
34     }
35
36     /**
37         The getZ method returns the Z coordinate.
38         @return The value of the Z coordinate.
39     */
40
41     public T getZ()
42     {
43         return zCoordinate;
44     }
45 }

```

Take a closer look at the class header in line 6:

```
public class Point3D<T extends Number> extends Point<T>
```

The first part of this statement reads as follows:

```
public class Point3D<T extends Number>
```

This indicates that the class we are defining is `Point3D`, and the type parameter is `T`. The notation `<T extends Number>` is used to define the type parameter, so any type argument passed to `T` must be `Number` or a subclass of `Number`. The next part of the statement is as follows:

```
extends Point<T>
```

This indicates that the class we are defining (`Point3D`) extends the `Point` class, with `T` passed as a type argument to the `Point` class. [Figure 17-4](#) illustrates the parts of this code.




Figure 17-4 Generic subclass header

The program in [Code Listing 17-9](#) demonstrates the class. (This program is stored in the source code folder [Chapter 17\Point Class Version 2.](#))

Code Listing 17-9 (TestPoint3D.java)

```
1  /**
2   * This program demonstrates the Point3D class.
3   */
4
5  public class TestPoint3D
6  {
7      public static void main(String[] args)
8      {
9          // Create a Point3D object that can hold Integers.
10         Point3D<Integer> point = new Point3D<>(10, 20, 30);
11
12         // Display the coordinates stored in point.
13         System.out.println("Here are the values in point.");
14         System.out.println("X Coordinate: " + point.getX());
15         System.out.println("Y Coordinate: " + point.getY());
16         System.out.println("Z Coordinate: " + point.getZ());
17     }
18 }
```

Program Output

```
Here are the values in point.
X Coordinate: 10
Y Coordinate: 20
Z Coordinate: 30
```

The “is-a” relationship is in effect between the `Point3D` class and the `Point` class. A `Point3D` object *is a* `Point` object. So, we can assign a `Point3D` object to a `Point` reference variable, or pass a `Point3D` object to a method that accepts `Point` objects. The program in [Code Listing 17-10](#) demonstrates this. (This program is also stored in the source code folder [Chapter 17\Point Class Version 2.](#))

Code Listing 17-10 (PassPoint3D.java)

```
1  /**
2   * This program passes a Point3D object to a method
3   * with a Point parameter variable.
4   */
5
6  public class PassPoint3D
7  {
8      public static void main(String[] args)
9      {
```

```

10      // Create a Point3D object that can hold Integers.
11      Point3D<Integer> point = new Point3D<>(10, 20, 30);
12
13      // Display the X and Ycoordinates stored in point.
14      System.out.println("Here are the X and Y coordinates.");
15      printPoint(point);
16  }
17
18  /**
19       The printPoint method prints a Point object's
20       X and Y coordinates.
21       @param point A Point object.
22  */
23
24  public static void printPoint(Point<?> point)
25  {
26      System.out.println("X Coordinate: " + point.getX());
27      System.out.println("Y Coordinate: " + point.getY());
28  }
29  }

```

Program Output

```

Here are the X and Y coordinates.
X Coordinate: 10
Y Coordinate: 20

```

The `Point3D` class is an example of a generic class that extends another generic class. Both generic and non-generic classes may be used together in an inheritance hierarchy, in any of the following ways:

- A generic class may extend another generic class
- A generic class may extend a non-generic class
- A non-generic class may extend a generic class



Checkpoint

1. 17.18 Can a generic class be a superclass? Can it be a subclass?
2. 17.19 A generic class extends another generic class. Can you assign a reference to an instance of the subclass to a superclass variable?
3. 17.20 Can you mix generic and non-generic classes in a class hierarchy?

17.7 Defining Multiple Type Parameters

CONCEPT: A generic class or method can have multiple type parameters, allowing it to accept multiple types as arguments.

It is possible to define multiple type parameters in a generic class or method. Different type arguments can then be passed to each type parameter. Here is an example of a class that defines two type parameters:

```
public class MyClass<T, S>
{
    class code here...
}
```

In this example the type parameters are `T` and `S`. Inside the angled brackets the type parameters are listed, separated by commas. You can also apply constraints to the type parameters. Here is an example:

```
public class MyClass<T extends Number, S extends Date>
{
    class code here...
}
```

In this example the type parameters are `T` and `S`. Any type that is passed to `T` must be `Number` or a subclass of `Number`. Any type that is passed to `S` must be `Date` or a subclass of `Date`.

[Code Listing 17-11](#) shows the `Pair` class, which uses two type parameters. The `Pair` class is designed to hold a pair of items. Each of the items can be of a different type.

Code Listing 17-11 (`Pair.java`)

```
1  /**
2     The Pair class demonstrates a generic class
3     with two type parameters.
4  */
5
6  public class Pair<T, S>
7  {
8      private T first;           // The first item
9      private S second;         // The second item
10
11     /**
12         Constructor
13         @param firstArg  Assigned to the first item.
14         @param secondArg Assigned to the second item.
15     */
16 }
```

```

17     public Pair(T firstArg, S secondArg)
18     {
19         first = firstArg;
20         second = secondArg;
21     }
22
23     /**
24         getFirst method
25         @return The first item in the pair.
26     */
27
28     public T getFirst()
29     {
30         return first;
31     }
32
33     /**
34         getSecond method
35         @return The second item in the pair.
36     */
37
38     public S getSecond()
39     {
40         return second;
41     }
42 }

```

Take a closer look at the class header in line 6:

```
public class Pair<T, S>
```

The notation `<T, S>` defines two type parameters, `T` and `S`. Notice that the first field is of the `T` type and the second field is of the `S` type. When we instantiate the class we must pass a type argument for each of these parameters. [Code Listing 17-12](#) demonstrates the class.

Code Listing 17-12 (PairTest.java)

```

1  /**
2   * This program demonstrates the Pair class which
3   * has two type parameters.
4   */
5
6  public class PairTest
7  {
8      public static void main(String[] args)
9      {
10         // Create an Integer to hold an ID number.
11         Integer idNumber = new Integer(475);
12
13         // Create a String to hold a name.
14         String name = "Smith, Sally";
15
16         // Create a Pair object to hold the ID

```



```

17      // number and the name.
18      Pair<Integer, String> myPair =
19          new Pair<>(idNumber, name);
20
21      // Display the pair of items.
22      System.out.println("ID Number: " +
23                          myPair.getFirst());
24      System.out.println("Name: " +
25                          myPair.getSecond());
26  }
27  }

```

Program Output

```

ID Number: 475
Name: Smith, Sally

```

Line 11 creates an `Integer` object initialized with the value 475, referenced by the variable `idNumber`. Line 14 creates a `String` object referenced by the `name` variable. Lines 18 and 19 create a `Pair` object using the generic notation `<Integer, String>`. This passes `Integer` to the `T` parameter and `String` to the `S` parameter. This means that the first field will be an `Integer` and the second field will be a `String`.

The statement in lines 18 and 19 also pass `idNumber` and `name` to the `Pair` constructor, storing those items in the `myPair` object.

17.8 Generics and Interfaces

CONCEPT: Interfaces, as well as classes, may be generic.

So far, all of the generic types that we have examined have been classes. Interfaces can be generic too. For example, the Java API has an interface named `Comparable`, which is defined as follows:

```
public interface Comparable<T>
{
    int compareTo(T o);
}
```

This interface specifies a method named `compareTo`, which is used to compare the calling object with another object that is passed as an argument into the `o` parameter. Notice that the interface defines a type parameter, `T`. In the `compareTo` method header, `T` is the type of the `o` parameter.

Many of the classes in the Java API implement this interface, including `String` and the wrapper classes. You might recall from our discussion of the `String` class in [Chapter 3](#) how the `compareTo` method works:

- If the calling object is less than the object passed to `o`, the method returns a negative number.
- If the calling object is equal to the object passed to `o`, the method returns 0.
- If the calling object is greater than the object passed to `o`, the method returns a positive number.

Suppose we are writing a class named `Tree`, and we want the `Tree` class to implement the `Comparable` interface. We would write the class in the following manner:

```
public class Tree implements Comparable<Tree>
{
    public int compareTo(Tree o)
    {
        Method code here...
    }
    Other class code here...
}
```

In the class header we are passing `Tree` as the type argument to the `Comparable` interface. So, the header for the `compareTo` method in the `Tree` class must be written as follows:

```
int compareTo(Tree o)
```

Any `Tree` object that calls the `compareTo` method will have to pass another `Tree` object as an argument to the method.

Constraining a Type Parameter to a Type That Implements an Interface

It is common for a method to have a parameter variable that is an interface reference rather than a class reference. For example, suppose we wish to write a generic method named `greatest`, which will compare two objects and determine which one is the “greater” of the two. Of course, the notion of “greater” depends entirely on the objects that are being compared. If we are comparing two `LuxuryCar` objects, the one that is the most expensive might be considered the greater one. However, if we are comparing two `RaceCar` objects, the one that has the highest top speed might be considered the greater one. Any objects that are passed as arguments to our method will have to support a standard method of comparison so we can determine which is the greater.

The `Comparable` interface helps in a situation like this. For example, let’s assume that the `LuxuryCar` class implements the `Comparable` interface. The class header would be written as follows:

```
public class LuxuryCar implements Comparable<LuxuryCar>
{
    Class code here...
}
```

If we have two `LuxuryCar` objects, `mercedes` and `jaguar`, we can compare the two objects with an expression such as:

```
mercedes.compareTo(jaguar)
```

This expression will return a negative number, 0, or a positive number, indicating that the `mercedes` object is less than, equal to, or greater than the `jaguar` object.

The way we should write our `greatest` method is to make sure that only objects that implement the `Comparable` interface are passed to it as arguments. That way, we can safely use one of the objects to call the `compareTo` method, and determine which object is the greatest. With generics, we could write the method as follows:

```
public static < T extends Comparable<T> > T greatest(T arg1, T arg2)
{
    if (arg1.compareTo(arg2) > 0)
        return arg1;
    else
        return arg2;
}
```

This method defines a type parameter `T` with the following notation:

```
< T extends Comparable<T> >
```

In generic notation, the key word `extends` also means “implements.” This specifies that the type passed to `T` must implement the `Comparable` interface. The method has two parameter variables,

`arg1` and `arg2`. Because these parameter variables are of type `T`, any objects passed to them must implement the `Comparable` interface. Notice that the method's return type is also `T`, meaning that the method will return a reference to an object of type `T`.

[Code Listing 17-13](#) shows a program that uses a generic method to search an array for a value. Both the array and the search value must be of a type that implements the `Comparable` interface. This is because the method uses the `compareTo` method to determine whether an element contains the search value.

Code Listing 17-13

(GenericSearchArray.java)

```
1  /**
2   * This program uses a generic method to sequentially
3   * search an array for a value.
4   */
5
6  public class GenericSearchArray
7  {
8      public static void main(String[] args)
9      {
10         int position; // To hold a string's position in the array
11
12         // Array of strings to search
13         String[] names = { "Jack", "Kelly", "Beth",
14                             "Chris", "Kenny", "Britainy" };
15
16         // Search the array for Chris.
17         position = sequentialSearch(names, "Chris");
18
19         // Determine whether Chris was found.
20         if (position == -1)
21             System.out.println("Chris is not in the array. ");
22         else
23             System.out.println("Chris is at position " + position);
24     }
25
26     /**
27      * The sequentialSearch method searches an array for
28      * a value.
29      * @param array The array to search.
30      * @param value The value to search for.
31      * @return The subscript of the value if found in the
32      *         array, otherwise -1.
33     */
34
35     public static < E extends Comparable<E> >
36         int sequentialSearch(E[] array, E value)
37     {
38         int index;           // Loop control variable
39         int position;        // Position the value is found at
40         boolean found;       // Flag indicating search results
```

```

41
42 // Position 0 is the starting point of the search.
43 index = 0;
44
45 // Store the default values for position and found.
46 position = -1;
47 found = false;
48
49 // Search the array.
50 while (!found && index < array.length)
51 {
52     if (array[index].compareTo(value) == 0)
53     {
54         found = true;
55         position = index;
56     }
57     index++;
58 }
59 return position;
60 }
61 }

```

Program Output

Chris is at position 3

Take a closer look at the `sequentialSearch` method header, which appears in lines 35 and 36:

```

public static < E extends Comparable<E> >
    int sequentialSearch(E[] array, E value)

```

This method header defines a type parameter, `E`, which is constrained to types that implement the `Comparable` interface. The method's first parameter variable, `array`, is an array of type `E`. The second parameter variable, `value`, is a reference to an object of type `E`. If `value` is found in the array, the method returns the subscript of the element containing it. If `value` is not found in the array, the method returns `-1`.



Checkpoint

1. 17.21 Can interfaces be defined as generic?
2. 17.22 Assume that you must define the type parameter `T` in a generic method. How would you constrain the type parameter `T` so it only accepts type arguments that implement the `Comparable` interface?

In the Spotlight:



Generic Subsets

Consider the task of generating a random subset of a set of some type. You might need to do this if you are writing a game of chance or if you are simulating certain types of real-world phenomena.

Let us suppose the set whose random subset we are to generate is in the form of a generic list,

```
List<T> argList
```

where T is a type parameter. There are different strategies we can use to pick a random subset: Here is one approach.

If the set has size N , it will have 2^N subsets. Each subset will correspond to a bit string of length N ; and in turn, each bit string will correspond to a binary number in the range $0 \text{--} 2^N - 1$. For example, if the set is $S = \{A, B, C\}$, then $N = 3$ and each of the eight subsets of S corresponds to a bit string of length 3. This three-way correspondence between subsets, bit strings, and numbers is illustrated in the following table.

Subset	$\{A, B, C\}$	$\{A, B\}$	$\{A, C\}$	$\{A\}$	$\{B, C\}$	$\{B\}$	$\{C\}$	$\{\}$
Bit string	111	110	101	100	011	010	001	000
Number	7	6	5	4	3	2	1	0

Briefly, each element of the original set is assigned a position. A bit string has a 1 in that position whenever the corresponding subset includes the element at that position. For a set of size N , we can use the `nextInt()` method of the `java.util.Random` class to select a random integer k in the range $0 \text{--} 2^N - 1$. We can next convert k to the binary string that represents it and then map this binary string to its corresponding subset. For example, if the random number selected is 5, then the bit string is 101, and the selected random subset is $\{A, C\}$. In this way, we can select random subsets of a set of size N by selecting random integers in the range $0 \text{--} 2^N - 1$.

There are two static methods of the `Integer` class that will prove useful. The first is as follows:

```
int rotateLeft(int i, int N)
```

Calling this method with a value of 1 for the first parameter i and a value of N less than 32 will return $2N$, the number of subsets in a set of size N . This is shown in Line 41 of the following code listing. To convert an integer to the corresponding bit string, we use the method

```
String toBinaryString(int i);
```

as shown in line 47. Contrary to what you might expect, the length of the string returned, `subsetStr`, may not be the same as the size of the `argList`, the input set. This is because leading zeroes are omitted. For example, the bit string corresponding to $\{B, C\}$ will be 11 rather than 011. To compensate, the code that maps 1 bits to list elements must shift the index into `argList` by the number of zero bits omitted. This code appears in lines 50–55.

The following code listing demonstrates a generic class with a generic method for constructing a random subset. The `main` method exercises the method by calling the method three times, passing it the same list of strings on each call. As the output of the program shows, each call to the method is likely to return a different subset.

Code Listing 17-14

(GenericRandomSubset.java)

```
1 import java.util.*;
2
3 /**
4  This program shows how to generate a random subset of a set S
5  when S is given as a generic list.
6  */
7
8 public class GenericRandomSubset <T>
9 {
10     public static void main(String[] args)
11     {
12         // Create a random subsets object specialized for strings
13         Subsets<String> strSubsets = new Subsets<>();
14
15         String [] people = {"Moe", "Curly", "Larry", "Bobby", "Stinky"};
16
17         // Transform the array into a list
18         List<String> peopleList = Arrays.asList(people);
19         System.out.println(strSubsets.getRandomSubset(peopleList));
20         System.out.println(strSubsets.getRandomSubset(peopleList));
21         System.out.println(strSubsets.getRandomSubset(peopleList));
22     }
23 }
24
25 class Subsets<T>
26 {
27     Random rand = new Random();
28
29     /**
30      This method computes a random subset of an input set.
```

```

31     @param argList: represents the input set
32     @return a list representing a random subset of a set
33 */
34
35 List<T> getRandomSubset(List<T> argList)
36 {
37     List<T> resultList = new ArrayList<T>();
38     int argListSize = argList.size();
39
40     // Compute the size of the power set of the input set.
41     int powerSetSize = Integer.rotateLeft(1, argListSize);
42
43     // Select a random integer in the range 0 to powerSetSize-1.
44     int randomInt = rand.nextInt(powerSetSize);
45
46     // Convert the random integer to a subset
47     String subsetStr = Integer.toBinaryString(randomInt);
48
49     // Interpret each bit in the string as an element
50     int length = subsetStr.length();
51     for (int k = 0; k < length; k++)
52     {
53         if (subsetStr.charAt(k) == '1')
54             resultList.add(argList.get(k + argList.size() - length));
55     }
56     return resultList;
57 }
58 }

```

Program Output

```

[Moe, Curly, Larry, Bobby]
[Moe, Curly, Bobby]
[Bobby, Stinky]

```


17.9 Erasure

CONCEPT: When the Java compiler processes a generic class or method, it erases the generic notation and substitutes an actual type for each type parameter.

When the Java compiler encounters generic code, it uses a process known as *erasure* to compile the code. The process is called erasure because the compiler erases the generic notation and substitutes an actual type for each type parameter. Here's how erasure works: When the compiler encounters a generic class, interface, or method with an unbound type parameter, such as `<T>` or `<E>`, it replaces all occurrences of the type parameter with `Object`. For example, the following class:

```
public class Point<T>
{
    private T xCoordinate;
    private T yCoordinate;

    public Point(T x, T y)
    { code... }

    public void setX(T x)
    { code... }

    public void setY(T y)
    { code... }

    public T getX()
    { code... }

    public T getY()
    { code... }
}
```

would be converted by the process of erasure to this code:

```
public class Point
{
    private Object xCoordinate;
    private Object yCoordinate;

    public Point(Object x, Object y)
    { code... }

    public void setX(Object x)
    { code... }

    public void setY(Object y)
    { code... }

    public Object getX()
    { code... }
}
```

```

    public Object getY()
    { code... }
}

```

And, the following generic method:

```

public static <E> void displayArray(E[] array)
{
    for (E element : array)
        System.out.println(element);
}

```

would be converted to:

```

public static void displayArray(Object[] array)
{
    for (Object element : array)
        System.out.println(element);
}

```

When the compiler encounters a class, interface, or method with a bound type parameter, such as `<T extends Number>` or `<E extends Comparable>`, it replaces all occurrences of the type parameter with the bound that is applied to the parameter. For example, the following class:

```

public class Point<T extends Number>
{
    private T xCoordinate;
    private T yCoordinate;

    public Point(T x, T y)
    { code... }

    public void setX(T x)
    { code... }

    public void setY(T y)
    { code... }

    public T getX()
    { code... }

    public T getY()
    { code... }
}

```

would be converted by the process of erasure to this code:

```

public class Point
{
    private Number xCoordinate;
    private Number yCoordinate;

    public Point(Number x, Number y)
    { code... }
}

```

```

public void setX(Number x)
{ code... }

public void setY(Number y)
{ code... }

public Number getX()
{ code... }

public Number getY()
{ code... }
}

```

And, the following generic method:

```

public static <E extends Comparable<E>>
    int sequentialSearch(E[] array, E value)
{
    method code here...
}

```

would be converted to:

```

public static
    int sequentialSearch(Comparable[] array, Comparable value)
{
    method code here...
}

```

As you can see, generic notation exists only in the source code that is written by the programmer. Once the compiler confirms that a generic type is being used safely, it converts it into a raw type, which is simply a non-generic type. At the bytecode level, only raw types exist.

When working with generic types, the compiler also inserts cast operators, where necessary, in the code that uses the generic type. For example, suppose we use the following code to create an instance of the `Point` class:

```

Integer x = new Integer(1);
Integer y = new Integer(2);
Point<Integer> myPoint = new Point<>(x, y);

```

Then, later in the same program we write the following statement:

```

Integer tempX = myPoint.getX();

```

Because the translated `getX` method returns an `Object` reference, the compiler will modify this statement to read the following:

```

Integer tempX = (Integer)myPoint.getX();

```

Generics were introduced in Java 5. Older versions of Java did not support generics, so programmers could only write raw classes. The process of erasure ensures that newer code utilizing generics can

coexist with older code that was written before Java 5.



Checkpoint

1. 17.23 During the process of erasure, when the compiler encounters a generic class, interface, or method with an unbound type parameter, such as `<T>` or `<E>`, what does it replace the type parameter with?
2. 17.24 During the process of erasure, when the compiler encounters a class, interface, or method with a bound type parameter, such as `<T extends Number>` or `<E extends Comparable>`, what does it replace the type parameter with?
3. 17.25 Do generic types still exist at the bytecode level?

17.10 Restrictions on the Use of Generic Types

CONCEPT: Type parameters cannot be used in every way that a real type can be used.

Because of the erasure process, there are restrictions on how and when generic types and type parameters can be used. In this section we will examine some of those restrictions.

You cannot create an instance of a type parameter.

Let's say that a generic class has a type parameter named `T`. In the class you cannot create an instance of `T`. For example, the following code will not compile:

```
1 public class MyClass<T>
2 {
3     public MyClass()
4     {
5         // The following statement causes an ERROR.
6         T myObject = new T();
7     }
8 }
```

In this code the statement in line 6 causes an error at compile time. The expression `new T()` implies that `T` is the name of a constructor, but it is not.

You cannot create an array of generic class objects.

The Java compiler will not allow you to create an array of objects that are instances of a generic type. For example, the following statement will not compile:

```
ArrayList<String>[] a = new ArrayList<>[100];
```

At runtime, arrays in Java carry internal information about the data type of their elements. Because of erasure, an array of generic class instances would not have the necessary runtime type information, so the compiler will not allow such arrays to be created.

A generic class's type parameter cannot be the

type of a static field, and cannot be referred to in a static method.

You cannot use a type parameter as the type of a static field. For example, the following code will not compile:

```
1 public class MyClass<T>
2 {
3     // The following statement will cause an ERROR.
3     private static T value;
4 }
```

Regardless of the number of instances of `MyClass`, there is only one instance of the static variable `value`, and therefore `value` can have only one type. The compiler also prohibits a static method from referring to a class's type parameter. The following code is illegal:

```
1 public class MyClass<T>
2 {
3     public static void doSomething()
4     {
5         // The following statement causes an ERROR.
6         T myValue;
7     }
8 }
```

You cannot make an exception class generic.

Let's imagine what would have to take place in order for a Java application to use generic exception classes. First, you would have to be able to use generic type notation in a `catch` clause. Then, when a generic exception was thrown, the system would have to check the type argument used to create the exception class to see if it matches the type specified in the `catch` clause. This is not possible because at runtime, generic type information has already been erased by the compiler. Because the compiler cannot know at compile time where an exception will come from, or what type argument will be used to create it, you are not allowed to create a generic exception class. Specifically, you are not allowed to create a generic subclass of `Throwable`. You are also not allowed to use generic type notation in a `catch` clause.

17.11 Common Errors to Avoid

- Forgetting to pass a type argument when instantiating a generic class. When you instantiate a generic class without passing a type argument, you are using the class as a raw type. When you do this, you give up the benefits of type-safety that you normally have with a generic class.
- Using a primitive type as a type argument. Generic classes and methods can accept only reference types as type arguments.
- Misusing inheritance when passing an instance of a generic class as an argument to a method. If a method's parameter variable is declared as the `Point<Number>` type, you cannot pass an instance of `Point<Integer>` to the parameter. Even though `Integer` is a subclass of `Number`, `Point<Integer>` is not a subclass of `Point<Number>`. When a parameter variable is declared as `Point<Number>`, the only arguments that can be passed to it are instances of `Point<Number>`.
- Being too type-restrictive when using a generic class to declare a parameter variable. Suppose in a method you declare a parameter variable as `Point<Integer> myPoint`. The only arguments that can be passed to this parameter are those of the `Point<Integer>` type. If you try to pass a `Point<Double>` object, an error will occur. By declaring the parameter as `Point<?> myPoint`, you will be able to pass any `Point` object to the method. By declaring the parameter variable as `Point<? extends Number> myPoint`, you will be able to pass any `Point` object with a type argument that is `Number` or a subclass of `Number`.
- Trying to use the `implements` key word to constrain a type parameter to types that implement an interface. You do not use the `implements` key word in generic type notation. To constrain a type parameter to types that implement a particular interface, you use the `extends` key word. In generic notation, the `extends` key word means “extends” or “implements.”
- Forgetting to define a type parameter in a generic method. If a generic method has its own type parameter, you must write a definition for the type parameter in the method header.
- Defining a type parameter in the wrong place in a generic method. Type parameter definitions in a generic method must appear just before the method's return type.

Review Questions and Exercises

Multiple Choice and True/False

1. When you create an instance of a generic class, what types can you pass as arguments to the class's type parameter?
 1. primitive types only
 2. reference types only
 3. interface types only
 4. primitive, reference, and interface types
2. In a generic method, a type parameter is defined ☐ .
 1. inside the parentheses, along with the method's parameter variables
 2. after the method's return type, and before the method's name
 3. before the method's return type
 4. inside the body of the method, but before any local variables are declared

3. Look at the following method header:

```
void displayPoint(Point<Number> myPoint)
```

Which of the following objects would we be allowed to pass as an argument to the `displayPoint` method? (Select all that apply.)

1. `Point<Number> p;`
 2. `Point<Integer> p;`
 3. `Point <Double> p;`
 4. `Point <String> p;`
4. Look at the following method header:

```
void displayPoint(Point<?> myPoint)
```


Which of the following objects would we be allowed to pass as an argument to the `displayPoint` method? (Select all that apply.)

1. `Point<Number> p;`
2. `Point<Integer> p;`
3. `Point <Double> p;`
4. `Point <String> p;`

5. Look at the following method header:

```
void displayPoint(Point<? extends Number> myPoint)
```

Which of the following objects would we be allowed to pass as an argument to the `displayPoint` method? (Select all that apply.)

1. `Point<Number> p;`
2. `Point<Integer> p;`
3. `Point <Double> p;`
4. `Point <String> p;`

6. Look at the following method header:

```
void displayPoint(Point<? super Double> myPoint)
```

Which of the following objects would we be allowed to pass as an argument to the `displayPoint` method? (Select all that apply.)

1. `Point<Number> p;`
2. `Point<Integer> p;`
3. `Point <Double> p;`
4. `Point <String> p;`

7. In the generic type notation `<T extends Number>` to what type of bound is `T` constrained?

1. an upper bound
2. a lower bound
3. both an upper bound and lower bound

4. `T` is not constrained to a bound
8. In the generic type notation `<T super Integer>` to what type of bound is `T` constrained?
 1. an upper bound
 2. a lower bound
 3. both an upper bound and lower bound
 4. `T` is not constrained to a bound
9. Which of the following generic type notations uses a wildcard?
 1. `Point<W>`
 2. `Point<T extends Number>`
 3. `Point<T super Integer>`
 4. `Point<?>`
10. The process used by the Java compiler to remove generic notation and substitute an actual type for type parameters is known as .
 1. erasure
 2. removal
 3. substitution
 4. masking
11. **True or False:** It is better to discover an error at runtime than at compile time.
12. **True or False:** It is possible to instantiate a generic class without specifying a type argument.
13. **True or False:** Type parameters must be single character identifiers, written in uppercase.
14. **True or False:** You can constrain a type parameter to both an upper bound and lower bound.
15. **True or False:** A generic class can extend a non-generic class.
16. **True or False:** You cannot create an array of generic class objects.
17. **True or False:** A generic class's type parameter can be the type of a static field.
18. **True or False:** An exception class cannot be generic.

Find the Error

1. `ArrayList myList = new ArrayList<String>();`

2. Assume the following is a method header in a class:

```
public < T implements Comparable<T> > T greatest(T arg1, T arg2)
```

3. `public class MyClass<T>`
`{`
 `public static void displayValue(T value)`
 `{`
 `System.out.println(value);`
 `}`
`}`

4. `public class Point<T extends Number super Integer>`
`{`
 `Class code . . .`
`}`

Algorithm Workbench

1. Assume there is a class named `Customer`. Write a statement that creates an `ArrayList` that can hold only `Customer` objects.
2. Assume `names` references an object of the `ArrayList<String>` class. Write a statement that adds the string "William Jefferson" to the `names` object.
3. Write the header for a generic class named `MyType`. The class should have one type parameter. The type parameter's upper bound should be the `String` class.
4. Write the header for a generic class named `MyType`. The class should have one type parameter. The type parameter's lower bound should be the `Integer` class.
5. Modify the following `max` method so its arguments can be of any type that implements the `Comparable` interface. The method's return type should be the same as the type of its parameter variables.

```
public static int max(int a, int b)
{
    if (a > b)
        return a;
    else
        return b;
}
```

6. Write the header for a generic class named `MyType`. The class should have two type parameters.
7. Write the header for a generic class named `MyType`. The class should have two type parameters. The first type parameter's upper bound should be the `Number` class. The second type parameter's upper bound should be the `String` class.

Short Answer

1. When a generic class is used in a program, when does type checking take place?
2. Look at the following method header:

```
public < T extends Comparable<T> > T least(T arg1, T arg2)
```

What constraint is placed on the type parameter `T`?

3. What is erasure?
4. Do generic types exist at the bytecode level?
5. When the compiler encounters a generic class, interface, or method with an unbound type parameter, such as `<T>` or `<E>`, it replaces all occurrences of the type parameter with what type?
6. When the compiler encounters a class, interface, or method with a bound type parameter, such as `<T extends Number>` or `<E extends Comparable>`, it replaces all occurrences of the type parameter with what type?

Programming Challenges

1. MyList Class

Write a generic class named `MyList`, with a type parameter `T`. The type parameter `T` should be constrained to an upper bound: the `Number` class. The class should have as a field an `ArrayList` of `T`. Write a public method named `add`, which accepts a parameter of type `T`. When an argument is passed to the method, it is added to the `ArrayList`. Write two other methods, `largest` and `smallest`, which return the largest and smallest values in the `ArrayList`.

2. MyList Modification

Modify the `MyList` class that you wrote for Programming Challenge 1 so the type parameter `T` should accept any type that implements the `Comparable` interface. Test the class in a program that creates one instance of `MyList` to store `Integers`, and another instance to store `Strings`.

3. PointList Class

Write a class named `PointList` that keeps a list of `Point` objects in an `ArrayList`. The `PointList` class should accept any object that is an instance of the `Point` class, or a subclass of `Point`. Demonstrate the class in an application.

4. Generic Insertion Sort

In Programming Challenge 1 of [Chapter 16](#) you wrote an `ObjectInsertionSorter` class that can sort `Comparable` objects. Modify the sorting method so it is a generic method. The method should use generic notation to accept an array of any objects that implement the `Comparable` interface. Demonstrate the class in an application.

5. Generic Binary Search

In Programming Challenge 3 of [Chapter 16](#) you wrote an `ObjectBinarySearcher` class. Modify the search method so it is a generic method. The method should use generic notation to accept an array of any objects that implement the `Comparable` interface. Demonstrate the class in an application.

6. Highest and Lowest Elements



Highest and Lowest Elements

Write a generic class with a type parameter constrained to any type that implements `Comparable`. The constructor should accept an array of such objects. The class should have methods that return the highest and lowest values in the array. Demonstrate the class in an application.

7. Number Analyzer

Write a generic class with a type parameter constrained to the `Number` class or any subclass of `Number`. The constructor should accept an array of such objects. The class should have methods that return the highest and lowest values in the array, the total of the elements, and the average value of all the elements. Demonstrate the class in an application.

8. Generic Random Permutation

Write a generic class with a generic method that returns a random permutation of a list passed to it as a parameter. Demonstrate the working of the method in an application that displays random permutations of lists of strings and of lists of numbers.

Chapter 18 Collections and the Stream API

TOPICS

1. [18.1 Introduction to the Java Collections Framework](#)
2. [18.2 Lists](#)
3. [18.3 Sets](#)
4. [18.4 Maps](#)
5. [18.5 The Collections Class](#)
6. [18.6 The Stream API](#)
7. [18.7 Common Errors to Avoid](#)

18.1 Introduction to the Java Collections Framework

CONCEPT: The Java API provides numerous interfaces and classes that allow you to create collections, which are objects that act as containers for groups of other objects.

It is common for an application to create a group of similar objects and work with that group of objects throughout the application's execution. Suppose, a teacher uses an application to manage student grades. That application might create several `Student` objects, with each object holding information about a particular student. Another example would be a payroll application that creates a group of `Employee` objects, with each object holding information about a specific employee.

Working with groups of objects is so common in programming that the Java API provides a set of interfaces and classes specifically for this purpose. These interfaces and classes, which are in the `java.util` package, are known as the *Java Collections Framework (JCF)*. You can use these interfaces and classes to create *collections*, which are objects that act as *containers* of other objects. You can use a collection to store a group of objects. Each object that is stored in a collection is referred to as an *element*. Collections allow you to add, retrieve, and delete objects.

Lists, Sets, and Maps

There are three general types of collections: *lists*, *sets*, and *maps*. Each of these is described here:

- **List**—A *list* is an *ordered collection* that orders its elements by position. The *position* of an element x within the list, also called its *index*, is the number of elements that precede x within the list. A list permits duplicate elements: two objects that are equal can be added to the list. The list distinguishes equal objects by their index.
- **Set**—A *set* is an *unordered collection* of objects that does not permit duplicates. Unlike a list, a set has no concept of position for its elements. Implementations of sets are usually optimized for searching so you can quickly find a specific stored value.
- **Map**—A *map* is a collection that stores every object x according to a property of x called its *key*. This means that each element in a map is a *key-value pair*. The *key* is a property of the object and the *value* is the object itself. A map allows a value to be quickly retrieved based on its key.

Java Functional Interfaces

The `java.util.function` package defines a number of functional interfaces that are referred to by methods defined in JCF interfaces and classes. In practice, the functional interfaces are used to declare the types of lambda expressions to be passed as a parameter to a collection class method. Knowing the naming system used to name these functional interfaces will help you to better understand the JCF methods that refer to them. The naming system uses the words *supplier*, *consumer*, *predicate*, *operator*, and *function* as root words, and employs a set of prefixes that includes *binary*, *bi*, *unary*, and the names of primitive types to modify the meaning of the root words.

A *function* takes at least one parameter and returns a value. The generic interface `Function<T, R>` takes a parameter of type `T` and returns a value of type `R`. A *bi-function* takes two parameters and returns a value: the generic interface `BiFunction<T, U, R>` takes two parameters of type `T` and `U`, in that order, and returns a value of type `R`.

It is useful to have a short hand notation for functional types. We will use the notation $T \rightarrow R$ for the type `Function<T, R>` and $(T, U) \rightarrow R$ for the type `BiFunction<T, U, R>`. Similar shorthand notation will be used for other functional types.

There are also functional interfaces that involve the primitive types `int`, `long`, and `double`. For example, the interface `IntFunction<R>` is the type $\text{int} \rightarrow R$, the interface `LongFunction<R>` is the type $\text{long} \rightarrow R$, and `IntToLongFunction` is the type $\text{int} \rightarrow \text{long}$.

An *operator* is a function whose parameters and return value have the same type. A *binary operator* is an operator that takes two parameters, and a *unary operator* is an operator that takes a single parameter. Thus `BinaryOperator<T>` has type $(T, T) \rightarrow T$, the interface `UnaryOperator<T>` has type $T \rightarrow T$, and `LongUnaryOperator` has type $\text{long} \rightarrow \text{long}$.

A *consumer* takes at least one parameter and returns no result. The interface `Consumer<T>` has type $T \rightarrow \text{void}$, while the interface `BiConsumer<T, U>` has type $(T, U) \rightarrow \text{void}$.

A *supplier* is the opposite of a consumer: it takes no parameters and returns a value. The interface `Supplier<T>` has type $() \rightarrow T$ and `IntSupplier` has type $() \rightarrow \text{int}$.

A *predicate* is a function that returns a boolean value. Thus `Predicate<T>` has type $T \rightarrow \text{boolean}$, while `BiPredicate<T, U>` has type $(T, U) \rightarrow \text{boolean}$.

JCF Interfaces

The Java Collection Framework defines interfaces and classes to facilitate the creation and use of collections. The classes and interfaces are generic so they can be used with different element types. Four of these interfaces are shown in [Figure 18-1](#).




Figure 18-1 The main JCF interfaces

[Figure 18-1 Full Alternative Text](#)

Iterating a Collection

[Figure 18-1](#) shows that lists and sets are special types of collections, and that every collection is iterable. *Iteration* is the successive retrieval of elements from a collection until all elements have been examined. A collection that provides methods for iteration is said to be *iterable*. The JCF `Iterable<E>` interface defines methods for iterating a collection that stores elements of type `E`. The `Collection` interface extends `Iterable` and adds many methods for working with collections. The `List` interface is implemented by all collections that order their elements by position. It defines methods for adding, removing, and retrieving elements by index. Finally, the `Set` interface serves as the interface for all unordered collections. The `Set` interface does not add any methods to those that it inherits from `Collection`.

There is also a `Map<K, V>` interface that defines methods for a map that stores key-value pairs with `K` as the type of key and `V` as the type of value. This interface does not extend `Collection`, and will be covered later.

The `Iterable` Interface and the `forEach` method

The `Iterable<E>` interface provides two methods that can be used to iterate through a collection. The first of these is the method:

```
void forEach(Consumer<? super E> action)
```

The `action` parameter must be an object that implements the `Consumer<? Super E>` interface. In

practice, this will usually be a lambda expression that takes a single parameter and returns no value. Moreover, the type of the parameter must be a superclass of $\mathbb{E}$, the collection's element type.

We can illustrate the use of `forEach()` using `ArrayList`, a class that you learned about in [chapter 7](#). The program in [Code Listing 18-1](#) creates a list with four names, and then uses `forEach()` to invoke the consumer function defined by the lambda expression

```
x ->
{
    System.out.printf("%s %d\n", x, x.length());
}
```

on each element of the list. The effect is to print the name and length of each element in the list.

Code Listing 18-1 (ForEachDemo.java)

```
1 import java.util.ArrayList;
2
3 /**
4  * This program demonstrates Iterable forEach().
5  */
6
7 public class ForEachDemo
8 {
9     public static void main(String [] args)
10    {
11        // Array of names.
12        String [] names = {"Anna", "Bob", "Carlos", "Debby"};
13
14        // Create list and add names.
15        ArrayList<String> nameList = new ArrayList<>();
16        for (String name : names)
17        {
18            nameList.add(name);
19        }
20
21        // Use forEach with lambda expression to print.
22        nameList.forEach(
23            x ->
24            {
25                System.out.printf("%s %d\n", x, x.length());
26            });
27    }
28 }
```

Program Output

```
Anna 4
Bob 3
Carlos 6
```

Iterators

Another method in the `Iterable<E>` is as follows:

```
Iterator<E> iterator()
```

Calling this method on a collection returns an iterator for that collection. An *iterator* for a collection *C* is an object that provides controlled iteration through the elements of *C*. Unlike the `forEach()` method, which applies an action to *all* elements of a collection, an iterator provides finer control, letting you call for the next element only when you need it, and allowing you to stop the iteration at any stage of the process. An iterator is defined by the `Iterator<E>` interface, whose methods are defined in [Table 18-1](#).

Table 18-1 Methods of the `Iterator<E>` interface

<code>boolean hasNext()</code>	Returns <code>true</code> if there are elements that have not yet been returned by the iterator, and <code>false</code> otherwise.
<code>E next()</code>	Returns the next element from the collection, or throws <code>NoSuchElementException</code> if no more elements remain.
<code>void remove()</code>	Removes from the underlying collection the element that was returned by the last call to <code>next()</code> . You can only call this method once after each call to <code>next()</code> .
<code>void forEachRemaining(Consumer<? super E> action)</code>	Calls the specified action method for each element that has not yet been returned by the iterator.

You get an iterator by calling the `iterator()` method on a collection. This method will take a snapshot of the elements in the collection at the time of the call and return an iterator that is ready to provide access to those elements. You typically use the iterator in a loop, calling `hasNext()` to see if

there are more elements yet to be examined, and calling `next()` to fetch the next available element. The `remove()` method is optional, so it is not supported by all iterators. If supported, it will remove the element returned by the last call to `next()` from the underlying collection. The `forEachRemaining()` method applies an action to every element that has not yet been returned by the iterator.

The program in [Code Listing 18-2](#) is a variation of the one in [Code Listing 18-1](#). It prints names and lengths until "Bob" has been printed, then it prints the name without the length for each remaining list element.

To make the code more readable, the program defines the lambda expression in a separate statement:

```
Consumer<String> action = x ->
    {
        System.out.printf("%s\n", x);
    };
```

Notice that the functional interface must be specialized to the element type of the collection, which is `String`. The iterator is created like this:

```
Iterator<String> iter = nameList.iterator();
```

The call `iter.hasNext()` checks to if there are still elements to be returned by the iterator, and the call to `iter.next()` returns the next available element.

Code Listing 18-2 (IteratorDemo.java)

```
1 import java.util.ArrayList;
2 import java.util.Iterator;
3 import java.util.function.Consumer;
4
5 /**
6  This program demonstrates Iterators.
7  */
8
9 public class IteratorDemo
10 {
11     public static void main(String [] args)
12     {
13         // Array of names.
14         String [] names = {"Anna", "Bob", "Carlos", "Debby"};
15
16         // Create list and add names.
17         ArrayList<String> nameList = new ArrayList<>();
18         for (String name : names)
19         {
20             nameList.add(name);
21         }
22
23         // Define an action for the "remaining" elements.
24         Consumer<String> action = x ->
```

```

25         {
26             System.out.printf("%s\n", x);
27         };
28
29         // Get the iterator to the list.
30         Iterator<String> iter = nameList.iterator();
31
32         // Process list elements with the iterator.
33         while (iter.hasNext())
34         {
35             String name = iter.next();
36             System.out.printf("%s %d\n", name, name.length());
37             if (name.equals("Bob"))
38             {
39                 // Act differently for names after "Bob".
40                 iter.forEachRemaining(action);
41             }
42         }
43     }
44 }

```

Program Output

```

Anna 4
Bob 3
Carlos
Debby

```

You can think of an iterator as having a list of all elements in the collection and a cursor that is initially positioned just before the first element of that list. Every call to `next()` returns the element at the cursor position and moves the cursor past the element returned. The `hasNext()` method returns false when the cursor gets past the last element in the list.

NOTE:

The `next()` method throws `NoSuchElementException` if it is called when there are no elements left to return.

The Enhanced For Loop

In [Chapter 7](#), you learned that the enhanced `for` loop can be used with arrays and `ArrayList` collections. In fact, the enhanced `for` loop will work with any object of any class that implements the `Iterable` interface. In particular, it can be used with any collection.

The Collection Interface

The `Collection<E>` interface defines methods for working with collections. Some of these methods, with brief explanations of their use, are shown in [Table 18-2](#).

Table 18-2 Methods of the `Collection<E>` interface

<code>boolean add(E e)</code>	Adds the object referenced by <code>e</code> to this collection. Some collections restrict the objects that can be added as elements. This method returns <code>true</code> if the object was added to the collection. It returns <code>false</code> if the collection did not allow the object to be added. This would be the case if the collection is a set and the <code>add()</code> is attempting to add a duplicate.

`boolean addAll(Collection<? extends E> c)`

The parameter `c` is another `Collection` object. This method adds all of the elements of `c` to this collection. The method returns `true` if the elements of `c` were added to this collection. It returns `false` otherwise.

<code>void clear()</code>	Removes all of the elements in this collection.

<code>boolean contains(Object o)</code>	Returns <code>true</code> if this collection contains the specified object, or <code>false</code> otherwise.

`boolean containsAll(Collection<?> c)`

Returns `true` if this collection contains all of the elements in the collection `c`, and returns `false` otherwise.

<code>boolean equals(Object o)</code>	Returns <code>true</code> if this collection is equal to the object <code>o</code> , and <code>false</code> otherwise.

<code>boolean isEmpty()</code>	Returns <code>true</code> if this collection contains no elements, and <code>false</code> otherwise.
<code>boolean remove(Object o)</code>	Removes a single instance of the object <code>o</code> from this collection if found, and returns <code>true</code> . If <code>o</code> is not found, the method returns <code>false</code> .
<code>boolean removeAll(Collection<?> c)</code>	Removes every element in the collection <code>c</code> from this collection. Returns <code>true</code> if at least one element was removed from this collection, and <code>false</code> otherwise.
<code>boolean removeIf(Predicate<? super E> filter)</code>	Removes all of the elements of this collection that satisfy the given predicate. Returns <code>true</code> if at least one element was removed from this collection, and <code>false</code> otherwise.
<code>boolean retainAll(Collection<?> c)</code>	The parameter <code>c</code> is another collection. This collection will retain all of its elements that are also elements of the collection <code>c</code> . All other elements in this collection are removed. Returns <code>true</code> if at least one element was removed.
<code>int size()</code>	Returns the number of elements in this collection.
<code>Stream<E> stream()</code>	Returns a sequential stream with this collection as its source. Streams are discussed later in this chapter.
<code>Object[] toArray()</code>	This method converts the collection into an array of <code>Object</code> and returns the array.

The parameter to this method is an array of some reference type `T`. The


```
<T> T[]
toArray(T[] a)
```

elements of this collection are stored as objects of type `T` in the array `a` and the array `a` is returned. If this collection contains more elements than the array `a`, then a new array is created to hold the elements of this collection and the new array is returned. If this collection holds fewer elements than the array `a`, then the elements of this collection are stored in `a` and the element of `a` that immediately follows the last collection element is set to `null`.

The `removeIf` method uses a predicate as a filter. This method filters out all elements that satisfy the collection, leaving in the collection only those elements on which the predicate returns a value of `false`. [Code Listing 18-3](#) uses this method to remove from a list of names all names whose length is at most 4. First, we define a lambda expression to be used as the filter:

```
Predicate<String> filter = x -> x.length() <= 4;
```

The lambda expression returns the value of the boolean expression `x.length() <= 4`. The filter function will return `true` when the length of the string `x` is at most 4, and `false` when the length of `x` is greater than 4.

Code Listing 18-3 (FilterDemo.java)

```
1 import java.util.ArrayList;
2 import java.util.function.Predicate;
3
4 /**
5  This program demonstrates the Collection removeIf.
6  */
7
8 public class FilterDemo
9 {
10     public static void main(String [] args)
11     {
12         // Array of names.
13         String [] names = {"Anna", "Bob", "Carlos", "Debby"};
14
15         // Create list and add names.
16         ArrayList<String> nameList = new ArrayList<>();
17         for (String name : names)
18         {
19             nameList.add(name);
20         }
21         // Use forEach with lambda expression to print.
22         Predicate<String> filter = x -> x.length() <= 4;
23
24         // Remove strings with length at most 4 from nameList
25         nameList.removeIf(filter);
26
27         // Print the array list to show remaining strings.
28         System.out.println(nameList);
```

```
29     }  
30 }
```

Program Output

[Carlos, Debby]



Checkpoint

1. 18.1 The interfaces and classes in the Java Collections Framework are part of what package?
2. 18.2 What are the three general types of collections?
3. 18.3 Describe the differences between a list and a set.
4. 18.4 A map's elements are comprised of how many parts? What are they?
5. 18.5 The `List` and `Set` interfaces extend what other interface?
6. 18.6 Name two methods in the `Iterable` interface.
7. 18.7 What is an iterator?

18.2 Lists

CONCEPT: There are two different ways to implement the `List` interface. The `ArrayList` class uses an implementation based on contiguous memory allocation, while the `LinkedList` implementation is based on linked allocation.

A list is a collection that stores its elements as a sequence. Each element in a sequence is assigned a nonnegative integer called its index. The first element in the sequence is assigned an index of 0, the next an index of 1, and so on.

The `List` Interface

The `List` interface specifies the operations that a list must support. It extends the `Collection` interface, so it inherits all of the methods specified by `Collection`. It also adds additional methods, shown in [Table 18-3](#), specifically for working with lists. In a nutshell, the `List` interface specifies methods for the following:

Table 18-3 Methods of the `List<E>` interface

<code>void add(int index, E element)</code>	Adds <code>element</code> to this list at the position specified by <code>index</code> .

<code>boolean addAll(int index, Collection<? extends E> c)</code>	

The parameter `c` is a `Collection` object. This method adds all of the elements of `c` to this list, starting at the position specified by `index`. Any elements that are currently stored in this list, beginning at the position specified by `index`, are shifted into positions of higher index. The method returns `true` if the elements of `c` were added to this list. It returns `false` otherwise.

<code>E get(int index)</code>	Returns a reference to the element at the position specified by <code>index</code> .

<pre>int indexOf(Object o)</pre>	Returns the position of the first occurrence of the object <code>o</code> in the list. If <code>o</code> does not exist in the list, the method returns <code>-1</code>.
<pre>int lastIndexOf(Object o)</pre>	Returns the position of the last occurrence of the object <code>o</code> in the list. If <code>o</code> does not exist in the list, the method returns <code>-1</code>.

<pre>E remove(int index)</pre>	This method removes the element at the position specified by <code>index</code>, and returns a reference to the element that was removed.

<pre>E set(int index, E element)</pre>	This method stores the object referenced by <code>element</code> at the position specified by <code>index</code>. The object currently stored at the specified position is replaced.

```
void sort(Comparator<? super E> c)
```

Sorts this list according to the order induced by the specified comparator.

```
List<E> subList(int fromIndex, int toIndex)
```

This method returns a `List` object containing the elements of this `List`, starting with the element at `fromIndex` and ending with the element at `toIndex`.

- Adding an element at a specific index
- Retrieving an element from a specific index
- Changing the value of an element at a specific index
- Removing an element that is at a specific index
- Searching for an element in the list and returning its index
- Retrieving a segment of the list as a sublist

Notice that many of the methods in [Table 18-3](#) take an integer parameter for an index. The value of the index must be within the proper bounds, or the method will throw an `IndexOutOfBoundsException`. Elements already in the list have nonnegative indices that are less than the length of the list, so `get()`,

`remove()`, and `set()` will throw this exception if the index passed to them is negative or greater or equal to the length of the list. The `add()` and `addAll()` methods throw the exception if the index is less than 0, or greater than the length of the list.

The ArrayList and LinkedList Classes

The purpose of an interface is to allow different classes to provide different implementations for the methods defined in the interface. JCF provides two implementations of the `List` interface. One of these is the `ArrayList` class, which you encountered in [Chapter 7](#). An array list internally uses a backing array to keep track of references to its elements. The `ArrayList` class overcomes the fixed-size limitation of arrays by automatically allocating a bigger array whenever its backing array is about to overflow, and copying all its elements to the bigger array. The `ArrayList` implementation is said to use *contiguous memory allocation* because all the elements in the backing array occupy successive locations in memory.

Another implementation of the `List` interface is the `LinkedList` class. This class uses *linked allocation*: it keeps track of a reference to the first element in the list, and each element contains a reference to its successor. These successor references are called *links*. Thus, the element at index 0 contains a link to the element at index 1, which contains a link to the element at index 2, and so on. The last element in the list has no successor, so its successor reference is `null`. A list that uses linked allocation is called a *linked list*. Linked lists are studied further in [Chapter 19](#).

Because of the nature of linked allocation, it takes longer to access an element near the end of the list than one near the beginning. This is because to get to an element at position k , you must start at zero and follow the successor links through all elements of lower index. In contrast, `ArrayList` can jump directly to the desired element by using k to index the backing array. Because of this, `ArrayList` is said to provide *random access* to its elements.

The `ArrayList` implementation performs best when most of the elements added to the list are added at the end (the position with the highest index), and most of the elements being removed are removed from the end. This is because inserting an element in the middle of the backing array requires space to be created for the new element by moving all elements with index equal or greater to the insertion point up by one. Similarly, deleting an element in the middle leaves a gap in the backing array, and requires all elements after the deletion point to move down by one to close the gap.

The `LinkedList` implementation performs best when many elements need to be added or removed at a position in the middle of the list. Such an operation can be carried out by reassigning a few successor links, unlike contiguous allocation where lots of elements may need to be shifted.

Creating and Initializing Lists

To create a list using the `ArrayList<E>` implementation, you need to use one of the following constructors:

```
ArrayList()  
ArrayList(int initialCapacity)  
ArrayList(Collection<? extends E> c)
```

The first constructor creates an empty list with the default initial capacity of 10, while the second constructor creates a list with the specified initial capacity. The *capacity* of an array list is the maximum number of elements it can hold before its backing array overflows. The capacity should not be confused with the *size* of the list, which is the number of elements currently stored in the list. The last constructor creates a list containing the same elements as the specified collection, in the order returned by that collection's iterator.

To create a list using the `LinkedList<E>` implementation, you need to use one of its two constructors:

```
LinkedList()  
LinkedList(Collection<? extends E> c)
```

These constructors allow you to create an empty list, or a list whose elements come from a specified collection.

If you need to initialize a list to a long list of values, you can save yourself some typing by first initializing an array and then copying the array elements to a list using a loop:

```
// Initialize array of names.  
String [] namesArray = {"Anna", "Bob", "Carlos", "Debby"};  
  
// Create empty list.  
LinkedList<String> list1 = new LinkedList<>();  
  
// Populate the list using a loop.  
for (String s : namesArray)  
{  
    list1.add(s);  
}  
System.out.println(list1);
```

Here we have created a linked list, added some elements to it, and printed the list. All implementations of collections in Java provide a `toString()` method, so collections can be printed without resorting to a loop to print the individual elements of the collection.

We can simplify this code further by using the constructor that takes an existing collection to initialize a list. The `Arrays` class has a method

```
static <T> List<T> asList(T ... a)
```

that takes an array of a reference type `T` and returns a fixed-size list backed by that array. So you can create a list from an array like this:

```
List<String> aList = Arrays.asList(namesArray);
```

You cannot add elements or remove elements from a fixed-size list, but you can use it as a parameter to a `LinkedList` or `ArrayList` constructor. Thus, you can initialize a list with names from an array,

and a new name to the list, and print the list like this:

```
String [] namesArray = {"Anna", "Bob", "Carlos", "Debby"};
List<String> fixedList = Arrays.asList(namesArray);
LinkedList list2 = new LinkedList<>(fixedList);
list2.add("Ephraim");
System.out.println(list2);
```

This code will print the line:

```
[Anna, Bob, Carlos, Debby, Ephraim]
```

Because the `asList` method takes a var-args parameter, you can shorten the code to create and initialize the list further:

```
LinkedList list2 =
    new LinkedList<>(Arrays.asList("Anna", "Bob", "Carlos", "Debby"));
```

List Iterators

For lists, the iterator inherited from `Collection` returns list elements in order of increasing index, and its `next()` method always moves “forward” through the list. The `List<E>` interface defines a method:

```
ListIterator<E> listIterator()
```

When called on a list, this method returns a special *list iterator* that is able to move both forward and backward through a list. In addition to a `next()` method, a list iterator has a `previous()` method that returns an element while moving backward. A list iterator has a cursor position that always lies between the element that would be returned by a call to `previous()` and the element that would be returned a call to `next()`. Initially, the cursor is just before the element with index 0. When `next()` is called, the cursor returns the element at 0 and moves forward past the returned element to the position between elements 0 and 1. At this point, a call to `previous()` will return the element 0 and move the cursor back to the beginning, but a call to `next()` will return element 1 and move the cursor forward to the position between elements 1 and 2. If the cursor is already past the last element, a call to `next()` will throw `NoSuchElementException`. Likewise, a call to `previous()` will throw the same exception if the cursor position is before the element 0.

The `remove()` method in a list iterator removes from the list the element that was returned by the last call to `next()` or `previous()`. The `remove()` method can only be called once after a call to either of those two methods. A list iterator also has an `add()` method that adds a new element to the list at the current cursor position. The new element goes just after the element that would be returned by a call to `previous()`, and just before the element that would be returned by a call to `next()`.

The methods of a list iterator are defined in the `ListIterator` interface. This interface extends `Iterator` and adds the methods shown in [Table 18-4](#).

Table 18-4 Methods of the `ListIterator<E>` interface

<code>void add(E e)</code>	Adds the specified element into the list at the current cursor position.
<code>E previous()</code>	Returns the element just before the cursor position and moves the cursor position backward.
<code>boolean hasPrevious()</code>	Returns <code>true</code> if there is an element before the current cursor position, <code>false</code> otherwise.
<code>int nextIndex()</code>	Returns the index of the element that would be returned by a subsequent call to <code>next()</code> .
<code>int previousIndex()</code>	Returns the index of the element that would be returned by a subsequent call to <code>previous()</code> .
<code>void set(E e)</code>	Replaces the last element returned by <code>next()</code> or <code>previous()</code> .

The program in [Code Listing 18-4](#) demonstrates how to use a list iterator to traverse a list and add a new element.

Code Listing 18-4 (`ListIteratorDemo.java`)

```
1 import java.util.*;
2
3 /**
4  * Demonstrates the list iterator
5  */
6
7 public class ListIteratorDemo
8 {
9     public static void main(String [] args)
10    {
11        // Create an array list to hold strings.
12        String []names = {"Chris", "David", "Katherine", "Kenny"};
13        List<String> nameList = new ArrayList<>(Arrays.asList(names));
14    }
```



```

15 // Display the names in the list.
16 System.out.println("Here are the original names.");
17 System.out.println(nameList);
18
19 // Get a list iterator.
20 ListIterator<String> it = nameList.listIterator();
21
22 // Add "Darlene" to the list right after "Chris".
23 while (it.hasNext())
24 {
25     String str = it.next();
26     // If the last name retrieved was "Chris"
27     // then insert "Darlene".
28     if (str.equalsIgnoreCase("Chris"))
29     {
30         it.add("Darlene");
31         // We are done
32         break;
33     }
34 }
35 //Display the names in the list again.
36 System.out.println("\nHere are the new names now.");
37 System.out.println(nameList);
38 }
39 }

```

Program Output

```

Here are the original names.
[Chris, David, Katherine, Kenny]
Here are the new names now.
[Chris, Darlene, David, Katherine, Kenny]

```

The program in [Code Listing 18-5](#) demonstrates using the `set()` method to replace an element, and the `remove()` method to remove an element. In this program, the first name in the list is replaced with "Alfonso" and the last name is removed.

Code Listing 18-5 (ListIteratorDemo2.java)

```

1 import java.util.*;
2
3 /**
4  * This program demonstrates setting and removing list
5  * elements with a ListIterator.
6  */
7
8 public class ListIteratorDemo2
9 {
10     public static void main(String[] args)
11     {
12         // Create an ArrayList to hold String objects.

```

```

13 List<String> nameList = new ArrayList<>();
14
15 // Add some names to the ArrayList.
16 nameList.add("Chris");
17 nameList.add("David");
18 nameList.add("Katherine");
19 nameList.add("Kenny");
20
21 // Display the names in the list
22 System.out.println("Here are the original names.");
23 System.out.println(nameList);
24
25 // Get a list iterator.
26 ListIterator<String> it = nameList.listIterator();
27
28 // Replace the first element with "Alfonso".
29 if (it.hasNext())
30 {
31     // Get the first element.
32     it.next();
33
34     // Replace it.
35     it.set("Alfonso");
36 }
37 // Move the iterator to the end of the list.
38 while (it.hasNext())
39 {
40     it.next();
41 }
42
43 // Remove the last element that was retrieved.
44 it.remove();
45
46 // Display the names in the list again. The list
47 // should read Alfonso, David, and Katherine.
48 System.out.println("\nHere are the names now.");
49 System.out.println(nameList);
50 }
51 }

```

Program Output

```

Here are the original names.
[Chris, David, Katherine, Kenny]
Here are the names now.
[Alfonso, David, Katherine]

```

Methods of the `LinkedList` Class

Linked allocation makes it easy to move *forward* in a list; you can easily pass from one element to the next by following the successor link. The Java implementation of a linked list is actually a doubly

linked list. In addition to each element having a link to its successor, every element has a link to its predecessor. The predecessor links make it easy to move *backward* in the list. Also, in addition to keeping a reference to the first element in the list, the Java implementation keeps a reference to the last element of the list. This makes it as easy to locate the last element as it is to locate the first.

To understand why a linked list should maintain predecessor as well as successor links, imagine an application that prints the element of a linked list of length 5 in reverse order. Suppose that the application accesses the elements by index, using the `get()` method of the `List` interface.

```
String [] namesArray = {"Anna", "Bob", "Carlos", "Debby", "Ephraim"};
LinkedList list2 = new LinkedList<>(Arrays.asList(namesArray));

for (int k = list2.size()-1; k >= 0; k--)
{
    System.out.printf("%s ", list2.get(k));
}
```

With each call to `get(k)`, the list would start at index 0 and traverse k links to get to the node at index k . Altogether, this code would require the traversal of $4+3+2+1=10$ links. For a long list, the code would need to traverse a lot of links.

To facilitate the backward and forward traversals of a linked list, and to guarantee quick access to the first and last elements, the `LinkedList` class adds several methods not found in `List`.

Table 18-5 Methods of the `LinkedList<E>` class

<code>void addFirst(E e)</code>	Inserts the specified element at the beginning of this list.
<code>void addLast(E e)</code>	Appends the specified element to the end of this list.
<code>E getFirst()</code>	Returns the first element in this list.
<code>E getLast()</code>	Returns the last element in this list.
<code>E removeFirst()</code>	Removes and returns the first element in this list.
<code>E removeLast()</code>	Removes and returns the last element in this list.

```
Iterator<E>  
descendingIterator()
```

Returns an iterator that traverses the list in reverse order.



The methods of [Table 18-5](#) for getting and removing the first and last elements will throw `NoSuchElementException` if called on an empty list.

Using an Interface Variable to Reference a Collection Object

Take a closer look at the statement in line 13 of [Code Listing 18-5](#):

```
List<String> nameList = new ArrayList<>();
```

Notice that `nameList` is a `List` variable, not an `ArrayList` variable. Because the `ArrayList` class implements the `List` interface, we can use a `List` variable to reference an `ArrayList` object. This is a common practice. By always using interface variables to reference collection objects, our code will be easier to maintain if it becomes necessary to use a different concrete class to create the object.

Either one of the `ArrayList` and `LinkedList` class can be used to create a list. Internally, however, there are differences in the way that these classes handle list operations. Suppose we write a program that at first uses an `ArrayList` to store a large number of `String` objects. We use the following statement to instantiate the `ArrayList`:

```
List<String> myList = new ArrayList<>();
```

Also, assume that in the program we pass the `ArrayList` object as an argument to numerous methods. Each of the methods uses a `List` variable as a parameter, such as the following:

```
public void displayList(List<String> somelist)  
{  
    method code here...  
}
```

Once we put the program into use, we determine that the `ArrayList` is performing slowly because a large number of objects are constantly being inserted in the middle. After doing a little research, we determine that the `LinkedList` class performs faster under these circumstances, and we decide to use it instead of the `ArrayList` class. Because the `LinkedList` class also implements the `List` interface, and we have only used `List` variables to reference the `ArrayList`, making the change is simple. In

the statement that instantiates the list object, we only have to change `ArrayList` to `LinkedList`, as shown here:

```
List<String> myList = new LinkedList<>();
```

Changes in other parts of the program are not necessary because we only used `List` variables to reference the list object.

Checkpoint

1. 18.8 What kind of memory allocation is used by the `ArrayList` class?
2. 18.9 What kind of memory allocation is used by the `LinkedList` class?
3. 18.10 Why is it a common practice to use an interface variable to reference a collection object?
4. 18.11 `ListIterator` is a subinterface of what other JCF interface?
5. 18.12 When it comes to moving through the elements of a collection, describe a difference between an `Iterator` and a `ListIterator`.
6. 18.13 When you need a list iterator, how do you get one?
7. 18.14 When stepping through the elements of a list with a list iterator's `next` method, how do you determine whether there are more elements in the list to visit?
8. 18.15 Which method should you call to determine `previous()` without throwing an exception?
9. 18.16 What element does the `remove` method (in the `Iterator` interface) remove?
10. 18.17 What element does the `set` method (in the `ListIterator` interface) change?
11. 18.18 What is the difference between the *size* and *capacity* of an `ArrayList`?
12. 18.19 Which concrete list class stores its elements in a `linked list`?
13. 18.20 Which concrete list class would be a good choice to use when your application will be storing a large amount of data in a list, and a lot of insertions and/or deletions in the middle of the list are likely to take place?

18.3 Sets

CONCEPT: A set is a collection that lets you quickly find a value in a large amount of data. A set stores elements in an unordered sequence, and does not allow duplicates.

Sometimes when data is stored in a list, you know the exact location of specific items. For example, suppose we have a list that holds the total sales for each month of the year. The list has twelve elements, with the indices 0 through 11. We know that January's sales amount is stored at index 0, February's sales amount is stored at index 1, and so forth. If we want to retrieve the sales amount for December, we know that we have to retrieve the element at index 11. In the list, the sales for month n is stored at index $n-1$.

Sometimes, we need to retrieve a particular element from a collection, but we do not know the location of that element. For example, suppose we have a collection of `Customer` objects, and we need to search for a particular object by the customer's name. `Customer` objects are added to and removed from the collection on a daily basis, and we have no way of knowing where a particular object is stored. If the objects are stored in a list, finding a particular customer means sequentially searching the list until we find the element we are looking for. This could be a time-consuming process if there are a large number of objects in the list.

Perhaps a set is the best type of collection to use in this situation—if fast retrieval is important. Sets are useful when you have a large collection of data, and you must retrieve an element from the collection by searching for its value. The elements in a set are stored in a data structure that is optimized for this type of retrieval.

It is helpful to compare sets to lists. A list will allow you to append an element at the end, add an element at the beginning, or insert an element at a specific position within the list. In contrast, a set element is not associated with a position within the set. Sets are also different from lists in another way: they do not allow duplicate elements. Each element in a set must be unique.

The Set Interface

The `Set` interface extends the `Collection` interface, so it inherits all of the methods specified by `Collection`. (See [Table 18-2](#) for a description of the methods specified by the `Collection` interface.) The `Set` interface does not add any additional methods of its own, but it does place restrictions on many of the methods inherited from `Collection`. These restrictions enforce the requirement that a set cannot contain duplicate elements.

The HashSet Class

The first concrete set class that we will study is the `HashSet` class, which implements the `Set` interface. The `HashSet` class uses hash codes to store its elements in a way that makes searching faster. To understand how a `HashSet` works, first you must understand the concept of a hash code.

A *hash code* is an integer value that is characteristic of an object. You use a *hashing algorithm* to calculate a hash code for an object. There are many different hashing algorithms, ranging from the simple to the complex. Although we will not conduct an in-depth study of hashing algorithms, later in this section we will look at some simple techniques for calculating hash codes.

Hashing algorithms typically use the data that is stored in an object to calculate the object's hash code. For example, if we were to calculate the hash code for a string, we would somehow use the characters making up the string in the calculation. Perhaps we would use each character's Unicode value in a mathematical formula that would give us an integer value.

Ideally, our hashing algorithm would produce different hash codes for different strings. For example, the hash code for "Fred" would be different from the hash code for "Wilma".

In addition to producing different hash codes for different objects, an ideal hashing algorithm should produce the same hash code for objects that are equal in value. For example, suppose we have the following two `String` objects, `str1` and `str2`:

```
String str1 = new String("Fred");  
String str2 = new String("Fred");
```

Although `str1` and `str2` are different objects, they contain equal data and a hashing algorithm should produce the same hash code for both.

In Java, every object has a `hashCode` method that is inherited from the `Object` class. This method returns a hash code for the object. The `hashCode` method that is inherited from the `Object` class, however, is rarely adequate because it returns a value that is based on the object's memory address. The value that is returned from the `Object` class's `hashCode` method is unique for every object, even objects that contain the same data. For this reason, the `Object` class's `hashCode` method must usually be overridden so that it returns the same hash code for equivalent objects.

Hash codes are at the heart of how the `HashSet` class works. When an object is added to a `HashSet`, the `HashSet` calls the object's `hashCode` method to get its hash code. Then it checks to see if an object with the same hash code has already been stored in the set. (Remember, a set does not allow duplicates.) If it finds that an object with the same hash code is already in the set, it performs one more test: it calls the object's `equals` method to see if it reports that the two objects are equal. If they are, the `HashSet` does not store the object in the set.

You might be wondering why a `HashSet` does not simply refuse to store an object as soon as it finds another object with the same hash code already in the set. Earlier, we said that an ideal hashing algorithm would produce different hash codes for different values. That is true, but hashing algorithms are not usually ideal. Occasionally, a hashing algorithm will produce the same hash code for two objects that are not equal. This is called a *collision*.

Because collisions can occur, the `HashSet` performs two tests, one using the `hashCode` method and the other using the `equals` method, to determine whether an object's value is already in the set.

Because collisions sometimes occur with hashing algorithms, it is possible that when a new object is added to a `HashSet` it will have the same hash code as an object already in the set, but it will not be equal to that object (the two objects will contain different data). In this case, the `HashSet` must allow the new object to be added to the set.

One popular way of thinking about hash codes is that they identify “buckets” in the `HashSet`. Each individual hash code has its own bucket, and all objects in the set that have that hash code are stored in that hash code's bucket. So, some buckets will have only one object stored in them, while other buckets will contain multiple objects. In pseudocode, the following is the procedure followed by the `HashSet` class to add a new object:

```
Call the new object's hashCode method to get the object's hash code.
If that hash code's bucket is empty, Then
    Store the object in the bucket for that hash code.
Else
    Use the object's equals method to compare it to each object in the bucket.
    If it is not equal to any other object in the bucket, Then
        Store the object in the bucket.
    Else
        Do not store the new object.
    End If
End If
```

TIP:

Because the `hashCode` and `equals` methods are used together like this, any time you override the `Object` class's `equals` method you should also override the `hashCode` method.

The `HashSet` class has four constructors, which are summarized in [Table 18-6](#). Note that when a `HashSet` is constructed, it has an initial capacity and a load factor. The *initial capacity* is the number of “buckets” that the `HashSet` initially has. The *load factor* is the percentage of the buckets that have to be occupied before the capacity is increased. For example, with a load factor of 0.75, the capacity will be increased when 75 percent of the buckets have a value stored in them.

Table 18-6 `HashSet` constructors

<code>HashSet()</code>	This constructor creates an empty <code>HashSet</code> object with an initial capacity of 16 and a load factor of 0.75.


```
HashSet(Collection<? extends E> c)
```

This constructor creates a `HashSet` object and stores the elements of the collection `c` in it. The initial capacity will be at least enough to hold the elements of `c`, and the load factor will be 0.75. Throws a `NullPointerException` if `c` is a null reference.

```
HashSet(int initialCapacity, float loadFactor)
```

This constructor creates an empty `HashSet` object with the specified initial capacity and load factor.

```
HashSet(int initialCapacity)
```

This constructor creates an empty `HashSet` object with the specified initial capacity and a load factor of 0.75.

Now, let's look at a simple example of how we can store `String` objects in a `HashSet`. The program in [Code Listing 18-6](#) creates a `HashSet` and adds a group of `String` objects to it. Then it attempts to add a duplicate `String` object to the set.

Code Listing 18-6 (HashSetDemo1.java)

```
1 import java.util.*;
2
3 /**
4  * This program demonstrates how to add elements to a HashSet.
5  * It also shows that duplicate elements are not allowed.
6  */
7 public class HashSetDemo1
8 {
9     public static void main(String[] args)
10    {
11        // Create a HashSet to hold String objects.
12        Set<String> fruitSet = new HashSet<>();
13
14        // Add some strings to the set.
15        fruitSet.add("Apple");
16        fruitSet.add("Banana");
17        fruitSet.add("Pear");
18        fruitSet.add("Strawberry");
19
20        // Display the elements in the set.
```

```

21         System.out.println("Here are the elements.");
22         System.out.println(fruitSet);
23
24         // Try to add a duplicate element.
25         System.out.println("\nTrying to add Banana to"
26             + "the set again . . .");
27         if (!fruitSet.add("Banana"))
28         {
29             System.out.println("Banana was not added again.");
30         }
31
32         // Display the elements in the set.
33         System.out.println("\nHere are the elements once more.");
34         System.out.println(fruitSet);
35     }
36 }

```

Program Output

```

Here are the elements.
[Apple, Pear, Strawberry, Banana]
Trying to add Banana to the set again...
Banana was not added again.
Here are the elements once more.
[Apple, Pear, Strawberry, Banana]

```

Take a closer look at lines 27–30:

```

if (!fruitSet.add("Banana"))
{
    System.out.println("Banana was not added again.");
}

```

The `HashSet` class's `add` method does not throw an exception if it is called with a duplicate item. It simply does not add the item. The method does return a boolean value, however, indicating whether the item was added. The statement shown here calls the `add` method, passing "Banana" as an argument. Because "Banana" is already in the set, the method returns false indicating that it was not added again.

The program in [Code Listing 18-7](#) shows another demonstration. It uses the second constructor shown in [Table 18-6](#) to create a `HashSet` from a previously created collection of names, and then gets an iterator to step through the set and print its elements. The program also calls the `HashSet` class's `contains` method to determine whether an object has been stored in the set.

Code Listing 18-7

```

1 import java.util.*;
2
3 /**

```

```

4   This program creates a HashSet with names in it,
5   gets an iterator for the set, and searches the
6   set for names.
7   */
8
9   public class HashSetDemo2
10  {
11      public static void main(String[] args)
12      {
13          // Create a HashSet to hold names.
14          List<String> nameList =
15              Arrays.asList("Chris", "David", "Katherine", "Kenny");
16          Set<String> nameSet = new HashSet<>(nameList);
17
18          // Get an iterator for the set.
19          Iterator<String> it = nameSet.iterator();
20
21          // Display the elements in the set.
22          System.out.println("Here are the names in the set.");
23          while (it.hasNext())
24          {
25              System.out.println(it.next());
26          }
27          System.out.println();
28
29          // Search for "Katherine". We should find this
30          // name in the set.
31          if (nameSet.contains("Katherine"))
32          {
33              System.out.println("Katherine is in the set.");
34          } else
35          {
36              System.out.println("Katherine is NOT in the set.");
37          }
38
39          // Search for "Bethany". We should not find
40          // this name in the set.
41          if (nameSet.contains("Bethany"))
42          {
43              System.out.println("Bethany is in the set.");
44          } else
45          {
46              System.out.println("Bethany is NOT in the set.");
47          }
48      }
49  }

```

Program Output

```

Here are the names in the set.
Chris
Katherine
David
Kenny
Katherine is in the set.

```

Bethany is NOT in the set.

Implementing Your Own hashCode Method

When you design a class, you will need to override the `hashCode` and `equals` methods if you want to store objects of the class in a `HashSet`. Here are three guidelines that you should follow when writing your own `hashCode` and `equals` methods:

1. If you call an object's `hashCode` method more than once during the execution of an application, the method should consistently return the same value. The only exception is when something in the object has changed, which would affect the way the `equals` method makes comparisons.
2. If the `equals` method reports that two objects are equal, then both objects' `hashCode` methods should return the same value.
3. If the `equals` method reports that two objects are *not* equal, it is permissible for both objects to have the same hash code. You should be aware, however, that having too many unequal objects with the same hash code will degrade the performance of an algorithm that uses the hash codes for searching.

The `String` class uses the following formula to calculate a hash code for a string:

$$\text{hashCode} = S_0 * 31^{(n-1)} + S_1 * 31^{(n-2)} + \dots + S_{n-1}$$

In the formula, n is the number of characters in the string, S_0 is the Unicode value of the character at position 0, S_1 is the Unicode value of the character at position 1, and S_{n-1} is the Unicode value of the character at position $n-1$.

The algorithm used by the `String` class is moderately complex. Other classes use a simpler approach. For example, the `Character` class simply uses the character's Unicode value as its hash code. The `Integer` class uses the actual integer value of the object as its hash code.

[Code Listing 18-8](#) shows a `Car` class that we will use to hold a VIN (vehicle identification number) and description of a car. These fields are stored as the `String` objects `vin` and `description`. In this class, we have overridden the `hashCode` method to simply return the hash code of `vin`. This should be adequate because each car has its own unique VIN. We have also overridden the `equals` method to return `true` if both objects have the same VIN.

Code Listing 18-8 (Car.java)

```
1 /**
2  The Car class stores a VIN (Vehicle Identification Number)
3  and a description for a car.
4  */
```

```

5 public class Car
6 {
7     private String vin; // Vehicle Identification Number
8     private String description; // Car description
9
10    /**
11     Constructor
12
13     @param v The VIN for the car.
14     @param desc The description of the car.
15     */
16    public Car(String v, String desc)
17    {
18        vin = v;
19        description = desc;
20    }
21
22    /**
23     getVin method
24
25     @return The car's VIN.
26     */
27    public String getVin()
28    {
29        return vin;
30    }
31
32    /**
33     getDescription method
34
35     @return The car's description.
36     */
37    public String getDescription()
38    {
39        return description;
40    }
41
42    /**
43     toString method
44
45     @return A string containing the VIN and description.
46     */
47    public String toString()
48    {
49        return "VIN:" + vin
50            + "\tDescription:"
51            + description;
52    }
53
54    /**
55     hashCode method
56
57     @return A hash code for this car.
58     */
59    public int hashCode()
60    {
61        return vin.hashCode();

```

```

62     }
63
64     /**
65         equals method
66
67         @param obj Another object to compare this object to.
68         @return true if the two objects are equal, false otherwise.
69     */
70     public boolean equals(Object obj)
71     {
72         // Make sure the other object is a Car.
73         if (obj == null || !(obj instanceof Car) )
74             return false;
75         else
76         {
77             // Both are cars: Get a reference to other
78             // object as a car and check if their VIN
79             // numbers are the same.
80             Car tempCar = (Car) obj;
81             return vin.equalsIgnoreCase(tempCar.vin);
82         }
83     }
84 }

```

The program in [Code Listing 18-9](#) tests the class by adding several `Car` objects to a `HashSet`. Then the `HashSet` is searched for two `Car` objects: one that is in the set, and one that is not.

Code Listing 18-9 (CarHashSet.java)

```

1  import java.util.*;
2
3  /**
4   This program stores Car objects in a HashSet and then
5   searches for various objects.
6  */
7  public class CarHashSet
8  {
9      public static void main(String[] args)
10     {
11         // Create a HashSet to store Car objects.
12         Set<Car> carSet = new HashSet<>();
13
14         // Add some Car objects to the HashSet.
15         carSet.add(new Car("227H54", "1997 Volkswagen"));
16         carSet.add(new Car("448A69", "1965 Mustang"));
17         carSet.add(new Car("453B55", "2007 Porsche"));
18         carSet.add(new Car("177R60", "1980 BMW"));
19
20         // Display the elements in the HashSet.
21         System.out.println("Here are the cars in the set:");
22         for (Car c : carSet)
23         {
24             System.out.println(c);
25         }

```

```

26      System.out.println();
27
28      // Search for a specific car. This one is in the set.
29      Car mustang = new Car("448A69", "1965 Mustang");
30      System.out.println("Searching for" + mustang);
31      if (carSet.contains(mustang))
32      {
33          System.out.println("The Mustang is in the set.");
34      } else
35      {
36          System.out.println("The Mustang is NOT in the set.");
37      }
38
39      // Search for another car. This one is not in the set.
40      Car plymouth = new Car("911C87", "2000 Plymouth");
41      System.out.println("Searching for" + plymouth);
42      if (carSet.contains(plymouth))
43      {
44          System.out.println("The Plymouth is in the set.");
45      } else
46      {
47          System.out.println("The Plymouth is NOT in the set.");
48      }
49  }
50 }

```

Program Output

```

Here are the cars in the set:
VIN: 177R60   Description: 1980 BMW
VIN: 227H54   Description: 1997 Volkswagen
VIN: 448A69   Description: 1965 Mustang
VIN: 453B55   Description: 2007 Porsche
Searching for VIN: 448A69 Description: 1965 Mustang
The Mustang is in the set.
Searching for VIN: 911C87 Description: 2000 Plymouth
The Plymouth is NOT in the set.

```

The LinkedHashSet Class

You might have noticed that in the program output of [Code Listings 18-6](#), [18-7](#), and [18-9](#), the order in which the `HashSet` elements are displayed does not match the order in which they were added to the set. Sets have no concept of internal position of their elements, and they do not keep track of the order of insertion. When you iterate through a `HashSet`, you should not expect the iterator to return elements in the order of their insertion.

If you need to access set elements in the order that they were inserted, you can use the `LinkedHashSet` class instead of `HashSet`. The `LinkedHashSet` class, which extends `HashSet`, keeps an internal linked list that references the elements of the set in the order that they were inserted. This gives you the ability of iterating through the set in that order.

The `LinkedHashSet` class's constructors and methods are identical to those of the `HashSet` method. The program in [Code Listing 18-10](#) creates two sets: a `HashSet` and a `LinkedHashSet`. The same `String` objects are stored in each set.

Code Listing 18-10 (`LinkedHashSetDemo.java`)

```
1 import java.util.*;
2
3 /**
4  * This program demonstrates how a LinkedHashSet
5  * preserves the order in which elements are added.
6  */
7 public class LinkedHashSetDemo
8 {
9     public static void main(String[] args)
10    {
11        // List of fruits to put in hash sets.
12        List<String> fruitList =
13            Arrays.asList("Apple", "Strawberry", "Pear", "Banana");
14
15        // Create the HashSet.
16        Set<String> set1 = new HashSet<>(fruitList);
17
18        // Now create the LinkedHashSet.
19        Set<String> set2 = new LinkedHashSet<>(fruitList);
20
21        // Display the elements in the HashSet.
22        System.out.println("Here are the elements in"
23            + "the HashSet:");
24        System.out.println(set1);
25
26        // Now display the elements in the LinkedHashSet.
27        System.out.println("\nAnd here are the elements"
28            + "in the LinkedHashSet:");
29        System.out.println(set2);
30    }
31 }
```

Program Output

```
Here are the elements in the HashSet:
[Apple, Pear, Strawberry, Banana]
And here are the elements in the LinkedHashSet:
[Apple, Strawberry, Pear, Banana]
```

The `SortedSet` Interface and the `TreeSet` Class

The `Set` interface has a subinterface called `SortedSet` whose methods are shown in [Table 18-7](#).

Table 18-7 SortedSet<E> methods

`Comparator <? super E> comparator()`

This method returns a reference to the comparator used with this `SortedSet`. We discuss comparators later in this section.

`E first()` Returns the first element in this set. This will be the element with the lowest value.

`SortedSet<E> headSet(E toElement)`

Returns a `SortedSet` that serves as a “view” into this `SortedSet`. The elements in the set that are returned are all less in value than the element specified by `toElement`. The elements in the set that are returned are merely references to elements in this set, so any changes made to the elements in the returned set are reflected in this set.

`E last()` Returns the last element in this set. This will be the element with the highest value.

`SortedSet<E> subSet(E fromElement, E toElement)`

Returns a `SortedSet` that serves as a “view” into this `SortedSet`. The elements in the set that are returned are greater than or equal to the element specified by `fromElement`, and less than the element specified by `toElement`. The elements in the set that are returned are merely references to elements in this set, so any changes made to the elements in the returned set are reflected in this set.

`SortedSet<E> tailSet(E fromElement)`

Returns a `SortedSet` that serves as a “view” into this `SortedSet`. The elements in the set that are returned are all greater in value than the element specified by `fromElement`. The elements in the set that are returned are merely references to elements in this set, so any changes made to the elements in the returned set are

reflected in this set.

A `SortedSet` is a set in which the elements are sorted according to their natural order. A class of objects has a *natural order* if it is possible to determine for any two objects of the class whether one object is equal to, or “less than,” or “greater than” the other. In this context, *less than* and *greater than* refers to which element should precede the other on the natural ordering of that class. For example, the natural order of numbers is to order them in increasing order. For strings, the natural order is alphabetic order.

A class imposes a natural order on its objects by implementing the `Comparable` interface. We looked at this interface in [Chapter 16](#) when we were studying generics. Recall that the `Comparable` interface specifies a method named `compareTo`, which is used to compare the calling object with another object that is passed as an argument to the method. This is how the `compareTo` method works:

- If the calling object is less than the object passed as an argument, the method returns a negative number.
- If the calling object is equal to the object passed as an argument, the method returns 0.
- If the calling object is greater than the object passed as an argument, the method returns a positive number.

There is a JCF class called `TreeSet` that implements the `SortedSet` interface. If you create a `TreeSet` using that class’s no-arg constructor, the tree set will sort objects added to it by their natural order. The program in [Code Listing 18-11](#) demonstrates this by creating a `TreeSet` and using it to sort strings. This works because the `String` class implements the `Comparable` interface.

Code Listing 18-11 (`TreeSetDemo1.java`)

```
1 import java.util.*;
2
3 /**
4  This program demonstrates how a TreeSet sorts
5  its elements in ascending order.
6  */
7 public class TreeSetDemo1
8 {
9     public static void main(String[] args)
10    {
11        // Create a TreeSet and store some values in it.
12        SortedSet<String> mySet = new TreeSet<>();
13        mySet.add("Pear");
14        mySet.add("Apple");
15        mySet.add("Strawberry");
16        mySet.add("Banana");
17
18        // Display the elements in the TreeSet.
```

```

19         System.out.println("Here are the TreeSet elements"
20                               + "in ascending order:");
21         System.out.println(mySet);
22
23         // Add a new element to the TreeSet.
24         System.out.println("\nAdding Blueberry to the set.");
25         mySet.add("Blueberry");
26
27         // Display the elements again.
28         System.out.println("\nHere are the TreeSet elements"
29                               + "again:");
30         System.out.println(mySet);
31     }
32 }

```

Program Output

```

Here are the TreeSet elements in ascending order:
[Apple, Banana, Pear, Strawberry]
Adding Blueberry to the set.
Here are the TreeSet elements again:
[Apple, Banana, Blueberry, Pear, Strawberry]

```

Using a Comparator to Compare Objects

If the objects that you wish to store in a `TreeSet` do not implement the `Comparable` interface, you can still store them in the `TreeSet` by creating a comparator. A *comparator* is an object that can compare two other objects and determine whether one is greater than, less than, or equal to the other. A comparator implements the `Comparator` interface, which is in the `java.util` package. The `Comparator` interface has two abstract methods:

```

public interface Comparator<T>
{
    public int compare(T obj1, T obj2);
    public boolean equals(Object obj);
}

```

NOTE:

Beginning with Java 8, the `Comparator` interface contains other methods besides `compare` and `equals`.

[Table 18-8](#) describes the two abstract methods specified by the `Comparator` interface.

Table 18-8 Abstract methods of the `Comparator<T>` interface

<pre>public int compare(T obj1, T obj2)</pre>	This method compares <code>obj1</code> with <code>obj2</code> and returns 0 if the two objects are equal, a negative value if <code>obj1</code> is less than <code>obj2</code> , and a positive value if <code>obj1</code> is greater than <code>obj2</code> .
<pre>public boolean equals(Object obj)</pre>	This method determines whether this comparator is equal to the object <code>obj</code> . If the two are equal, it returns <code>true</code> .

[Code Listing 18-12](#) shows an example of a `Comparator` class that can compare `Car` objects. The `compare` method accepts references to two `Car` objects. It gets each car's VIN, which is a `String`, and then uses the `String` class's `compareToIgnoreCase` method to compare the two VINs.

Code Listing 18-12 (`CarComparator.java`)

```
1 import java.util.Comparator;  
2  
3 public class CarComparator  
4     implements Comparator<Car>  
5 {  
6     public int compare(Car car1, Car car2)  
7     {  
8         // Get the two cars' VINs.  
9         String vin1 = car1.getVin();  
10        String vin2 = car2.getVin();  
11  
12        // Compare the VINs and return the  
13        // result of the comparison.  
14        return vin1.compareToIgnoreCase(vin2);  
15    }  
16 }
```

NOTE:

We did not provide an `equals` method in the `CarComparator` class, even though it is specified in the `Comparator` interface. The `Object` class provides an `equals` method, so we will not get an error when we compile this class. If we wanted to be able to compare objects of this class with other comparators, we would have provided an `equals` method.

The program in [Code Listing 18-13](#) demonstrates how to create a `TreeSet` that uses a `CarComparator` to compare its elements. If you pass a comparator object as an argument to the `TreeSet` constructor, the `TreeSet` will use the comparator to perform comparisons.

Code Listing 18-13 (`TreeSetDemo2.java`)

```
1 import java.util.*;
2
3 /**
4  * This program demonstrates how a TreeSet can use
5  * a Comparator to sort its elements.
6  */
7 public class TreeSetDemo2
8 {
9     public static void main(String[] args)
10    {
11        // Create a TreeSet and pass an instance of
12        // CarComparator to it.
13        SortedSet<Car> carSet
14            = new TreeSet<>(new CarComparator());
15
16        // Add some Car objects to the TreeSet.
17        carSet.add(new Car("227H54", "1997 Volkswagen"));
18        carSet.add(new Car("453B55", "2007 Porsche"));
19        carSet.add(new Car("177R60", "1980 BMW"));
20        carSet.add(new Car("448A69", "1965 Mustang"));
21
22        // Display the elements in the TreeSet.
23        System.out.println("Here are the cars sorted in"
24            + " order of their VINs:");
25        for (Car car : carSet)
26        {
27            System.out.println(car);
28        }
29    }
30 }
```

Program Output

```
Here are the cars sorted in order of their VINs:
VIN: 177R60   Description: 1980 BMW
VIN: 227H54   Description: 1997 Volkswagen
VIN: 448A69   Description: 1965 Mustang
VIN: 453B55   Description: 2007 Porsche
```

[Table 18-9](#) summarizes some of the most frequently used `TreeSet` constructors.

Table 18-9 `TreeSet<E>` constructors

<code>TreeSet()</code>	This constructor creates an empty <code>TreeSet</code> object. Any objects added to the <code>TreeSet</code> must implement the <code>Comparable</code> interface.
<code>TreeSet(Collection<? extends E> c)</code>	This constructor creates a <code>TreeSet</code> object and stores the elements of the collection <code>c</code> in it. The elements of <code>c</code> must implement the <code>Comparable</code> interface.
<code>TreeSet(Comparator<? extends E> c)</code>	This constructor creates a <code>TreeSet</code> object and uses <code>c</code> , a comparator, to compare its elements.



Checkpoint

1. 18.21 What is a hash code?
2. 18.22 What method do you call to get an object's hash code?
3. 18.23 In hashing, what is a collision?
4. 18.24 What does the `HashSet` class's `add` method do when you try to add a duplicate item?
5. 18.25 Any time you override the `Object` class's `equals` method, what other method should you also override?
6. 18.26 Which concrete set class should you use if you need to access its elements in the order that they were inserted?
7. 18.27 Which concrete set class should you use if you need to retrieve its elements in their natural order?
8. 18.28 What is a comparator? What interface must it implement?

18.4 Maps

CONCEPT: Each element in a map has two parts: a key and a value. Each key is associated with a specific value, and can be used to locate that value.

Each element in a *map* has two parts: a key and a value. Each key is associated with a specific value, and can be used to locate that value. For example, suppose we wish to store employee names in a map, and we want to search for specific employees by their ID numbers. Each time we store an employee's name in the map as a value, we also store that employee's ID number as a key. The map creates an association between the value and the key, and allows us to retrieve the name by searching for the ID number. The elements in a map are referred to as *mappings*.

There are a number of similarities between sets and maps. In fact, you can think of a map as storing a set of keys, with each key in the map's *key set* being associated with a value. Just as the JCF has a `Set` interface, there is a `Map` interface. Just as the `HashSet` and `LinkedHashSet` classes implement the `Set` interface; there are `HashMap` and `LinkedHashMap` classes that implement the `Map` interface. Just as `Set` has a `SortedSet` subinterface; `Map` has a `SortedMap` subinterface. Finally, there is a `TreeMap` class that implements the `SortedMap` interface, mirroring the way that `TreeSet` implements the `SortedSet` interface.

The map interface is generic, and is parameterized by two type variables `K` and `V`. These respectively stand for the classes representing the key and value objects. [Table 18-10](#) lists some of the more commonly used methods specified by the `Map<K, V>` interface.

Table 18-10 Methods specified by the `Map<K, V>` interface

<code>void clear()</code>	Clears the map of all its elements. (This is an optional method. If the map does not support it, it throws an <code>UnsupportedOperationException</code> .)

<code>boolean containsKey(Object key)</code>	

	Returns <code>true</code> if this map contains a value that is associated with the specified key.

<code>boolean containsValue(Object value)</code>	

Returns `true` if this map contains one or more keys mapped to the specified value.

```
Set< Map.Entry<K,V> > entrySet()
```

Returns a set that serves as a “view” into this map, containing all of the map’s mappings. The elements in the set that is returned all implement the `Map.Entry` interface. (`Entry` is defined as an inner interface in the `Map` - interface.) The elements in the set that is returned are merely references to elements in this set, so any changes made to the elements in the returned set are reflected in this map.

```
void forEach(BiConsumer<? super K,? super V> action)
```

Calls the action method on each key and value pair in this map.

```
V get(Object key)
```

Returns the value that is associated with the specified key. If the key does not map to a value, the method returns `null`.

```
boolean  
isEmpty()
```

Returns `true` if this map contains no mappings.

```
Set<K>  
keySet()
```

Returns a set that contains all of the keys in this map. The set that is returned serves as a “view” into this map. The elements in the set are merely references to the keys in this map, so any changes made to the elements in the returned set are reflected in this map.

```
V put(K key,  
V value)
```

This method puts a mapping in the map. The specified key is mapped to the specified value. If an existing mapping for this key exists, it is replaced by this mapping. The method returns the previous value that was associated with this key, or `null` if there was no value associated with this key previously.

```
V  
remove(Object  
key)
```

If a mapping for the specified key exists, it is removed. A reference to the value that the key was previously mapped to is returned.

`int size()` Returns the number of mappings in this map.

`Collection<V>
values()` Returns a collection that contains all of the values in this map. The collection that is returned serves as a “view” into this map. The elements in the collection are merely references to the values in this map, so any changes made to the elements in the returned collection are reflected in this map.



The `HashMap` Class

The `HashMap` Class

The `HashMap` class is a concrete class that implements the `Map` interface. Like a `HashSet`, a `HashMap` uses hash codes to store and retrieve data. Specifically, the `HashMap` class stores keys by their hash codes. [Table 18-11](#) summarizes the `HashMap` constructors. Some of these constructors will look familiar because they are similar to the `HashSet` constructors. Because the `HashMap` class uses hash codes to store keys, it provides constructors that allow you to specify initial capacity and load factor. (See the previous section on the `HashSet` class for a discussion of initial capacity and load factor.)

Table 18-11 `HashMap<K, V>` constructors

`HashMap()` This constructor creates an empty `HashMap` object with an initial capacity of 16 and a load factor of 0.75.

`HashMap(Map<? Extends K, ? extends V> m)`

This constructor creates a `HashMap` object and stores the mappings of the map `m` in it. The initial capacity will be at least enough to hold the elements of `m`, and the load factor will be 0.75. Throws a `NullPointerException` if `m` is a `null` reference.

`HashMap(int initialCapacity, float loadFactor)`

This constructor creates an empty `HashMap` object with the specified initial capacity and load factor.

```
HashMap(int initialCapacity)
```

This constructor creates an empty `HashMap` object with the specified initial capacity and a load factor of 0.75.

The program in [Code Listing 18-14](#) demonstrates a simple `HashMap` for storing `Car` objects. This application uses VINs (vehicle identification numbers) as keys, and `Car` objects as values. As a result, the `HashMap` will allow us to search for a specific `Car` object by specifying its VIN.

Code Listing 18-14 (CarHashMap1.java)

```
1 import java.util.*;
2
3 /**
4  * This program stores mappings in a HashMap and then
5  * searches for various objects.
6  */
7
8 public class CarHashMap1
9 {
10     public static void main(String[] args)
11     {
12         // Create a HashMap to store Car objects.
13         Map<String, Car> carMap = new HashMap<>();
14
15         // Create some Car objects.
16         Car vw = new Car("227H54", "1997 Volkswagen");
17         Car mustang = new Car("448A69", "1965 Mustang");
18         Car porsche = new Car("453B55", "2007 Porsche");
19         Car bmw = new Car("177R60", "1980 BMW");
20
21         // Put some mappings into the HashMap. In each
22         // mapping, the car's VIN is the key and the
23         // Car object containing that VIN is the value.
24         carMap.put(vw.getVin(), vw);
25         carMap.put(mustang.getVin(), mustang);
26         carMap.put(porsche.getVin(), porsche);
27         carMap.put(bmw.getVin(), bmw);
28
29         // Search for the Mustang by its VIN.
30         System.out.println("\nSearching for the car with" +
31                             "VIN" + mustang.getVin());
32         Car foundCar = carMap.get(mustang.getVin());
33
34         // If the car was found, display it.
35         if (foundCar != null)
```

```

36         System.out.println(foundCar);
37     else
38         System.out.println("The Mustang is NOT in the set.");
39
40     // Search for another VIN. This one is not in the set.
41     System.out.println("\nSearching for the car with" +
42         "VIN 911C87");
43     foundCar = carMap.get("911C87");
44
45     // If the car was found display it.
46     if (foundCar != null)
47         System.out.println(foundCar);
48     else
49         System.out.println("That car is NOT in the set.");
50 }
51 }

```

Program Output

```

Searching for the car with VIN 448A69
VIN: 448A69 Description: 1965 Mustang
Searching for the car with VIN 911C87
That car is NOT in the set.

```

The statement in line 13 creates the `HashMap`:

```
Map<String, Car> carMap = new HashMap<>();
```

Note that the `Map<K, V>` interface has been specialized to `Map<String, Car>` because `String` is the data type of the keys and `Car` is the data type of the values.

The statements in lines 16 through 19 create four `Car` objects: one for a VW, one for a Mustang, one for a Porsche, and one for a BMW. Then, the statements in lines 24 through 27 use the `put` method to create several mappings in the `HashMap`. For example, look at the statement in line 24:

```
carMap.put(vw.getVin(), vw);
```

The `put` method puts a mapping into the `HashMap`. The first argument is the key, and the second argument is the value to associate with the key. In this program, we are creating mappings between VINs and `Car` objects. In each mapping, the key is a VIN and the value is the `Car` object that goes with that VIN. In this statement, we are creating a mapping that associates the VW's VIN with the VW's `Car` object. The statements in lines 24 through 27 create similar mappings for the other `Car` objects.

Line 32 shows how we can search for a specific `Car` object:

```
Car foundCar = carMap.get(mustang.getVin());
```

This statement creates a `Car` reference variable, `foundCar`, and assigns to it the value returned from the `carMap.get` method. The `get` method accepts a key as its argument and it returns the value that

the key is mapped to. In this statement, we are passing the Mustang's VIN as the argument. We are basically saying to the `HashMap` "Give me the `Car` object that is associated with this VIN." If the key is found in the `HashMap`, the method returns a reference to the value that is associated with that key. Because the Mustang's VIN is indeed found as a key in the `HashMap`, this statement returns a reference to the `mustang` object and assigns to `foundCar`.

In line 43 we aren't so lucky. Here we pass to the `get` method a VIN that does not exist in the `HashMap`:

```
foundCar = carMap.get("911C87");
```

As a result, the method will return `null`.

Unlike the `Set<E>` and `List<E>` interfaces, `Map<K, V>` does not extend the `Collection<E>` interface. And unlike the other collections, maps do not provide iterators. One way to iterate over the contents of a map is to convert its mappings to some other type of collection. For example, you can use the `keySet()` method to get a set that contains all of the keys in a map, and you can use the `values()` method to get a collection that contains all of the map's values. Then you can iterate over the resulting set or collection. [Code Listing 18-15](#) shows an example.

Code Listing 18-15 (CarHashMap2.java)

```
1 import java.util.*;
2
3 /**
4     This program retrieves a set of keys and a
5     collection of values from a HashMap.
6 */
7
8 public class CarHashMap2
9 {
10     public static void main(String[] args)
11     {
12         // Create a HashMap to store Car objects.
13         Map<String, Car> carMap = new HashMap<>();
14
15         // Create some Car objects.
16         Car vw = new Car("227H54", "1997 Volkswagen");
17         Car mustang = new Car("448A69", "1965 Mustang");
18         Car porsche = new Car("453B55", "2007 Porsche");
19         Car bmw = new Car("177R60", "1980 BMW");
20
21         // Put some mappings into the HashMap. In each
22         // mapping, the car's VIN is the key and the
23         // Car object containing that VIN is the value.
24         carMap.put(vw.getVin(), vw);
25         carMap.put(mustang.getVin(), mustang);
26         carMap.put(porsche.getVin(), porsche);
27         carMap.put(bmw.getVin(), bmw);
28
29         // Get a set containing the keys in this map.
```

```

30     Set<String> keys = carMap.keySet();
31
32     // Iterate through the keys, printing each one.
33     System.out.println("Here are the keys:");
34     for (String k : keys)
35         System.out.println(k);
36
37     // Get a collection containing the values.
38     Collection<Car> values = carMap.values();
39
40     // Iterate through the values, printing each one.
41     System.out.println("\nHere are the values:");
42     for (Car c : values)
43         System.out.println(c);
44 }
45 }

```

Program Output

Here are the keys:

```

177R60
227H54
448A69
453B55

```

Here are the values:

```

VIN: 177R60      Description: 1980 BMW
VIN: 227H54      Description: 1997 Volkswagen
VIN: 448A69      Description: 1965 Mustang
VIN: 453B55      Description: 2007 Porsche

```

NOTE:

The set and the collection that are returned from the `keySet` and `values` methods serve as “views” into the map. The elements in the set and collection are merely references to the entries in the map, so any changes made to the elements in the set or collection are reflected in the map, and vice versa.

In [Code Listing 18-15](#), the keys and values of the map are retrieved separately. You can also retrieve a single set that contains all the key-value pairs stored in the map. This is done by the map’s `entrySet` method.

The `entrySet` method returns a set of `Map.Entry` objects. `Entry` is an inner interface in the `Map` interface. Two of the methods specified by `Map.Entry` are `getKey` and `getValue`. The `getKey` method returns the key and `getValue` returns the value. The program in [Code Listing 18-16](#) demonstrates the `entrySet` method and how to access the keys and values of the set elements it returns.

Code Listing 18-16 (CarHashMap3.java)

```
1 import java.util.*;
2
3 /**
4     This program retrieves the mappings from a HashMap
5     as a Set of Map.Entry objects.
6 */
7
8 public class CarHashMap3
9 {
10     public static void main(String[] args)
11     {
12         // Create a HashMap to store Car objects.
13         Map<String, Car> carMap = new HashMap<>();
14
15         // Create some Car objects.
16         Car vw = new Car("227H54", "1997 Volkswagen");
17         Car mustang = new Car("448A69", "1965 Mustang");
18         Car porsche = new Car("453B55", "2007 Porsche");
19         Car bmw = new Car("177R60", "1980 BMW");
20
21         // Put some mappings into the HashMap. In each
22         // mapping, the car's VIN is the key and the
23         // Car object containing that VIN is the value.
24         carMap.put(vw.getVin(), vw);
25         carMap.put(mustang.getVin(), mustang);
26         carMap.put(porsche.getVin(), porsche);
27         carMap.put(bmw.getVin(), bmw);
28
29         // Get a set containing the mappings in this map.
30         Set<Map.Entry<String, Car>> cars = carMap.entrySet();
31
32         // Iterate through the mappings, printing each one.
33         System.out.println("Here are the mappings:");
34         for (Map.Entry<String, Car> entry : cars)
35         {
36             System.out.println("Key =" + entry.getKey());
37             System.out.println("Value =" + entry.getValue());
38             System.out.println();
39         }
40     }
41 }
```

Program Output

```
Here are the mappings:
Key = 177R60
Value = VIN: 177R60      Description: 1980 BMW
Key = 227H54
Value = VIN: 227H54      Description: 1997 Volkswagen
Key = 448A69
Value = VIN: 448A69      Description: 1965 Mustang
```

Key = 453B55
Value = VIN: 453B55 Description: 2007 Porsche

NOTE:

The set that is returned from the `entrySet()` method also serves as a “view” into the map. The elements in the set are merely references to the entries in the map, so any changes made to the elements in the set are reflected in the map, and vice versa.

Java 8 added to the `Map<K, V>` interface a default method

```
void forEach(BiConsumer<? super K, ? super V> action)
```

that can be used to invoke an action on each key-value pair stored in a map. The action must implement the `BiConsumer<? super K, ? super V>` interface. This means that the `action` function must have a type that is appropriately related to the type $(K, V) \rightarrow \text{void}$; its parameters are a key and a value, and its return type is `void`. As usual, we can use a lambda expression for the action. If we adopt this technique, Lines 29-39 of [Code Listing 18-16](#) can be replaced with:

```
System.out.println("Here are the mappings:");
carMap.forEach( (key, value) ->
    {
        System.out.println("Key = " + key);
        System.out.println("Value = " + value);
        System.out.println();
    }
);
```

The LinkedHashMap Class

`HashMap` objects, like sets, have no concept of order for their elements. You will notice, for example, that the order in which the mappings are displayed in [Code Listing 18-16](#) does not match the order in which they were put into the `HashMap`.

The `LinkedHashMap` class, which extends `HashMap`, allows you to access its mappings in one of two orders:

- Insertion order, which is the order in which the mappings were inserted into the map
- Access order, which is the order in which the mappings were last accessed

A `LinkedHashMap` keeps an internal linked list that references mappings. If you are using the map in insertion order mode, the linked list references the mappings in the order that they were inserted. If you are using the map in access order mode, the linked list references the mappings in the order that they were accessed with either `put` or `get`.

In access order mode, the mappings that have been accessed the most are placed at the end of the linked list, and the mappings that have been accessed the least are placed at the beginning. This mode is useful in applications that need to know which data items are being accessed the most or the least.

You indicate which sorting order you prefer with the appropriate `LinkedHashMap` constructor. [Table 18-12](#) summarizes them.

Table 18-12 `LinkedHashMap` constructors

<code>LinkedHashMap()</code>	This constructor creates an empty <code>LinkedHashMap</code> using insertion-order mode, an initial capacity of 16, and a load factor of 0.75.

<code>LinkedHashMap(int initialCapacity)</code>	

	This constructor creates an empty <code>LinkedHashMap</code> using insertion-order mode, an initial capacity of the specified amount, and a load factor of 0.75.

<code>LinkedHashMap(int initialCapacity, float loadFactor)</code>	

	This constructor creates an empty <code>LinkedHashMap</code> using insertion-order mode and the specified values for initial capacity and a load factor.

<code>LinkedHashMap(int initialCapacity, float loadFactor, boolean accessOrder)</code>	

	This constructor creates an empty <code>LinkedHashMap</code> using the specified values for initial capacity and a load factor. If <code>accessOrder</code> is true, the map will use access-order mode. If <code>accessOrder</code> is false, the map will use insertion-order mode.

<code>LinkedHashMap(Map<? extends K, ? extends V> m)</code>	

	This constructor creates a <code>LinkedHashMap</code> and stores the mappings of the map <code>m</code> in it. The initial capacity will be at least enough to hold the

elements of `m`, and the load factor will be 0.75. It throws a `NullPointerException` if `m` is a null reference.

The program in [Code Listing 18-17](#) is a modification of [Code Listing 18-16](#). In this version, a `LinkedHashMap` is used in insertion-ordering mode, and the `forEach()` method is used to display the mappings. In the program output, the mappings are displayed in the order that they were inserted into the map.

Code Listing 18-17 (CarHashMap4.java)

```
1 import java.util.*;
2
3 /**
4     This program retrieves the mappings from a HashMap
5     by using the forEach method.
6 */
7
8 public class CarHashMap4
9 {
10     public static void main(String[] args)
11     {
12         // Create a HashMap to store Car objects.
13         Map<String, Car> carMap = new LinkedHashMap<>();
14
15         // Create some Car objects.
16         Car vw = new Car("227H54", "1997 Volkswagen");
17         Car mustang = new Car("448A69", "1965 Mustang");
18         Car porsche = new Car("453B55", "2007 Porsche");
19         Car bmw = new Car("177R60", "1980 BMW");
20
21         // Put some mappings into the HashMap. In each
22         // mapping, the car's VIN is the key and the
23         // Car object containing that VIN is the value.
24         carMap.put(vw.getVin(), vw);
25         carMap.put(mustang.getVin(), mustang);
26         carMap.put(porsche.getVin(), porsche);
27         carMap.put(bmw.getVin(), bmw);
28
29         // Iterate through the mappings, printing each one.
30         System.out.println("Here are the mappings:");
31         carMap.forEach( (key, value) ->
32             {
33                 System.out.println("Key =" + key);
34                 System.out.println("Value =" + value);
35                 System.out.println();
36             }
37         );
38     }
```

Program Output

Here are the mappings:
Key = 227H54
Value = VIN: 227H54 Description: 1997 Volkswagen
Key = 448A69
Value = VIN: 448A69 Description: 1965 Mustang
Key = 453B55
Value = VIN: 453B55 Description: 2007 Porsche
Key = 177R60
Value = VIN: 177R60 Description: 1980 BMW

The SortedMap Interface and the TreeMap Class

`SortedMap<K, V>` is a subinterface of `Map<K, V>` that allows maps to store mappings by key order. The `TreeMap` class implements `SortedMap`. [Table 18-13](#) lists some of its constructors.

Table 18-13 `TreeMap<K, V>` constructors

<code>TreeMap()</code>	This constructor creates a tree map that orders keys by their natural order. The key type <code>K</code> must implement the <code>Comparable</code> interface.

<code>TreeMap(Comparator<? super K> c)</code>	This constructor creates a tree map that uses the comparator <code>c</code> to order the keys.

Using `TreeMap` is no different from using the other types of maps. For example, we can replace Line 13 of [Code Listing 18-17](#) with:

```
SortedMap<String, Car> carMap = new TreeMap<>();
```

The program will create a tree map that will order its entries in alphabetical order of the VIN number. If we run the resulting program, the output will be as shown below.

Program Output for `TreeMap` version of Code Listing 18-17

Here are the mappings:
Key = 177R60
Value = VIN: 177R60 Description: 1980 BMW

Key = 227H54
Value = VIN: 227H54 Description: 1997 Volkswagen

Key = 448A69
Value = VIN: 448A69 Description: 1965 Mustang

Key = 453B55
Value = VIN: 453B55 Description: 2007 Porsche



1. 18.29 Each element in a map is comprised of what two parts?
2. 18.30 What part of a mapping does the `HashMap` store by its hash code?
3. 18.31 The `LinkedHashMap` class allows you to retrieve mappings in what two orders?
4. 18.32 How are the keys in a `SortedMap` sorted?

18.5 The Collections Class

CONCEPT: The `Collections` class provides numerous static methods for performing algorithms and utility operations on collections.

The JCF has a class named `Collections` which provides numerous static methods that perform useful operations on collections. Note that the name of this class, `Collections`, is plural, and should not be confused with the `Collection` interface.

Many of the methods in the `Collections` class perform operations on lists. There are methods for copying, sorting, searching, randomly shuffling, filling with a value, and many other procedures.

[Table 18-14](#) describes several of the `Collections` methods.

Table 18-14 Some of the static methods in the `Collections` class

```
static <T> int binarySearch(List<? extends Comparable<? super T>> list,
T key)
```

Performs a binary search for the element specified by `key` in the list specified by `list`.

Note that the elements in the list must implement the `Comparable` interface and be sorted according to their natural order. If the element is found, the method returns its index. Otherwise it returns a negative number.

```
static <T> int binarySearch(List<? extends T> list, T key, Comparator<?
super T> c)
```

Performs a binary search for the element specified by `key` in the list specified by `list`.

This version of the method uses the comparator `c` to make comparisons. If the element is found, the method returns its index. Otherwise, it returns a negative number. The list must be sorted in comparator order before it is passed to this method.

```
static boolean disjoint (Collection<?> c1, Collection<?> c2)
```

Returns `true` if the collections `c1` and `c2` have no elements in common.

```
static <T extends Object & Comparable<? super T>> T max(Collection<? extends T> coll)
```

Returns the maximum element in the collection `coll`. This method requires that all elements of `coll` implement the `Comparable` interface. There is a similar method `min` that returns the minimum.

```
static <T> T max(Collection<? extends T>coll, Comparator<? super T> comp)
```

Returns the maximum element in the collection `coll`. This method uses the comparator `comp` to compare the elements. There is a similar method `min` that returns the minimum.

```
static void shuffle(List<?> list)
```

This method rearranges the elements of `list` in a random order.

```
public static <T extends Comparable<? super T>> void sort(List<T> list)
```

Sorts the elements in `list` by their natural order. This method requires that the elements of `list` implement the `Comparable` interface.

```
public static <T> void sort(List<T> list, Comparator<? super T> c)
```

Sorts the elements in `list` according to the order imposed by the comparator `c`.



In the Spotlight:

Using Anonymous Subclasses to Initialize Collections

Normally, you create anonymous classes from interfaces or abstract superclasses because you want to override their methods. You can also create an anonymous class from a concrete class and override some of its methods. In particular, you can do this with sets, lists, and maps. For example, here is the syntax for creating an anonymous subclass of `HashMap`:

```
Map<String, Integer> myMap = new HashMap<String, Integer>()
{
    // optional initialization block
    {
        // initialization code
    }
    // optional method overrides
};
```

As the example syntax shows, you can have a block of code to initialize the instance that is being created from the anonymous class. If you are instantiating a set, list, or map, you can use the initialization block to add elements. Here is an example that initializes a map and prints its contents.

```
Map<String, Integer> myMap = new HashMap<String, Integer>()
{
    {
        put("Mary", 4);
        put("Sam", 3);
    }
};
// Print the map
System.out.println(myMap);
```

18.6 The Stream API

CONCEPT: A stream is an object that permits a pipeline of operations to be performed on a sequence of elements drawn from a source.

A *stream* is an object that permits a pipeline of operations to be performed on a sequence of elements drawn from a source. The *source* of elements may be an array, a collection, or a function that is able to generate an element each time it is called. By a *pipeline* of operations, we mean a sequence of operations where the output of one operation becomes the input to the next operation in the sequence.

The operations in a stream pipeline may be of two kinds. An *intermediate* operation acts on the stream and turns it into another stream with the same source: this resulting stream may then be acted on by the next operation in the pipeline. A *terminal* operation, also called a *reduction*, acts on the stream and yields a non-stream result. A stream pipeline always begins with a source of elements, followed by zero or more intermediate operations, and ends with a single terminal operation.

The Stream Package

The `java.util.stream` package defines a number of interfaces and classes. Among them is `Stream<T>`, a generic interface that defines the methods supported by streams whose elements are objects of type `T`. Unlike the collection API, which only supports reference types and requires primitive types to be boxed, the stream API has direct support for streams of primitive types through the `IntStream`, `DoubleStream`, and `LongStream` interfaces. There is also a `Collectors` class that has static methods for implementing useful reduction operations that accumulate stream elements into different type of collections or summary values.

Creating a Stream

There are many ways to create a stream. The `Stream` interfaces, the `Arrays` class, the `Collection` interface, and the `Random` class, all have methods that can be used to create streams.

1. The `Arrays` class has several static `stream()` methods that take an array parameter and return a stream whose source is the array:

```
static Stream<T> stream(T[] array)
static DoubleStream stream(double[] array)
static IntStream stream(int[] array)
static LongStream stream(long[] array)
```

The first of these methods is used to obtain a stream from an array whose elements are reference types. The other three `stream()` methods return streams of primitive types.

2. The `Collection<E>` interface defines an instance method

```
Stream<E> stream()
```

that returns a stream whose source is the collection.

3. All stream interfaces have a static `of( )` method with a var-args parameter that lets you create a stream from a specified list of values:

```
static Stream<T> of(T ... values)
static IntStream of(int ... values)
static DoubleStream of(double ... values)
static LongStream of(long ... values)
```

4. All Stream interfaces have a static `generate()` method that creates a potentially infinite stream of elements by repeatedly calling on a *supplier* function:

```
static Stream<T> generate(Supplier<T> s)
static IntStream generate(IntSupplier s)
static DoubleStream generate(DoubleSupplier s)
static LongStream generate(LongSupplier s)
```

5. Each of the `IntStream` and `LongStream` interfaces has two static methods that create a stream from a sequence of consecutive integers beginning at `m` and ending at `n`:

```
static IntStream range(int m, int n)           // n excluded
static IntStream rangeClosed(int m, int n)     // n included
static LongStream range(long m, long n)        // n excluded
static LongStream rangeClosed(long m, long n)  // n included
```

The `rangeClosed()` methods include the upper bound `n`, while the `range()` methods exclude it.

6. Objects of the `java.util.Random` class have methods that can return `IntStream`, `DoubleStream`, and `LongStream` streams of random numbers. As an example, we list the following methods of `Random` that return an `IntStream` object:

```
IntStream ints()
IntStream ints(int m, int n)
IntStream ints(long count)
IntStream ints(long count, int m, int n)
```

The `ints()` method returns a stream that can generate infinitely many random integers. The `ints(m,n)` method returns a stream that can generate infinitely many integers bounded below by `m` and above by `n`. The versions of `ints()` that take a `count` parameter are similar, except that they generate a finite stream of `count` elements.

In addition to these methods for creating a stream, the stream interfaces define many intermediate operations that transform an existing stream into another stream. We will look at those later.

Terminal Operations

A *terminal* operation, also called a *reduction*, operates on a stream and produces a non-stream result. The result produced by a reduction is often an aggregate of some type, or a summary of some characteristics of the stream elements. Some terminal operations may perform actions on stream elements (such as printing or updating other program variables) without returning a value.

One of the simplest reduction operators is `long count()`, which returns the number of elements in a stream. The following code creates a stream of string elements `stooges` and invokes the reduction operator `long()` on it.

```
Stream<String> stooges = Stream.of("Larry", "Curly", "Moe");
System.out.println(stooges.count());
```

[Table 18-15](#) lists some of the simpler stream reduction operators.

Table 18-15 Stream reduction operators

<code>long count()</code>	Returns the number of elements in this stream.

<code>void forEach(Consumer<? super T> action)</code>

Calls the specified action for each element of this stream.

<code>Optional<T> max(Comparator<? super T> c)</code>

Returns the maximum of all elements in this stream according to the comparator <code>c</code> .

<code>Optional<T> min(Comparator<? super T> c)</code>

Returns the minimum of all elements in this stream according to the comparator <code>c</code> .

<code>boolean noneMatch(Predicate<? super T> p)</code>

Returns `true` if no elements of this stream satisfy the predicate `p`, `false` otherwise.

```
boolean allMatch(Predicate<? super T> p)
```

Returns `true` if all elements of this stream satisfy the predicate `p`, `false` otherwise.

```
boolean anyMatch(Predicate<? super T> p)
```

Returns `true` if at least one element of this stream satisfies the predicate `p`, `false` otherwise.

```
<A> A[] toArray(IntFunction<A[]> generator)
```

Returns an array containing the elements of this stream. The array is created using a function that takes an `int` parameter `x` and returns an array of `x` elements of type `A`.

Consider the reduction operator:

```
Optional<T> max(Comparator<? super T> c)
```

This operator requires a comparator for elements of type `T` as parameter and returns a value of type `Optional<T>`. This type represents a value of type `T` that may not exist. The maximum of all elements in a stream does not exist if the stream has *no* elements. The class `Optional<T>` has a method

```
boolean isPresent()
```

that returns `true` if the optional value is actually present. The class also has a method

```
T get()
```

that returns the value if it is indeed present. If `get()` is called when the value is not present, it will throw `NoSuchElementException`.

Now, suppose we want to find and print the string in the `stooges` stream that comes first in alphabetic order. We can use the statements:

```
stooges = Stream.of("Larry", "Curly", "Moe");
Optional<String> min = stooges.min( (x, y) -> x.compareTo(y) );
System.out.printf("First stooge by alphabetic order: %s\n", min.get());
```

Here we have used the lambda expression

```
(String x, String y) -> s.compareTo(y)
```

to create an object that implements the `Comparator` interface for string types. It turns out that Java 8 allows a shorter way to write such a lambda expression: you can just write `String::compareTo`. The call to the `min` reduction can therefore be written like this:

```
Optional<String> min1 = stooges.min(String::compareTo);
```

In the code above, we do not bother to call `min.isPresent()` before we call `min.get()` because we know our `stooges` stream is not empty. You should call `isPresent()` before calling `get()` whenever there is a possibility that the stream is empty.

The reduction operator

```
<A> A[] toArray(IntFunction<A[]> generator)
```

collects stream elements into an array that is created by an array *generator*. This is a function that takes an integer `x` and returns an array of length `x`, whose elements have type `A`. In our case, the type `A` is `String`, so we want the lambda expression:

```
x -> new String [x]
```

Again, Java 8 provides a convenient notation for such a function: simply write `String[]::new`. In general, for any type `A`, the expression `A[]::new` represents a function that takes an integer parameter `x` and returns an array of type `A` that has `x` elements.

[Code Listing 18-18](#) illustrates how to use some of the reduction operators in [Table 18-15](#).

Code Listing 18-18

(StreamReductionsDemo.java)

```
1 import java.util.Arrays;
2 import java.util.Optional;
3 import java.util.stream.Stream;
4 /**
5  This program illustrates the use of a
6  number of stream reduction operators.
7  */
8
9 public class StreamReductionsDemo
10 {
11     public static void main(String[] args)
12     {
13         // Create the source.
14         String[] stoogesArray0 =
15         {
```

```

16         "Larry", "Curly", "Moe"
17     };
18     System.out.printf("The Stream source is the array: %s \n",
19         Arrays.toString(stoogesArray0));
20     // Variable for the stream.
21     Stream<String> stooges;
22
23     // Get a stream and reduce with count.
24     stooges = Stream.of(stoogesArray0);
25     System.out.printf("The count is %d\n", stooges.count());
26
27     // Get a stream and reduce with forEach.
28     stooges = Stream.of(stoogesArray0);
29     stooges.forEach(x ->
30     {
31         System.out.printf("%s has length %d\n", x, x.length());
32     });
33
34     System.out.println("The stream reduced to an array (2 ways):");
35
36     // Get a stream and reduce with toArray().
37     stooges = Stream.of(stoogesArray0);
38     String[] stoogesArray1 = stooges.toArray(x -> new String[x]);
39     System.out.println(Arrays.toString(stoogesArray1));
40
41     // Get a stream and reduce with toArray().
42     stooges = Stream.of(stoogesArray0);
43     String[] stoogesArray2 = stooges.toArray(String[]::new);
44     System.out.println(Arrays.toString(stoogesArray2));
45
46     // Get a stream and reduce with min.
47     stooges = Stream.of(stoogesArray0);
48     Optional<String> min = stooges.min((x, y) -> x.compareTo(y));
49     System.out.printf("First stooge by alphabetic order: %s\n",
50         min.get());
51
52     // Get a stream and reduce with min.
53     stooges = Stream.of(stoogesArray0);
54     Optional<String> min1 = stooges.min(String::compareTo);
55     System.out.printf("First stooge by alphabetic order: %s\n",
56         min1.get());
57
58     // Get a stream and reduce with max.
59     stooges = Stream.of(stoogesArray0);
60     Optional<String> max = stooges.max(String::compareTo);
61     System.out.printf("Last stooge by alphabetic order: %s\n",
62         max.get());
63 }

```

Program Output

```

The Stream source is the array: [Larry, Curly, Moe]
The count is 3
Larry has length 5
Curly has length 5
Moe has length 3

```

```
The stream reduced to an array (2 ways):  
[Larry, Curly, Moe]  
[Larry, Curly, Moe]  
First stooge by alphabetic order: Curly  
First stooge by alphabetic order: Curly  
Last stooge by alphabetic order: Moe
```

In [Code Listing 18-18](#), we have to recreate the stream from the source every time we execute a reduction operation on it:

```
stooges = Stream.of(stoogesArray0);
```

This is because a stream object can be reduced only one time. Once it is reduced, the stream is drained of all its elements and becomes unuseable.

The `IntStream`, `DoubleStream`, and `LongStream` have additional reduction operators. Because of the close similarity, we will list only the operators for `DoubleStream`. Analogous operators exist for the other two streams.

```
OptionalDouble average()  
OptionalDouble max()  
OptionalDouble min()  
Double sum()  
double[] toArray()
```

These operators respectively reduce a stream of numbers of type `double` to their average, maximum, minimum, and sum. The `max()` and `min()` are similar to the corresponding operators of the same name in `Stream<T>`, but do not require a comparator because the natural order of numbers is used. The last operator collects stream elements into an array of type `double`.

Intermediate Operations on Streams

An intermediate operation acts on an existing stream and produces a new stream. [Table 18-16](#) lists some of the intermediate operations defined in the `Stream<T>` interface.

Table 18-16 `Stream<T>` intermediate operators

<code>Stream<T></code> <code>distinct()</code>	Removes duplicates and returns the resulting stream.

```
Stream<T> filter(Predicate<? super T> predicate)
```

Returns a stream consisting of only those elements that satisfy the given predicate.

```
Stream<T> limit(long n)
```

Returns a stream consisting of at most the first *n* elements of this stream.

```
<R> Stream<R> map(Function<? super T,? extends R> mappingFunction)
```

Returns a stream whose elements are the result of calling the mapping function on elements of this stream.

```
Stream<T> peek(Consumer<? super T> action)
```

Returns a stream consisting of the same elements as this stream. The given action method will be called on each element of the resulting stream as the element is being drawn from the stream.

```
Stream<T> skip(long n)
```

Returns this stream after discarding the first *n* elements.

```
Stream<T> sorted()
```

Returns a stream with the same elements as this stream, but sorted according to their natural order.

```
Stream<T> sorted(Comparator <? super T> comparator)
```

Returns a stream with the same elements as this stream, but sorted according to the given comparator.

[Code Listing 18-19](#) illustrates some of these intermediate operators. The program creates a stream that is able to produce infinitely many random integers less than 10. A filter operation is then plugged

onto that stream to filter for even numbers, followed by a `limit(15)` operation that will ask for only 15 of those numbers. The last operation in the pipeline reduces the resulting stream to an array, which is then printed.

Code Listing 18-19

(StreamIntermediateOpsDemo.java)

```
1 import java.util.Arrays;
2 import java.util.Random;
3 import java.util.stream.IntStream;
4
5 /**
6  This program illustrates intermediate Stream operations.
7 */
8
9 public class StreamIntermediateOpsDemo
10 {
11     public static void main(String [] args)
12     {
13         // Random number generator.
14         Random randy = new Random();
15
16         // Potentially infinite stream of random
17         // numbers all of which are less than 10.
18         IntStream infRandom = randy.ints(0, 10);
19
20         // Filter out all the odd numbers, keep only even numbers.
21         IntStream intEvenRandom = infRandom.filter(x -> x % 2 == 0);
22
23         // Keep only the first 15 of the infinite
24         // stream of random even numbers.
25         IntStream first15 = intEvenRandom.limit(15);
26
27         // Reduce to an array and print.
28         int [] evenRandomArr = first15.toArray();
29         System.out.println("The result of stream operations is: ");
30         System.out.println(Arrays.toString(evenRandomArr));
31     }
32 }
```

Program Output

```
The result of stream operations is:
[0, 2, 6, 2, 2, 6, 8, 4, 8, 8, 6, 4, 6, 6, 4]
```

Intermediate Stream Operations are Lazy

You may well be wondering how the stream

```
randy.ints(0,10)
```

can represent an infinite stream of numbers as claimed. The reason is that all intermediate stream operations are *lazy*: the intermediate operations will produce an element only when the element is being demanded by a reduction operation. To see how this works, let us remove all the “temporary” variables that are keeping track of the streams at intermediate stages. We end up with the pipeline expression:

```
evenRandomArr = randy.ints(0, 10)
                  .filter(x -> x% 2 == 0)
                  .limit(15)
                  .toArray();
```

When this is executed, the `toArray()` reduction demands all elements in the aggregate stream.

The `limit(15)` operation will then demand 15 elements from the filter operation, which will keep demanding elements from `ints(0,10)` until it has passed 15 even numbers to `limits(15)`. At that point, `limits(15)` is satisfied and stops asking for elements from filter, which stops asking for elements from `ints(0,10)`.

You can think of a stream as a conduit through which objects flow, beginning at a source and ending at some terminal operation. The `peek()` operation executes an action on each element as it passes by. Consider [Code Listing 18-20](#), and the output that follows the code listing. We have modified the stream pipeline shown above by placing a `peek()` on elements as they leave the `ints(0,10)` stream, and another `peek()` as the elements leave the `filter()`. To cut down on the volume of the output, we limit the size of the final stream to 5 elements.

Code Listing 18-20 (PeekDemo.java)

```
1 import java.util.Arrays;
2 import java.util.Random;
3 import java.util.stream.IntStream;
4
5 /**
6  This program illustrates the working of the
7  peek intermediate operation.
8 */
9
10 public class PeekOpDemo
11 {
12     public static void main(String [] args)
13     {
14         // Random number generator.
15         Random randy = new Random();
16
17         // Potentially infinite stream of random
18         // numbers all of which are less than 10.
```



```

19      IntStream infRandom = randy.ints(0, 10);
20
21      int [] evenRandomArr =
22          randy.ints(0, 10)
23              .peek(x -> {System.out.printf("ints: %d\n", x);})
24              .filter(x -> x% 2 == 0)
25              .peek(x-> {System.out.printf("filter: %d\n", x);})
26              .limit(5)
27              .toArray();
28      System.out.println("The result of stream operations is: ");
29      System.out.println(Arrays.toString(evenRandomArr));
30  }
31  }

```

Program Output

```

ints: 9
ints: 0
filter: 0
ints: 9
ints: 1
ints: 5
ints: 6
filter: 6
ints: 3
ints: 9
ints: 4
filter: 4
ints: 4
filter: 4
ints: 6
filter: 6
The result of stream operations is:
[0, 6, 4, 4, 6]

```

Notice that `ints()` produces 11 elements altogether, 6 odd and 5 even. All odd elements make it out of the `int()` stream, but get filtered out by `filter()`.

The Stream map Operation

The `map` operation

```
<R> Stream<R> map(Function<? super T,? extends R> mappingFunction)
```

applies a function of type $T \rightarrow R$ to transform a stream of elements of type T into a stream of type R . Actually the input type of the function needs only be a superclass of T , and the return type only needs to be a subclass of R .

A simple example of this is the following. If you have a stream of consecutive integers

```
[1, 2, 3, 4, 5, 6]
```

you can apply the mapping function $x \rightarrow 2x-1$ to every element of the stream and get the stream of consecutive odd numbers

```
[1, 3, 5, 7, 9, 11]
```

The code for this is:

```
int n = sc.nextInt();
IntStream odds = IntStream.rangeClosed(1, n).map(x -> 2*x -1);
System.out.println(Arrays.toString(odds.toArray()));
```

Now, suppose you want to turn a stream of strings into a stream of integers by mapping each string to its length. The map operation $T \rightarrow R$ requires that R be a reference type. If you use `map`, the compiler will automatically box the integer length of each string, creating a stream of type `Stream<Integer>` instead of one of type `IntStream`.

```
Stream<String> names = Stream.of("Abby", "Bob", "Christopher",
                                "David", "Elizabeth");
Stream<Integer> lengths = names.map(x -> x.length());
```

This means that executing the lines of code above will result in `lengths` being a stream of `Integer` rather than `int` elements. If you want to map a `Stream<T>` object to a stream whose elements are primitive types, you need to use the following `mapToXXX()` methods of the `Stream<T>` interface.

```
DoubleStream mapToDouble(ToDoubleFunction<? super T> mapper)
IntStream mapToInt(ToIntFunction<? super T> mapper)
LongStream mapToLong(ToLongFunction<? super T> mapper)
```

These methods take as input a function that takes a parameter of type T and returns a double, an `int`, or `long` value. The proper way to map a stream of strings to an `IntStream` is like this:

```
Stream<String> names = Stream.of("Abby", "Bob", "Christopher",
                                "David", "Elizabeth");
IntStream lengths = names.mapToInt(x -> x.length());
System.out.println(Arrays.toString(lengths.toArray()));
```

The Stream reduce Methods

Many reduction operations can be viewed as a process of “accumulating” stream elements into a single value of the same type. The simplest example of this is when the elements of a stream of integers are processed to yield another integer that is the sum of the stream elements. Another example is the processing of a stream of integers to yield a product.

To perform an accumulation reduction on a stream of type T , you need an *accumulation operator* p of type $(T, T) \rightarrow T$ that will be used to combine two values of type T into a single value. You also need a value i of type T to represent the result of performing the accumulation reduction on an empty stream.

This value is called the *identity* value for p and must satisfy the relation $p(i, x)=x$ for all values x of type T . The accumulation process uses an internal variable a to represent the value of stream elements already accumulated. Initially, no elements have been accumulated and a is set to i . The process proceeds by setting $a=p(a, x)$ for each stream element x . The final value of a is returned at the end of the process.

As an example, suppose we want to reduce a stream of integers to the sum of its elements. The accumulation operator p is a function that returns the sum of two integers. For example, p might be the lambda expression $(x, y) \rightarrow x+y$. The identity value for addition is the integer zero. Initially, the internal variable a would be 0. At each stage of the reduction, the current stream element would be added to a by setting $a=a+x$.

All three interfaces for streams of primitive types define `reduce()` methods that implement accumulation reductions of this nature:

```
int reduce(int identity, IntBinaryOperator op)
long reduce(long identity, LongBinaryOperator op)
double reduce(double identity, DoubleBinaryOperator op)
```

When you are accumulating an integer stream to find the sum, you use $(x, y) \rightarrow x+y$ for the binary operator and take 0 for the identity. However, if you are accumulating by multiplying, the appropriate binary operator is $(x, y) \rightarrow x*y$ and the identity value is 1. The following code uses reductions to find the sum of the first 4 positive odd integers, and also the product of the first 4 positive integers.

```
int n = 4;
// Create streams of integers.
IntStream odds = IntStream.rangeClosed(1, n).map(x-> 2*x -1);
IntStream ints = IntStream.rangeClosed(1, n);

// Reduce the streams.
int oddsSum = odds.reduce(0, (x,y)->x+y );
int intsProduct = ints.reduce(1, (x, y)-> x*y);

// Print the results.
System.out.printf("The sum is %d and the product is %d\n",
    oddsSum, intsProduct);
```

When executed, this code will print 14 for the sum and 24 for the product.

The `Stream<T>` interface has a similar `reduce()` method:

```
T reduce(T identity, BinaryOperator<T> accumulator)
```

This method works similarly to the ones we have already discussed.

The `reduce()` methods require the binary accumulator operator p to be associative. A binary operator p is *associative* if for any sequence of values a, b, c , the two values $p(a, p(b, c))$ and $p(p(a, b), c)$ are equal. The requirement allows for the parallelization of the reduction operation. The implementation may split the stream into segments, reduce these segments in parallel, and later combine the result. The associativity of the accumulator operator ensures that final result will not be

affected by the order in which the results are combined.

There are times when you want to reduce the elements of a stream of type T to a single value of a different type U . To meet this need, the `Stream<T>` interface defines a method:

```
<U> U reduce(U identity, BiFunction<U,? super T,U> accumulator,  
            BinaryOperator<U> combiner)
```

This method requires an *accumulator* bi-function f of type $(U, T) \rightarrow U$ to “add” a current value of type T to the accumulated value of type U . It also requires a value i of type U to act as the initial accumulation value to get the process started. At each stream element x , the reduction replaces the current accumulation value a with $f(a, x)$ by setting $a = f(a, x)$. The reduction may split the stream into segments and reduce the segments in parallel, with each segment producing an accumulated value of type U . These values are combined two at a time by using a binary *combiner* operator c , which must be associative. Moreover, the accumulation value i must be an *identity* value for the operator c .

As an example, suppose that we want to reduce a stream of `String` elements by collecting them into a set. We will see shortly that there is a `streamCollectors` method that will do that, but let us push ahead for the sake of illustration. Here the type T is `String`, while the type U is `Set`. At each stage, the accumulated value will be a set, and the accumulator bi-function will be a function that takes a set and a string, adds the string to the set, and returns the resulting set:

```
(Set s, String x) -> { s.add(x); return s; }
```

The combiner operator will take two sets, add all the elements of the first set to the second set, and return the second set:

```
(Set s, Set t) -> { t.addAll(s); return t; }
```

Finally, the identity value that gets the set-building process started is an empty set:

```
new HashSet<String>();
```

[Code Listing 18-21](#) shows the complete program.

Code Listing 18-21 (ReduceToSetDemo.java)

```
1 import java.util.HashSet;  
2 import java.util.Set;  
3 import java.util.function.BiFunction;  
4 import java.util.function.BinaryOperator;  
5 import java.util.stream.Stream;  
6  
7 /**  
8    This program demonstrates how to use the Stream  
9    reduce method.  
10 */  
11
```

```

12 public class ReducetoSetDemo
13 {
14     public static void main(String [] args)
15     {
16         Stream<String> stooges = Stream.of("Larry", "Curly", "Moe");
17
18         // The accumulator for the reduction.
19         BiFunction<Set<String>, String, Set<String>> accumulator =
20             (s, x) -> { s.add(x); return s;};
21
22         // The combiner for the reduction
23         BinaryOperator<Set<String>> combiner =
24             (s, t) -> { t.addAll(s); return t;};
25
26         // Reduce the stream to a set.
27         Set<String> stoogeSet = stooges.reduce(new HashSet<String>(),
28                                             accumulator, combiner);
29         // Print the set
30         System.out.printf("The set of stooges is %s\n", stoogeSet);
31     }
32 }

```

Program Output

The set of stooges is [Moe, Larry, Curly]

Notice, we can easily convert a stream into a set using the `forEach()` reduction method that we learned about earlier. The advantage of the `reduce()` methods lies in the potential for parallelization. When working with large quantities of data, parallelization may result in significant savings in time.

Stream Collectors

Collectors are a slightly different approach to reductions. *Collectors* are typically used to accumulate stream elements into a collection, although they can perform other types of reductions.

The `Stream<T>` interface has a method

```
<R,A> R collect(Collector<? super T,A,R> collector)
```

that uses a collector object to reduce a stream of elements of type T into a single value of type R . The collector accumulates stream elements into a value of type A , and then transforms the type A value into a final result of type R when all stream elements have been processed.

The `Stream` package has a `Collectors` class that provides static methods that return different types of collectors. Among them are the two methods:

```

static <T> Collector<T,?,List<T>>  toList()
static <T> Collector<T,?,Set<T>>   toSet()

```

The `toList()` method returns a collector that will reduce a stream of elements of type T by collecting them into a list of type `List<T>`. Here the final type R is `List<T>`, and the intermediate type A is an unspecified implementation detail. This is because the user of the collector does not need to know the internal type used for the intermediate accumulated values. The `toSet()` method is similar.

Let us illustrate the use of the `toSet()` collector. The following code collects the elements of a stream of strings into a set and prints the set:

```
Stream<String> stooges = Stream.of("Larry", "Curly", "Moe");
Set<String> stoogesSet = stooges.collect(Collectors.toSet());
System.out.printf("The set of stooges is %s\n", stoogesSet);
```

As you can see, the `collect` method takes a lot less code than the program of [Code Listing 18-21](#).

Another interesting method is `toMap()`:

```
static <T,K,U> Collector<T,?,Map<K,U>>
    toMap(
        Function<? super T,? extends K> keyMapper,
        Function<? super T,? extends U> valueMapper
    )
```

This method returns a collector that can reduce a stream to a map. The method takes as parameters two functions: a key mapper that determines the key for a stream element, and a value mapper that determines the value associated with the key. It is these key-value pairs that are stored in the map.

To illustrate, suppose we have a stream of strings, and we want to create a map that stores the strings as values, with each string being keyed by its length. For example, a stream of the strings

```
"Larry", "Moe", "Al"
```

would result in a map that stores the following key-value pairs:

```
(5, Larry)
(3, Moe)
(2, Al)
```

In the syntax of the definition of `toMap()`, the stream element type T is `String`, the key type K is `Integer`, and the value type U is `String`. Our key-mapping function is the lambda expression,

```
(String x)-> x.length()
```

while the value-mapping function is the lambda expression

```
(String x) -> x
```

A program that reduces a stream of strings to such a map is shown in [Code Listing 18-22](#).

Code Listing 18-22 (ReduceToMapDemo1.java)

```

1 import java.util.Map;
2 import java.util.function.Function;
3 import java.util.stream.Collectors;
4 import java.util.stream.Stream;
5
6 /**
7     This program demonstrates the Collectors.toMap method.
8 */
9
10 public class ReduceToMapDemo1
11 {
12     public static void main(String [] args)
13     {
14         // This is the stream to reduce.
15         Stream<String> names = Stream.of("Larry", "Moe", "Al");
16
17         // key of a string is its length.
18         Function<String, Integer> keyMapper = x -> x.length();
19
20         // The value of a string is the string itself.
21         Function<String, String> valueMapper = x -> x;
22
23         // Reduce the stream to a map
24         Map<Integer, String> namesMap =
25             names.collect(Collectors.toMap(keyMapper, valueMapper));
26
27         // Print the map.
28         System.out.printf("The map of stooges is %s\n", namesMap);
29     }
30 }

```

Program Output

The map of stooges is {2=Al, 3=Moe, 5=Larry}

In the input to the program of [Code Listing 18-22](#), there is at most one string with a given length. This is because the lengths of the strings are keys of a map, and maps require keys to be unique. A more realistic reduction would allow any number of strings of a given length and groups the strings according to their length. Such a reduction would reduce a stream of strings to a map whose values are lists of strings, and the key of a value is the common length of all strings in the list. For example, the stream of strings

"Larry", "Curly", "Moe", "Al"

would be reduced to the map

```

(5, [Larry, Curly])
(3, [Moe])
(2, [Al])

```

The `Streams<T>` interface has such a method:

```

static <T,K,U> Collector<T,?,Map<K,U>>
    toMap(
        Function<? super T,? extends K> keyMapper,
        Function<? super T,? extends U> valueMapper,
        BinaryOperator<U> mergeFunction
    )

```

This method is just like the last one, except it has a *merge function* parameter that is a binary operator on the type *U*. The merge function *m* has type $(U, U) \rightarrow U$, meaning that *m* takes two parameters of type *U* and returns a single result of the same type. If two values *u1* and *u2* have the same key *k*, the `toMap()` method will call the merge function on the two values and use the result `m(u1, u2)` as the “merged” value for *k*.

Let’s look at how this plays out if we want to reduce a stream of strings into a map that groups strings by their lengths. First, note that the type *U* of the final result is no longer `String`, but `List<String>`. The key mapper is the same as before, but the value mapper now assigns to each string *x* a list whose only member is *x*:

```

Function<String, List<String>> valueMapper = x->
{
    List<String> listOfOne = new ArrayList<>();
    listOfOne.add(x);
    return listOfOne;
};

```

The merge function takes two lists that would have the same key and returns a single list that is the concatenation of the two lists:

```

BinaryOperator<List<String>> mergeFunction = (list1, list2) ->
{
    list1.addAll(list2);
    return list1;
};

```

[Code Listing 18-23](#) shows a complete program that implements the reduction.

Code Listing 18-23 (ReduceToDemo2.java)

```

1 import java.util.ArrayList;
2 import java.util.List;
3 import java.util.Map;
4 import java.util.function.BinaryOperator;
5 import java.util.function.Function;
6 import java.util.stream.Collectors;
7 import java.util.stream.Stream;
8
9 /**
10
11     This program demonstrates the Collectors.toMap method.
12 */
13

```



```

14 public class ReduceToMapDemo2
15 {
16     public static void main(String [] args)
17     {
18         // This is the stream to reduce to a map.
19         Stream<String> names = Stream.of("Larry", "Curly", "Moe", "Al");
20
21         // The key for a string is the length of the string.
22         Function<String, Integer> keyMapper = x -> x.length();
23
24         // The value for a string is a list of all strings in
25         // the stream with the same length. Every string is initially
26         // mapped to list containing just that string.
27         Function<String, List<String>> valueMapper = x->
28         {
29             List<String> listOfOne = new ArrayList<>();
30             listOfOne.add(x);
31             return listOfOne;
32         };
33
34         // When two values have the same key, the merge function
35         // is used to merge the two values and the merged result
36         // becomes the value of the key.
37         BinaryOperator<List<String>> mergeFunction = (list1, list2) ->
38         {
39             list1.addAll(list2);
40             return list1;
41         };
42
43         //Reduce the stream and get a map.
44         Map<Integer, List<String>> namesMap =
45         names.collect(Collectors.toMap(
46                                     keyMapper, valueMapper,
47                                     mergeFunction));
48
49         // Print the map.
50         System.out.printf("The map of names is %s\n", namesMap);
51     }
52 }

```

Program Output

The map of names is {2=[Al], 3=[Moe], 5=[Larry, Curly]}

The `Stream<T>` interface has a second `collect` method that takes as parameters the ingredients needed to build a collector. The method internally builds the collector and uses it to reduce the stream. You can use this method in the event that none of the collectors in the `Collectors` class suit your needs.

```

<R> R collect(
    Supplier<R> supplier,
    BiConsumer<R,? super T> accumulator,
    BiConsumer<R,R> combiner
)

```

Briefly, the supplier is a method that is called to provide an initial value of the final type R ; the role of the supplier is similar to the role of the identity in the reduce methods. The accumulator and combiners play roles similar to those of the accumulator and combiner in the reduce methods. You can learn more about this and other methods by studying the streams API documentation.



Checkpoint

1. 18.33 How do you define a stream of elements?
2. 18.34 How does a stream intermediate operation differ from a terminal operation?
3. 18.35 Name at least two different classes or interfaces that have methods for creating streams.
4. 18.36 What is a stream pipeline?
5. 18.37 What is the difference between a terminal operation and a reduction?
6. 18.38 What is a collector?
7. 18.39 How many reduction operations can be in a stream pipeline?

18.7 Common Errors to Avoid

- Not calling an iterator's `hasNext` method before calling `next`. To prevent an exception from being thrown, you should always call an iterator's `hasNext` method before calling its `next` method. This ensures that the `next` method will actually return a reference.
- Not calling an iterator's `next` method before calling `remove`. An iterator's `remove` method removes the last element that was returned by the `next` method. The `next` method must be called before the `remove` method can be called.
- Not calling a `ListIterator`'s `hasPrevious` method before calling `previous`. To prevent an exception from being thrown, you should always call a `ListIterator`'s `hasPrevious` method before calling its `previous` method.
- Not calling a `ListIterator`'s `hasNext` or `hasPrevious` method immediately before calling `set`. To prevent an exception from being thrown, you should always call a `ListIterator`'s `hasNext` or `hasPrevious` method before calling its `set` method.
- Modifying a collection via one of its methods after an iterator has been created for the collection. After an iterator has been created for a collection, you cannot modify the collection using one of the collection's methods and then continue to use the iterator. Doing so will cause an exception to be thrown. If you wish to modify a collection while it is being iterated over, you should use the iterator to perform the modifications. For example, the `Iterator` interface specifies a `remove` method, and the `ListIterator` interface specifies `add` and `set` methods.
- Designing a class so objects of the class can be stored in a `HashSet`, but not overriding both the `equals` and `hashCode` methods. When an object is stored in a `HashSet`, both the `hashCode` and the `equals` methods are used to find a location for the object in the `HashSet`. If you are designing a class so objects of the class can be stored in a `HashSet`, you should override both the `hashCode` and the `equals` methods.
- Storing objects that cannot be compared in a `TreeSet`. In a `TreeSet`, objects are compared using the `compareTo` method. An error will occur if you attempt to store in a `TreeSet` an object of a class that does not implement the `compareTo` method.
- Attempting to store duplicate keys in a map. The keys in a map must be unique.
- Trying to reuse a stream after it has been reduced. A stream becomes drained of its elements when it is reduced. Trying to use it after that will result in an error.

Review Questions and Exercises

Multiple Choice and True/False

1. This type of collection stores elements in an ordered sequence, and allows duplicates.
 1. list
 2. set
 3. map
 4. hash set
2. Each element in one of these is comprised of two parts: a key and a value.
 1. list
 2. set
 3. map
 4. hash set
3. This type of collection is optimized for searching, stores elements in an unordered sequence, and does not allow duplicates.
 1. list
 2. set
 3. map
 4. `ArrayList`
4. This interface specifies the basic operations of a list.
 1. `ListCollection`
 2. `List`
 3. `Set`
 4. `ArrayList`

5. A terminal operation in a stream pipeline is also called
 1. the stream end
 2. a reduction
 3. the sink
 4. an iterator
6. This is an object that can be used to retrieve the elements in a collection, one after the other.
 1. iterator
 2. manipulator
 3. comparator
 4. concrete collection class
7. `ListIterator` is a subinterface of
 1. `Iterator`
 2. `Collection`
 3. `Comparator`
 4. `List`
8. This `ListIterator` method replaces an existing element with a new element.
 1. `replace`
 2. `update`
 3. `set`
 4. `add`
9. A stream collector is an object that
 1. implements the `Collection` interface
 2. implements the `Map` interface
 3. is used to reduce a stream

4. uses lambda expressions
10. This list class stores its elements by connecting each element to its successor.
 1. ArrayList
 2. LinkedList
 3. HashList
 4. ConnectedList
11. This is an object that can compare two other objects and determine whether one is greater than, less than, or equal to the other.
 1. Comparator
 2. Comparer
 3. Sorter
 4. Sequencer
12. This class provides numerous static methods that perform useful operations on collections.
 1. Collection
 2. CollectionUtilities
 3. Collections
 4. Util
13. This JCF interface is not a subinterface of `Collection`.
 1. List
 2. Set
 3. Map
 4. SortedSet
14. This class is a good choice to use when your application is storing data in a list and is doing lots of insertions and deletions in the middle of the list.
 1. ArrayList

2. `Vector`
 3. `LargeList`
 4. `LinkedList`
15. **True or False:** Duplicates can be stored in a list.
 16. **True or False:** Duplicates can be stored in a set.
 17. **True or False:** `Collection` is a subinterface of `List`.
 18. **True or False:** A `ListIterator` can move backward through a list.
 19. **True or False:** A `ListIterator` can be used to modify the contents of a list.
 20. **True or False:** When the Java compiler processes an enhanced `for` loop, it actually converts it to a traditional loop that uses an iterator.
 21. **True or False:** You can randomly access the elements in a linked list.
 22. **True or False:** When you iterate through a `HashSet`, the objects will be retrieved in the same order that they were added to the `HashSet`.

Find the Error

1.

```
List<String> nameList = new LinkedHashSet<>();
```
2.

```
List<String> nameList = new ArrayList<>();

ListIterator it = new ListIterator(nameList);
```
3.

```
Set<String> nameSet = new HashSet<>();

nameSet.add("Chris");

nameSet.add("Kenny");

ListIterator it = nameSet.listIterator();
```
4.

```
List<String> nameList = new ArrayList<>();

ListIterator it = new ListIterator(nameList);

it.set("Herbert");
```

Algorithm Workbench

1. Write a statement that declares a `List` reference variable, instantiates an `ArrayList` object that can hold `String` objects, and assigns the address of the `ArrayList` object to the reference variable.
2. Assume that `myList` references a `List` object. Write a statement that gets a list iterator from the `myList` object.
3. Assume that `it` references a newly created iterator for a list of `String` objects. Write code that uses the iterator to display each of the `String` objects in the list.
4. In pseudocode, outline the procedure followed by a `HashSet` object to add a new object.

Short Answer

1. What do you call an object that provides the elements that make up a stream?
2. What is the essential difference between a stream intermediate operation and a stream terminal operation?
3. What is the benefit of using an interface variable to reference a collection object?
4. What is an iterator?
5. How do you get an iterator for a list?
6. Which element does the `ListIterator` method `set` replace in a list?
7. How does the Java compiler process an enhanced `for` loop?
8. Which is more efficient—inserting an element in the middle of an `ArrayList`, or inserting a node into the middle of a `LinkedList`?
9. Why do we need `IntStream` when we have `Stream<Integer>`?
10. What is the purpose of the `map()` method in the `Stream<T>` interface?
11. Look at the following code:

```
String str1 = new String("Barney");  
  
String str2 = new String("Barney");
```

Will the methods `str1.hashCode()` and `str2.hashCode()` return the same value?

12. In hashing, what is a collision?
13. How does a map work?
14. What are the three guidelines that you should follow when overriding the `hashCode` and `equals` methods in a class?

Programming Challenges

1. Word Set

Write an application that reads a line of input from the keyboard, and then displays each unique word that was entered, sorted in ascending order. You can do this by tokenizing the line of input and adding each token to an appropriate `Set` object.

2. Instructor Set

[Chapter 8](#) presented an `Instructor` class that holds information about an instructor. (See Section 8.7—Aggregation) Modify the `Instructor` class by overriding the `hashCode` and `equals` methods. Then write a class that stores several `Instructor` objects in a `HashSet`. The class should be able to display all the instructors in the set, and allow the user to search for an instructor. Demonstrate the class in an application.

3. EmployeeMap Class

Create an `Employee` class that stores an employee's ID number and name. Then create an `EmployeeMap` class that allows you to add `Employee` objects and look them up by their ID numbers. The `EmployeeMap` class should use a `Map` object to map ID numbers to `Employee` objects. Create an application to demonstrate the classes.

4. Stock Map

[Chapter 8](#) presented a `Stock` class that holds information about a stock. (See Section 8.4—The `toString` Method) Write a class that keeps `Stock` objects in a `Map`. The class should be able to retrieve a particular `Stock` object from the `Map` by searching on its stock symbol. Demonstrate the class in an application.

5. Prime Number List

A prime number is an integer (other than 1) that is only evenly divisible by itself and 1. For example, the number 5 is prime because it can only be evenly divided by 1 and 5. The number 6, however, is not prime because it can be divided evenly by 1, 2, 3, and 6. Write an application that calculates and stores the first 100 prime numbers (starting at 2) in a list. Use an iterator to display the numbers.

6. Dealing Cards

Write a class named `Card`, which can be used to represent a card from a deck of cards. The class should be able to store a card's suit and face value. A card's suit can be one of the following: Hearts, Diamonds, Spades, or Clubs. A card's face value can be Ace, Jack, Queen, King, or a value in the range of two through ten.

Next write a `Deck` class. This class constructor should create a list of 52 `Card` objects, each representing a valid card in a deck of cards. The class should have a `shuffle` method that randomly shuffles the `Card` objects in the list. It should also have a `deal` method that “deals” a card from the deck. It does this by removing the `Card` object at the beginning of the list and returning a reference to that object.

Next, write `CardPlayer` class. This class should keep a list of `Card` objects that have been dealt to it. This represents a hand of cards. A method named `getCard` should accept a reference to a `Card` object, which is added to the list. A method named `showCards` displays the `Card` objects in the list.

Demonstrate these classes in an application that creates a `Deck` object, shuffles the cards it contains, and deals five cards from the `Deck` to a `CardPlayer` object. The `CardPlayer` should then display its cards.

7. Word Frequency Count

Write a program that allows the user to specify a text file, opens the file, and prints a two-column table consisting of all the words in the file together with the number of times that each word appears. Words are space-delimited and case-sensitive. The table should list the words in alphabetical order.

8. groupingBy Collector

Consult the online Stream API documentation and read about the `groupingBy()` method of the `Collectors` class, and use that method to get a simpler and shorter solution to the problem of classifying strings that was solved in the program of [Code Listing 18-23](#).

Chapter 19 Linked Lists

TOPICS

1. [19.1 Introduction to Linked Lists](#)
2. [19.2 Operations on Linked Lists](#)
3. [19.3 Doubly Linked and Circularly Linked Lists](#)
4. [19.4 Recursion on Linked Lists](#)
5. [19.5 Common Errors to Avoid](#)

19.1 Introduction to Linked Lists

CONCEPT: A linked list is an implementation of a list in which elements are stored in nodes, and each node contains a link to its successor.

As you learned in [Chapter 18](#), a list is a collection that stores its elements in a position-based sequence. Each element in a list is associated with an index that corresponds to its position within the list. The first element has index 0, the second element has index 1, and so forth.

This position-based order gives rise to the notion of a predecessor and a successor for each element: if X is a list element with index k , then its *successor* is the element with index $k+1$ and its *predecessor* is the element with index $k-1$. The first element in the list has no predecessor and the last element in the list has no successor.

There are two ways to allocate storage for the elements of a list: contiguous allocation, and linked allocation. Array-based lists use *contiguous allocation*, which allocates storage for successive list elements in consecutive memory locations. A drawback of contiguous allocation is the necessity to move large numbers of elements when inserting elements into, or deleting elements from, the middle of the list.

Linked allocation does not need consecutive memory locations to store list elements. Instead, it stores with each element, a pointer to the memory location of its successor. A *linked list* is a list that uses linked allocation. In Java, a linked list is implemented as a sequence of *nodes*, where each node stores a list element and a reference to the next node in the list. For example, to implement a linked list of `String` elements, we can use a class such as the following:

```
class Node
{
    String value;           // Value of a list element
    Node next;             // Link to next node in the list

    Node(String val, Node n)
    {
        value = val;
        next = n;
    }
    Node(String val)
    {
        value = val;
        next = null;
    }
}
```

A linked list is represented by a reference to its first node. The first element of the list is called the *head* of the list, so this reference is often called the *head reference*. An empty linked list is indicated by setting its head reference to `null`.

Our examples in this chapter will involve linked lists of strings. Working with linked lists of other types of elements, and even with generic linked lists, is similar.

Creating Lists and Adding Nodes

Once we have a suitable node class, we can create linked lists by creating nodes and setting their successor links to reflect the order that the list imposes on its elements. A sequence of statements that builds a linked list is shown in [Figure 19-1](#). In the figure, a `Node` object is shown as a rectangular box with a compartment for the element and a second compartment for the successor reference `next`.




Figure 19-1 Adding nodes to a linked list

[Figure 19-1 Full Alternative Text](#)

The first statement in [Figure 19-1](#)

```
Node myList = new Node("Bob");
```

creates a linked list with the single element "Bob". At this point, `myList` refers to a `Node` object that contains a field `myList.next`, which can itself be made to point to another node:

```
myList.next = new Node("Carol");
```

This creates a linked list whose two elements form the sequence "Bob Carol". The newly added node contains its own `next` field, which can be accessed by using the expression `myList.next.next` and made to point to yet another node:

```
myList.next.next = new Node("Debby");
```

We are not limited to the end of a list when adding nodes. The statement


```
Node p = new Node("Allan", myList);
```

adds a new node "Allan" to the beginning of the list of elements already accumulated. This statement uses the two-parameter constructor, which allows the successor of a node to be specified. The result of executing this sequence of four statements is the list of four elements shown at the bottom of [Figure 19-1](#). The figure shows the `next` field of a `Node` object as an arrow pointing to a successor node. A backward diagonal drawn through the `next` field indicates a `null` value, and signifies that the node has no successor.

In addition to adding nodes at the beginning and end of a list, we can insert nodes into the middle. Given the final scenario shown in [Figure 19-1](#), we can add the string "Brad" immediately after "Bob" by writing

```
Node b = p.next;           // b points to "Bob"  
Node c = b.next;           // c points to "Carol"  
b.next = new Node("Brad", c);
```

Execution of these statements gives us the sequence of scenarios shown in [Figure 19-2](#).



```
b.next = new Node("Brad", c);
```




Figure 19-2 Inserting a node in the middle of a linked list

[Figure 19-2 Full Alternative Text](#)

In the example just shown, we used two auxiliary references, `b` and `c`, to insert a node in the middle of a list. In general, adding a new node at any place other than the beginning requires a reference to the node that will be the predecessor (in this case, `b`) and to a node that will be the successor of the new node (in this case, `c`). The auxiliary variables are not always needed, however. For example, instead of using the three statements shown in [Figure 19-2](#), we can add "Brad" to the list by writing

```
myList.next = new Node("Brad", myList.next);
```

Removing a Node

You remove a node from a linked list by “routing” the successor link around it. Basically, to remove a node *X* (say `Carol` in the first list shown in [Figure 19-3](#)) from the middle of a linked list, you must first obtain references to its predecessor (In this case, `Bob`) and successor (In this case, `Debby`).




Figure 19-3 Removing a node from the middle of a linked list

[Figure 19-3 Full Alternative Text](#)

```
Node b = p.next;          // Predecessor of Carol
Node d = b.next.next;     // Successor of Carol
```

Then you remove the targeted node *X* by manipulating references to bypass *X* and make the successor of *X* the successor of its predecessor:

```
b.next = d;
```

This sequence of events is illustrated in [Figure 19-3](#).

As you may already have guessed, the removal of a node does not require an explicit reference to the successor of the targeted node. The following code fragment also removes `Carol` from the list, and is shorter:

```
b = p.next;
b.next = b.next.next;
```

The procedure we have just outlined for removing a node from the middle of a list will not work when the node being targeted for removal is the first in the list. This is because the first node has no predecessor. You remove a node at the beginning of a list by moving the head reference one node forward. For instance, to remove "Allan" from the list at the beginning of [Figure 19-3](#), simply write:

```
p = p.next;
```

As we have seen, both addition and removal require a reference to the node just before the point of the addition or removal. You can obtain such a reference by setting a reference to the head of the list and then walking the reference forward, one node at a time, to get it to the desired place. The process of walking a reference along the nodes of a list is called traversing the list.

Traversing Linked Lists

A *traversal* of a collection is a systematic method of doing some processing at each element of the collection. To traverse a linked list, set a reference, say `ref`, to the first node in the list, process the value at that node (by accessing `ref.value`) and then move on to the next node by setting `ref` to the successor of the current node:

```
ref = ref.next;
```

In this way, the reference `ref` moves from each node to the next until it becomes `null` at the end of the list. The following code fragment traverses a linked list `myList` and prints all of its elements:

```
Node ref = myList;
while (ref != null)
{
    System.out.println(ref.value);
    ref = ref.next;
}
```

[Code Listing 19-1](#) builds a simple linked list and then traverses it to print its contents. The centerpiece of the program is the `LinkedList0` class. This class contains a private declaration of `Node`, which it uses to represent the linked list. `LinkedList0` uses a private field `first` to point to the first node of the linked list.

Code Listing 19-1 (`LinkedList0.java`)

```

1 public class LinkedList0
2 {
3     /**
4      * The Node class represents a list node.
5      */
6
7     private class Node
8     {
9         String value;
10        Node next;
11
12        /**
13         * Constructor.
14         * @param val The element to store in this node.
15         * @param n The reference to the next node.
16         */
17
18        Node(String val, Node n)
19        {
20            value = val;
21            next = n;
22        }
23
24        /**
25         * Constructor.
26         * @param val The element to store in this node.
27         */
28
29        Node(String val)
30        {
31            value = val;
32            next = null;
33        }
34    }
35
36    //Reference to the first node in the list
37    private Node first = null;
38
39    /**
40     * Constructor.
41     * Builds a linked list.
42     */
43
44    public LinkedList0()
45    {
46        // Incrementally build the list
47        // Chuck Debby Elaine Fred
48        first = new Node("Debby");
49        first.next = new Node("Elaine");
50        first.next.next = new Node ("Fred");
51        first = new Node ("Chuck", first);
52
53        String [ ] names = {"Bob", "Allan"};
54
55        // Use a loop to add all names in the array to the
56        // front of the linked list to build the list
57        // Allan Bob Chuck Debby Elaine Fred
58        for (String s : names)

```

```

59         first = new Node(s, first);
60     }
61
62     /**
63      * The print method traverses the list
64      * and prints all of its elements.
65      */
66
67     public void print()
68     {
69         Node ref = first;
70         while (ref != null)
71         {
72             System.out.print(ref.value + " ");
73             ref = ref.next;
74         }
75     }
76
77     /**
78      * The main method creates the linked list
79      * and invokes its print method.
80      */
81
82     public static void main(String [] args)
83     {
84         LinkedList0 ll = new LinkedList0();
85         String str = "The contents of the list are:";
86         System.out.println(str);
87         ll.print();
88     }
89 }

```

Program Output

The contents of the list are:
Allan Bob Chuck Debby Elaine Fred



Checkpoint

1. 19.1 Give a sequence of statements that creates a linked list of strings "red green blue" by creating a node with "green", then a node with "blue", then a node with "red", in that order.
2. 19.2 Consider the following statements:

```

Node colorList = new Node("Yellow", new Node("Purple"));
colorList = new Node("Brown", colorList);
System.out.print(colorList.value);
System.out.print(colorList.next.value);
System.out.print(colorList.next.next.value);

```

What is the output?

19.2 Operations on Linked Lists

CONCEPT: The most basic operations on linked lists are traversing the list, moving to a position determined by an index, adding an element, removing an element, and searching for an element.

Let us write a linked list class that implements a subset of the methods of the `java.util.List` interface. Our class will implement the following methods:

- `boolean isEmpty()`: check if the list is empty.
- `int size()`: return the number of elements in the list.
- `void add(String e)`: add an element `e` to the end of the list.
- `void add(int index, String e)`: add the specified element `e` to the list at the position given by `index`.
- `String remove(int index)`: remove and return the element at the position given by `index`.
- `boolean remove(String e)`: remove the specified element `e` and return true if the item was successfully removed.

The `LinkedList1` class uses the same `Node` class as `LinkedList0`. The class also uses two references

```
Node first;  
Node last;
```

to keep track of the first and last nodes in the list. The reference `first` serves as the head of the list, and `last` is used to enable elements to be added quickly to the end of the list. Both references will be set to `null` when the list is empty.

The `isEmpty` Method

We can check to see if the list is empty by evaluating the expression `first == null`. The `isEmpty` method simply returns the value of that expression:

```
boolean isEmpty()  
{  
    return first == null;  
}
```

The `size` Method

The number of elements in a linked list can be determined by initializing a counter to 0 and then incrementing the counter at each step of a traversal of the linked list.

The add Methods

The `add(String e)` method adds a new element `e` to the end of the linked list. If the list is empty, the new node containing `e` becomes both the first and last node:

```
first = new Node(e);
last = first;
```

If the list is nonempty, the new node is made the successor of the last node, and then the reference `last` is set to point to the new node:

```
last.next = new Node(e);
last = last.next;
```

The `add(int index, String e)` method is a little more involved. We consider several cases:

- The index is invalid. If the index is negative, or greater than the size of the list, the method throws an exception.
- The index is zero. The new node becomes the first node of the list. If the list was previously empty, then `last` must be set to `first`.
- The index is greater than zero. The method must find the right place to add the new node. It does this by skipping `index-1` nodes to get to the node that is to be the predecessor of the new node as follows:

```
Node pred = first;
for (int k = 1; k <= index - 1; k++)
{
    pred = pred.next;
}
```

Once the predecessor is found, the new node is created and spliced in as follows:

```
pred.next = new Node(e, pred.next);
```

At this point, the method checks to see if the new node is now the last node. If so, `last` must be updated to point to the new node as follows:

```
if (pred.next.next == null)
    last = pred.next;
```

The remove Methods

The remove methods follow logic similar to the corresponding add methods. Let us consider the `remove(int index)` method. Again, we consider a number of different cases:

- The index is invalid. If the index is negative, or greater or equal to the size of the list, the method throws an exception.
- The index is zero. Here the element to be removed is the first on the list. The method moves `first` down to remove the first node. If the list is now empty, then `last` must be set to `null`.
- The index is greater than zero. In this case the method sets a reference `pred` to the predecessor of the node being targeted for removal. The target node is then removed. If the removed node was at the end of the list, then `last` must be reset to point to its predecessor:

```
pred.next == pred.next.next;
if (pred.next == null)
    last = pred;
```

The toString Method

This method returns a string representation of the contents of the list, and is useful for displaying the list in a GUI component. The method appends the string value of each list element to a `StringBuilder` object and returns the string value representation of that object.

The `LinkedList1` class and all of its methods can be seen in [Code Listing 19-2](#).

Code Listing 19-2 (LinkedList1.java)

```
1  /**
2   * The LinkedList1 class implements a Linked list.
3   */
4
5  class LinkedList1
6  {
7      /**
8       * The Node class stores a list element
9       * and a reference to the next node.
10     */
11
12     private class Node
13     {
14         String value;
15         Node next;
16
17         /**
18          * Constructor.
19          * @param val The element to store in the node.
20          * @param n The reference to the successor node.
21         */
```

```

22
23     Node(String val, Node n)
24     {
25         value = val;
26         next = n;
27     }
28
29     /**
30     Constructor.
31     @param val The element to store in the node.
32     */
33
34     Node(String val)
35     {
36         // Call the other (sister) constructor.
37         this(val, null);
38     }
39 }
40
41 private Node first;    // list head
42 private Node last;    // last element in list
43
44 /**
45 Constructor.
46 */
47
48 public LinkedList1()
49 {
50     first = null;
51     last = null;
52 }
53
54 /**
55 The isEmpty method checks to see
56 if the list is empty.
57 @return true if list is empty,
58 false otherwise.
59 */
60
61 public boolean isEmpty()
62 {
63     return first == null;
64 }
65
66 /**
67 The size method returns the length of the list.
68 @return The number of elements in the list.
69 */
70
71 public int size()
72 {
73     int count = 0;
74     Node p = first;
75     while (p != null)
76     {
77         // There is an element at p
78         count ++;

```



```

79         p = p.next;
80     }
81     return count;
82 }
83
84 /**
85     The add method adds an element to
86     the end of the list.
87     @param e The value to add to the
88     end of the list.
89 */
90
91 public void add(String e)
92 {
93     if (isEmpty())
94     {
95         first = new Node(e);
96         last = first;
97     }
98     else
99     {
100         // Add to end of existing list
101         last.next = new Node(e);
102         last = last.next;
103     }
104 }
105
106 /**
107     The add method adds an element at a position.
108     @param e The element to add to the list.
109     @param index The position at which to add
110     the element.
111     @exception IndexOutOfBoundsException When
112     index is out of bounds.
113 */
114
115 public void add(int index, String e)
116 {
117     if (index < 0 || index > size())
118     {
119         String message = String.valueOf(index);
120         throw new IndexOutOfBoundsException(message);
121     }
122
123     // Index is at least 0
124     if (index == 0)
125     {
126         // New element goes at beginning
127         first = new Node(e, first);
128         if (last == null)
129             last = first;
130         return;
131     }
132
133     // Set a reference pred to point to the node that
134     // will be the predecessor of the new node
135     Node pred = first;

```

```

136     for (int k = 1; k <= index - 1; k++)
137     {
138         pred = pred.next;
139     }
140
141     // Splice in a node containing the new element
142     pred.next = new Node(e, pred.next);
143
144     // Is there a new last element ?
145     if (pred.next.next == null)
146         last = pred.next;
147 }
148
149 /**
150  * The toString method computes the string
151  * representation of the list.
152  * @return The string form of the list.
153  */
154
155 public String toString()
156 {
157     StringBuilder strBuilder = new StringBuilder();
158
159     // Use p to walk down the linked list
160     Node p = first;
161     while (p != null)
162     {
163         strBuilder.append(p.value + "\n");
164         p = p.next;
165     }
166     return strBuilder.toString();
167 }
168
169 /**
170  * The remove method removes the element at an index.
171  * @param index The index of the element to remove.
172  * @return The element removed.
173  * @exception IndexOutOfBoundsException When index is
174  *         out of bounds.
175  */
176
177 public String remove(int index)
178 {
179     if (index < 0 || index >= size())
180     {
181         String message = String.valueOf(index);
182         throw new IndexOutOfBoundsException(message);
183     }
184
185     String element; // The element to return
186     if (index == 0)
187     {
188         // Removal of first item in the list
189         element = first.value;
190         first = first.next;
191         if (first == null)
192             last = null;

```

```

193     }
194     else
195     {
196         // To remove an element other than the first,
197         // find the predecessor of the element to
198         // be removed.
199         Node pred = first;
200
201         // Move pred forward index - 1 times
202         for (int k = 1; k <= index - 1; k++)
203             pred = pred.next;
204
205         // Store the value to return
206         element = pred.next.value;
207
208         // Route link around the node to be removed
209         pred.next = pred.next.next;
210
211         // Check if pred is now last
212         if (pred.next == null)
213             last = pred;
214     }
215     return element;
216 }
217
218 /**
219  * The remove method removes an element.
220  * @param element The element to remove.
221  * @return true if the remove succeeded,
222  *         false otherwise.
223  */
224
225 public boolean remove(String element)
226 {
227     if (isEmpty())
228         return false;
229
230     if (element.equals(first.value))
231     {
232         // Removal of first item in the list
233         first = first.next;
234         if (first == null)
235             last = null;
236         return true;
237     }
238
239     // Find the predecessor of the element to remove
240     Node pred = first;
241     while (pred.next != null &&
242            !pred.next.value.equals(element))
243     {
244         pred = pred.next;
245     }
246
247     // pred.next == null OR pred.next.value is element
248     if (pred.next == null)
249         return false;

```

```

250
251         // pred.next.value is element
252         pred.next = pred.next.next;
253
254         // Check if pred is now last
255         if (pred.next == null)
256             last = pred;
257
258         return true;
259     }
260
261     public static void main(String [] args)
262     {
263         LinkedList1 ll = new LinkedList1();
264         ll.add("Amy");
265         ll.add("Bob");
266         ll.add(0, "Al");
267         ll.add(2, "Beth");
268         ll.add(4, "Carol");
269         System.out.println("The members of the list are:");
270         System.out.print(ll);
271     }
272 }

```

Program Output

```

The members of the list are:
Al
Amy
Beth
Bob
Caro

```

A Graphical Interface for the Linked List Program

Now we describe a graphical user interface that can be used to test the `LinkedList1` class and other classes like it. The user interface is shown in [Figure 19-4](#). The screenshot shows a scenario in which the user has given commands to add the strings `Amy`, `Bob`, `Carol`, and `David` to the list, followed by a command to remove the string at index 1 (`Bob`).




Figure 19-4 A graphical interface for testing the linked list class

The interface has a command entry text field at the bottom of the window, a text area in the center of the window, and an un-editable text field at the top. The text field at the top is used to display values returned by some of the list class methods.

The program that builds this user interface is shown in [Code Listing 19-3](#). It attaches an event handler to the command text field at the bottom of the window, causing the handler to be called when the user types a command into the text field and presses the `Enter` key. The event handler method retrieves the typed command, analyzes it to determine its meaning, and calls the appropriate `LinkedList1` class method.

The program recognizes simple commands that are closely patterned after the methods they are intended to test. The list of recognized commands:

- `isempty`: This command calls the `isEmpty()` method and displays the returned value in the result text field at the top of the frame.
- `size`: This command calls `size()` and displays the returned value in the result text field.
- `add s`: This command calls the `add(String e)` method and passes it `s` in place of `e`. For example, the user might type `add Amy`.
- `add i s`: This command calls the `add(int index, String e)` method and passes it the argument `i` for `index` and `s` for `e`. As an example, the user might type `add 2 Brad` to add "Brad" as the third element in the list.
- `remove s`: This command calls the `remove(String e)` method and passes `s` for `e`.
- `remove i`: This command calls the `remove(int index)` method and passes `i` for `index`.

The listener method works by separating the text retrieved from the command entry text field into a

command followed by its arguments. For example, the command text

```
add 2 Brad
```

is separated into three tokens: the command `add`, the integer index `2`, and the element `Brad`. The method creates a `Scanner` object that will read from the command string and then uses the usual `Scanner` methods to read strings and numbers from the command. This is done as follows:

```
String cmdText = cmdTextField.getText();
// Create a Scanner to read from the string
Scanner sc = new Scanner(cmdText);
```

The full details of the implementation can be found in [Code Listing 19-3](#). The program uses a separate, non-nested class for the event handler. The class uses constructor parameters to pass the linked list object, the text area to be used for the list view, and the text field used to display results. The text field used for command entry is accessed through the `getTarget()` method of the event parameter of the event handler method.

Code Listing 19-3 (LinkedListDemo.java)

```
1 import java.util.Scanner;
2 import javafx.application.Application;
3 import javafx.event.*;
4 import javafx.geometry.Insets;
5 import javafx.scene.Scene;
6 import javafx.scene.control.*;
7 import javafx.scene.layout.*;
8 import javafx.stage.Stage;
9
10 /**
11     This class is used to demonstrate
12     the operations in the LinkedList1 class.
13 */
14
15 public class LinkedListDemo extends Application
16 {
17     @Override
18     public void start(Stage stage) throws Exception
19     {
20         // Reference to linked list.
21         LinkedList1 ll = new LinkedList1();
22
23         // listView shows current list elements.
24         TextArea listView = new TextArea();
25         listView.setEditable(false);
26
27         // Input for user command.
28         TextField cmdTextField = new TextField();
29         cmdTextField.setPrefColumnCount(5);
30
31         // Displays result of list method that was
32         // invoked by the user command.
```

```

33     TextField resultTextField = new TextField();
34     resultTextField.setPrefColumnCount(5);
35     resultTextField.setEditable(false);
36
37     // Attach event handler to cmdTextField
38     EventHandler<ActionEvent> handler =
39         new CommandHandler(ll, listView, resultTextField);
40     cmdTextField.setOnAction(handler);
41
42     // HBox to contain command result label and text field.
43     HBox hbox1 = new HBox(10);
44     Label resultLabel = new Label("Command Result");
45     hbox1.getChildren().addAll(resultLabel, resultTextField);
46
47     // HBox to contain label and text field for command input.
48     HBox hbox2 = new HBox(10);
49     Label cmdLabel = new Label("Command: ");
50     hbox2.getChildren().addAll(cmdLabel, cmdTextField);
51
52     // VBox to contain the user interface components
53     VBox vbox = new VBox(10);
54     vbox.setPadding(new Insets(10));
55     vbox.getChildren().addAll(hbox1, listView, hbox2);
56
57     // Set up the scene and show the stage.
58     stage.setScene(new Scene(vbox));
59     stage.setTitle("Linked List Demo");
60     stage.show();
61 }
62
63 public static void main(String [] args)
64 {
65     launch();
66 }
67 }
68
69 // Event Handler class for the command text field.
70 class CommandHandler implements EventHandler<ActionEvent>
71 {
72     // Fields to hold information passed to the constructor.
73     private final LinkedList1 ll;
74     private final TextField resultTextField;
75     private final TextArea listView;
76
77     CommandHandler(LinkedList1 lList, TextArea lView,
78                   TextField rTfield)
79     {
80         ll = lList;
81         resultTextField = rTfield;
82         listView = lView;
83     }
84
85     @Override
86     public void handle(ActionEvent event)
87     {
88         // Get the command from the command textfield.
89         TextField cmdTextField = (TextField)event.getTarget();
90         String cmdText = cmdTextField.getText();

```

```

91
92 // Use a scanner to read the name of the linked list
93 // method and do a switch on it.
94 Scanner sc = new Scanner(cmdText);
95 String cmd = sc.next();
96 switch(cmd)
97 {
98     case "add":
99         if (sc.hasNextInt())
100         {
101             // add index element
102             int index = sc.nextInt();
103             String element = sc.next();
104             ll.add(index, element);
105         }
106         else
107         {
108             // add element
109             String element = sc.next();
110             ll.add(element);
111         }
112         listView.setText(ll.toString());
113         return;
114     case "remove":
115         if (sc.hasNextInt())
116         {
117             // remove index
118             int index = sc.nextInt();
119             String res = ll.remove(index);
120             resultTextField.setText(res);
121         }
122         else
123         {
124             // remove element
125             String element = sc.next();
126             boolean res = ll.remove(element);
127             String resText = String.valueOf(res);
128             resultTextField.setText(resText);
129         }
130         listView.setText(ll.toString());
131         return;
132     case "isempty":
133         String resText = String.valueOf(ll.isEmpty());
134         resultTextField.setText(resText);
135         return;
136     case "size":
137         String resText1 = String.valueOf(ll.size());
138         resultTextField.setText(resText1);
139         return;
140 }
141 }
142 }

```


19.3 Doubly Linked and Circularly Linked Lists

CONCEPT: A singly linked list makes it easy to move from any node to its successor. A doubly linked list makes it easy to move from any node to its successor or predecessor. A circularly linked list makes it easy to move from one end of the list to the other.

The singly linked lists we have been studying allow us to move from any node to its successor in one step. If we want to move from a node to its predecessor, though, we must start at the head of the list and then move all the way down to the predecessor of the current node. That gets costly if the current node is near the end of a long list. Moving to the predecessor can be done in one step if we add a predecessor reference to each node. Linked lists with successor and predecessor references are said to be *doubly linked*. Nodes in a doubly linked list look like this:

```
class Node
{
    String value; // Element
    Node next;    // Reference to successor
    Node prev;    // Reference to predecessor
    Node(String val, Node n, Node p)
    {
        value = val;
        next = n;
        prev = p;
    }

    Node(String val)
    {
        // Call the other constructor
        this(val, null, null);
    }
}
```

A node with `next` and `prev` references is shown in [Figure 19-5](#).




Figure 19-5 A node for a doubly linked list

[Figure 19-5 Full Alternative Text](#)

Working with doubly linked lists is not much different from working with singly linked lists. You use a `null` reference to represent an empty list and a reference to the head of the list to represent a

nonempty list. You set `next` to `null` if the node has no successor, and set `prev` to `null` if the node has no predecessor. An example of a doubly linked list with the names `Alan`, `Bob`, and `Carol` is shown in [Figure 19-6](#). The list uses a head pointer named `first` and a pointer to the last element named `last`.




Figure 19-6 A doubly linked list

[Figure 19-6 Full Alternative Text](#)

Adding a Node to a Doubly Linked List



VideoNote Adding a Node to a Doubly Linked List

Adding a node to a doubly linked list is similar to adding a node to a singly linked one. Say we want to add an element `e` to a doubly linked list. If the list is empty, the new node will have neither successor nor predecessor, so we can write the following:

```
first = new Node(e, null, null);
last = first;
```

Alternatively, we can use the one-arg constructor as follows:

```
first = new Node(e);
last = first;
```

There are two cases to consider when the list is nonempty: when the new node is to be added *before* the first element, and when the new node is to be added *after*. To add the new node before the first element, we can write

```
first = new Node(e, first, null);
first.next.prev = first;
```

Now let us consider the latter case. One way to proceed is first to obtain a reference to the node that will become the predecessor of the new node. Let us say such a reference, `pred`, has already been obtained. Now take another reference, `succ`, and set it to the node following `pred`:

```
Node succ = pred.next;
```

The new node to be added will go between `pred` and `succ`. Create this new node using a reference

```
named middle;
```

```
Node middle = new Node(e, succ, pred);
```

Executing this sequence results in the scenario shown in [Figure 19-7](#), with `Alex` playing the role of the element `e`. The goal is to insert `Alex` right after `Allan`.




Figure 19-7 The first step in inserting a node into a doubly linked list

[Figure 19-7 Full Alternative Text](#)

Once the situation in [Figure 19-7](#) is achieved, it only remains to set `pred.next` to `middle` and `succ.prev` to `middle`. But we must be careful: while we know that `pred` is not `null`, we cannot say the same about `succ`. Moreover, if `succ` is `null`, it means that the new node was added at the end of the list, and `last` must be set to point to it. The correct code to finish the job follows:

```
pred.next = middle;
if (succ == null)
    last = middle
else
    succ.prev = middle;
```

Removing a Node from a Doubly Linked List

Removing a node from a doubly linked list is not difficult, although it needs to be done with some care. First, obtain a reference to the node you want to remove. Given such a reference, say `target`, obtain references to the nodes on either side:

```
Node pred = target.prev;
Node succ = target.next;
```

Now you want to make `succ` the successor of `pred`, and `pred` the predecessor of `succ`:

```
pred.next = succ;
succ.prev = pred;
```

Unfortunately, this code does not work if the targeted node is either the first or the last in the list. For example, if the targeted node is the last in the list, then `succ` will be `null` and trying to set `succ.prev` to `pred` will cause an exception to be thrown. The correct code follows:

```
if (pred == null)
    first = succ;
else
    pred.next = succ;

if (succ == null)
    last = pred;
else
    succ.prev = pred;
```

Note that this code works correctly even when the node being removed is the only one in the list. An implementation of a doubly linked list class is given in [Code Listing 19-4](#).

Code Listing 19-4 (DLinkedList.java)

```
1  /**
2     The DLinkedList class implements a doubly
3     Linked list.
4  */
5
6  class DLinkedList
7  {
8      /**
9         The Node class stores a list element
10        and a reference to the next node.
11     */
12     private class Node
13     {
14         String value;           // Value of a list element
15         Node next;             // Next node in the list
16         Node prev;             // Previous element in the list
17
18         /**
19            Constructor.
20            @param val The element to be stored in the node.
21            @param n The reference to the successor node.
22            @param p The reference to the predecessor node.
23         */
24
25         Node(String val, Node n, Node p)
26         {
27             value = val;
28             next = n;
29             prev = p;
30         }
31     }
```

```

32     /**
33         Constructor.
34         @param val The element to be stored in the node.
35     */
36
37     Node(String val)
38     {
39         // Just call the other (sister) constructor
40         this(val, null, null);
41     }
42 }
43
44 private Node first;          // Head of the list
45 private Node last;          // Last element on the list
46
47 /**
48     Constructor.
49 */
50
51 public DLinkedList()
52 {
53     first = null;
54     last = null;
55 }
56
57 /**
58     The isEmpty method checks to see if the list
59     is empty.
60     @return true if list is empty, false otherwise.
61 */
62
63 public boolean isEmpty()
64 {
65     return first == null;
66 }
67
68 /**
69     The size method returns the length of the list.
70     @return The number of elements in the list.
71 */
72
73 public int size()
74 {
75     int count = 0;
76     Node p = first;
77     while (p != null)
78     {
79         // There is an element at p
80         count++;
81         p = p.next;
82     }
83     return count;
84 }
85
86 /**
87     The add method adds to the end of the list.
88     @param e The value to add.
89 */

```

```

90
91 public void add(String e)
92 {
93     if (isEmpty())
94     {
95         last = new Node(e);
96         first = last;
97     }
98     else
99     {
100         // Add to end of existing list
101         last.next = new Node(e, null, last);
102         last = last.next;
103     }
104 }
105
106 /**
107     This add method adds an element at an index.
108     @param e The element to add to the list.
109     @param index The index at which to add.
110     @exception IndexOutOfBoundsException
111     When the index is out of bounds.
112 */
113
114 public void add(int index, String e)
115 {
116     if (index < 0 || index > size())
117     {
118         String message = String.valueOf(index);
119         throw new IndexOutOfBoundsException(message);
120     }
121
122     // Index is at least 0
123     if (index == 0)
124     {
125         // New element goes at beginning
126         Node p = first;           // Old first
127         first = new Node(e, p, null);
128         if (p != null)
129             p.prev = first;
130         if (last == null)
131             last = first;
132         return;
133     }
134
135     // pred will point to the predecessor
136     // of the new node.
137     Node pred = first;
138     for (int k = 1; k <= index - 1; k++)
139     {
140         pred = pred.next;
141     }
142
143     // Splice in a node with the new element
144     // We want to go from pred-- succ to
145     // pred--middle--succ
146     Node succ = pred.next;

```

```

147     Node middle = new Node(e, succ, pred);
148     pred.next = middle;
149     if (succ == null)
150         last = middle;
151     else
152         succ.prev = middle;
153 }
154
155 /**
156     The toString method computes the string
157     representation of the list.
158     @return The string representation of the
159     linked list.
160 */
161
162 public String toString()
163 {
164     StringBuilder strBuilder = new StringBuilder();
165
166     // Use p to walk down the linked list
167     Node p = first;
168     while (p != null)
169     {
170         strBuilder.append(p.value + "\n");
171         p = p.next;
172     }
173     return strBuilder.toString();
174 }
175
176 /**
177     The remove method removes the element
178     at a given position.
179     @param index The position of the element
180     to remove.
181     @return The element removed.
182     @exception IndexOutOfBoundsException When
183     index is out of bounds.
184 */
185
186 public String remove(int index)
187 {
188     if (index < 0 || index >= size())
189     {
190         String message = String.valueOf(index);
191         throw new IndexOutOfBoundsException(message);
192     }
193
194     // Locate the node targeted for removal
195     Node target = first;
196     for (int k = 1; k <= index; k++)
197         target = target.next;
198
199     String element = target.value; // Element to return
200     Node pred = target.prev;       // Node before the target
201     Node succ = target.next;       // Node after the target
202
203     // Route forward and back pointers around
204     // the node to be removed

```

```

205         if (pred == null)
206             first = succ;
207         else
208             pred.next = succ;
209
210         if (succ == null)
211             last = pred;
212         else
213             succ.prev = pred;
214
215         return element;
216     }
217
218     /**
219      * The remove method removes an element from the list.
220      * @param element The element to remove.
221      * @return true if the element was removed, false otherwise.
222      */
223
224     public boolean remove(String element)
225     {
226         if (isEmpty())
227             return false;
228
229         // Locate the node targeted for removal
230         Node target = first;
231         while (target != null
232             && !element.equals(target.value))
233             target = target.next;
234
235         if (target == null)
236             return false;
237
238         Node pred = target.prev;           // Node before the target
239         Node succ = target.next;          // Node after the target
240
241         // Route forward and back pointers around
242         // the node to be removed
243         if (pred == null)
244             first = succ;
245         else
246             pred.next = succ;
247
248         if (succ == null)
249             last = pred;
250         else
251             succ.prev = pred;
252
253         return true;
254     }
255
256     public static void main(String [] args)
257     {
258         DLinkedList ll = new DLinkedList();
259         ll.add("Amy");
260         ll.add("Bob");
261         ll.add(0, "Al");

```



```

262         ll.add(2, "Beth");
263         ll.add(4, "Carol");
264         System.out.println("The elements of the list are:");
265         System.out.println(ll);
266     }
267 }

```

Program Output

```

The elements of the list are:
Al
Amy
Beth
Bob
Carol

```

This class can be used through a graphical user interface similar to that used with the `LinkedList1` class. The code that implements the interface is in the file *DLinkedListDemo.java*, and can be found in the book's source code, available at www.pearson.com/gaddis.

Circularly Linked Lists

A singly linked list is *circularly linked* if the successor of the last node is the first. A doubly linked list is *circularly linked* if the successor of the last node is the first, and the predecessor of the first node is the last. Working with circularly linked lists is similar to working with non-circularly linked lists: the main differences are when you are adding to an empty list and when you are removing the last node from a list.



Checkpoint

- 19.3 Consider adding an element e to a nonempty doubly linked list with head pointer `first` and a pointer `last` pointing to the last node. The new node is to be added before some node that is already in the list. Let `succ` be a reference to the node in the list that will be the successor of the new node. Write code that can be used to add e to the list.
- 19.4 Consider adding an element e to a nonempty doubly linked list with head pointer `first` and a pointer `last` pointing to the last node. Write code that can be used to add e to the end of the list.

19.4 Recursion on Linked Lists

CONCEPT: A linked list is an inherently recursive data structure, and many operations on linked lists have natural recursive implementations.

The *tail* of a nonempty linked list L is the linked list that remains when the head of L is removed. The fact that the tail of a linked list is itself a linked list means that linked lists have a recursive structure. Many operations on linked lists can be given recursive formulations that process the element at the head of the list and then recursively operate on the tail. The base case for such recursive formulations is usually the case in which the list to be operated on is empty. This fact is very useful in formulating recursive methods that work on linked lists, so we restate it for emphasis:

A linked list is either empty (base case), or consists of a head and a tail. The head is the first element of the list, and the tail is the linked list consisting of all the other elements.

Consider a linked list class `RLinkedList` whose outline is shown here:

```
class RLinkedList
{
    private Node
    {
        }
    private Node first = null;
}
```

This class differs from the `LinkedList1` class of [Section 19.2](#) in that it omits the reference that points to the last node in the list. We want to write recursive implementations of the `size`, `add`, and `remove` methods.

A Recursive Implementation of `size`

The recursive strategy to compute the size of a list mirrors the recursive nature of the linked list: If the list is empty, its size is 0; otherwise, the size of the list is the size of the tail plus 1. Note that the tail of a list `list` is `list.next`. A recursive implementation of `size` follows:

```
private int size(Node list)
{
    if (list == null)
        return 0; // Base case
    else
        return size(list.next) + 1; // Add 1 to size of tail
}
```

The public `size` method simply calls the private recursive version of `size`, passing it the reference to

the head of the list as follows:

```
public int size( )
{
    return size(first);
}
```

Other recursive methods of this class will similarly be private, and will be called by a corresponding public method.

A Recursive Implementation of the `add(String e, Node list)` Method

The `RLinkedList` class will have a recursive private method

```
Node add(String e, Node list)
```

that will be called by its public counterpart, `void add(String e)`. The recursive method will take a list `list` and an element `e` as parameters, and will return the list that results from adding `e` to the end of `list`. As in the case of `size`, the recursive strategy is based on the recursive nature of a linked list. If the list is empty, then the method returns a list of one item containing `e` as its only element. If the list is nonempty, then the element `e` is added to the end of the tail of the list as follows:

```
private Node add(String e, Node list)
{
    if (list == null)
    {
        // Base case
        return new Node(e);
    }
    else
    {
        // Add e to the end of the tail and use
        // the result to replace the tail
        list.next = add(e, list.next);
        return list;
    }
}
```

Notice the sequence of events in the nonbase case. First, the recursive call is used to return the result of adding `e` to the end of the tail:

```
add(e, list.next);
```

Second, this modified tail is used to replace the current tail in `list`:

```
list.next = add(e, list.next);
```

Third and finally, the method returns the original list with the tail thus replaced:

```
return list;
```

A Recursive Implementation of `add(int index, String e, Node list)`

This method inserts an element `e` at a given index within a list and returns the resulting list. First the method checks that `index` is within the proper range and throws an exception if it is not. Suppose therefore that the index is nonnegative, and that the size of the list is greater than or equal to the index. Consider two cases:

- The index is zero. The method needs to put `e` at the beginning of the list (which may or may not be empty) and return the result. This is accomplished by

```
return new Node(e, list);
```

- The index is greater than zero. Because the size of the list is greater than or equal to the index, the list cannot be empty, and therefore must have a head and a tail. The method needs to replace the tail with the result of adding `e` at position `index-1` in the tail and return the modified list:

```
list.next = add(index-1, e, list.next);  
return list;
```

In this method, the recursion is on the `index` parameter rather than on the `list`, and the base case is when the index is zero. The complete method is as follows:

```
private Node add(int index, String e, Node list)  
{  
    if (index < 0 || index > size())  
    {  
        String message = String.valueOf(index);  
        throw new IndexOutOfBoundsException(message);  
    }  
  
    if (index == 0)  
        return new Node(e, list);  
  
    list.next = add(index-1, e, list.next);  
    return list;  
}
```

A Recursive Implementation of the `remove` Methods

Each of the two recursive remove methods returns a pair of results: the node removed from the list, and the list that remains after said node has been removed. These two results can be embodied in a single object of type `RemovalResult`:

```
private class RemovalResult
{
    Node node;    // The node removed from the list
    Node list;    // The list remaining after the removal

    RemovalResult(Node remNode, Node remList)
    {
        node = remNode;
        list = remList;
    }
}
```

First, consider the method

```
RemovalResult remove(int index, Node list)
```

This method is supposed to return the node at the given index, and the remaining list, packaged together in a single `RemovalResult` object. Assume that `index` is nonnegative and is less or equal to one less than the size of list. Again, consider the following cases:

- *The index is zero.* Then the size of the list must be greater or equal to 1, so the list is nonempty. The method needs to remove the first node from `list`, package the node and remaining list (which will be `list.next`), and return the resulting `RemovalResult` object.

```
if (index == 0)
{
    // Remove the first node on list
    RemovalResult remRes;
    remRes = new RemovalResult(list, list.next);
    list.next = null;
    return remRes;
}
```

- *The index is nonzero.* In this case, the method needs to recursively remove the node at `index-1` from the tail:

```
RemovalResult remRes;
remRes = remove(index-1, list.next);
```

The method then replaces the tail with the results of the recursive call:

```
list.next = remRes.list;
```

Before returning to its caller, the method must modify the `RemovalResult` object returned from the recursive call by setting its remaining list field to `list` (as modified by the replacement of its tail):

```
remRes.list = list;
```

```
return remRes;
```

The second removal method is similar and can be found in [Code Listing 19-5](#).

Code Listing 19-5 (RLinkedList.java)

```
1  /**
2   * Linked List class with recursive methods.
3   */
4
5  class RLinkedList
6  {
7      /**
8       * The Node class will store a list element
9       * and a reference to the next node.
10     */
11
12     private class Node
13     {
14         String value;
15         Node next;
16
17         /**
18          * Constructor.
19          * @param val The element to store in the node.
20          * @param n The reference to the successor node.
21          */
22
23         Node(String val, Node n)
24         {
25             value = val;
26             next = n;
27         }
28
29         /**
30          * Constructor.
31          * @param val The element to be stored in the node.
32          */
33
34         Node(String val)
35         {
36             // Just call the other (sister) constructor
37             this(val, null);
38         }
39     }
40
41     private Node first; // List head
42
43     /**
44     * Constructor.
45     */
46
47     public RLinkedList()
48     {
```

```

49     first = null;
50 }
51
52 /**
53     The isEmpty method checks to see if the list is empty.
54     @return true if list is empty, false otherwise.
55 */
56
57 public boolean isEmpty()
58 {
59     return first == null;
60 }
61
62 /**
63     The size method returns the length of the list.
64     @return The number of elements in the list.
65 */
66
67 public int size()
68 {
69     return size(first);
70 }
71
72 /**
73     This size method recursively computes the
74     length of a linked list passed to it.
75     @param The linked list.
76     @return The number of items in the linked list.
77 */
78
79 private int size(Node list)
80 {
81     if (list == null)
82         return 0;
83     else
84         return size(list.next) + 1;
85 }
86
87 /**
88     The add method adds an element to the end of the list.
89     @param e The value to add to the end of the list.
90 */
91
92 public void add(String e)
93 {
94     // Replace first with result of adding e to first
95     first = add(e, first);
96 }
97
98 /**
99     This recursive private add method adds
100    an element e to the end of a list.
101    @param e The element to add to the list.
102    @param list The list to add e to.
103    @return The list resulting from adding e to its end.
104 */
105
106 private Node add(String e, Node list)

```

```

107 {
108     if (list == null)
109     {
110         // Base case
111         return new Node(e);
112     }
113     else
114     {
115         // Add e to the end of the tail and use
116         // the result to replace the tail
117         list.next = add(e, list.next);
118         return list;
119     }
120 }
121
122 /**
123  The add method adds an element e at place index
124  in this linked list.
125  @param index The place in the list to add an element.
126  @param e The element to add this the linked list.
127  @exception IndexOutOfBoundsException When index is
128          out of bounds.
129 */
130
131 public void add(int index, String e)
132 {
133     // Replace first with the result of adding
134     // e at index in first
135     first = add(index, e, first);
136 }
137
138 /**
139  This add method adds an element at an index in a list.
140  @param e The element to add to the list.
141  @param index The index at which to add the element.
142  @param list The list to add e to.
143  @return The list resulting from adding e.
144  @exception IndexOutOfBoundsException When index is
145          out of bounds.
146 */
147
148 private Node add(int index, String e, Node list)
149 {
150     if (index < 0 || index > size())
151     {
152         String message = String.valueOf(index);
153         throw new IndexOutOfBoundsException(message);
154     }
155     if (index == 0)
156         return new Node(e, list);
157
158     // 0 < index and index <= size so list is not empty
159     // Replace the tail with result of adding e at index - 1
160     // in the tail
161     list.next = add(index-1, e, list.next);
162     return list;
163 }

```



```

164
165 /**
166     The toString method computes the string
167     representation of the list.
168     @return The string representation of the linked list.
169 */
170
171 public String toString()
172 {
173     StringBuilder strBuilder = new StringBuilder();
174
175     // Use p to walk down the linked list
176     Node p = first;
177     while (p != null)
178     {
179         strBuilder.append(p.value + "\n");
180         p = p.next;
181     }
182     return strBuilder.toString();
183 }
184
185 /**
186     The RemovalResult class describes the results of
187     removing a node from a linked list.
188 */
189
190 private class RemovalResult
191 {
192     Node node;    // The node removed from the list
193     Node list;    // The list remaining after the removal
194     RemovalResult(Node remNode, Node remList)
195     {
196         node = remNode;
197         list = remList;
198     }
199 }
200
201 /**
202     The remove method removes the element at an index.
203     @param index The index of the element to remove.
204     @return The element removed.
205     @exception IndexOutOfBoundsException When index is
206             out of bounds.
207 */
208
209 public String remove(int index)
210 {
211     // Pass the job on to the recursive version
212     RemovalResult remRes = remove(index, first);
213     String element = remRes.node.value;    // Element to return
214     first = remRes.list;                  // Remaining list
215     return element;
216 }
217
218 /**
219     The private remove method recursively removes
220     the node at the given index from a list.
221     @param index The position of the node to remove.

```

```

222     @param list The list from which to remove a node.
223     @return The result of removing the node from the list.
224     @exception IndexOutOfBoundsException When index is
225             out of bounds.
226 */
227 private RemovalResult remove(int index, Node list)
228 {
229     if (index < 0 || index >= size())
230     {
231         String message = String.valueOf(index);
232         throw new IndexOutOfBoundsException(message);
233     }
234
235     if (index == 0)
236     {
237         // Remove the first node on list
238         RemovalResult remRes;
239         remRes = new RemovalResult(list, list.next);
240         list.next = null;
241         return remRes;
242     }
243
244     // Recursively remove the element at index-1 in the tail
245     RemovalResult remRes;
246     remRes = remove(index-1, list.next);
247
248     // Replace the tail with the results and return
249     // after modifying the list part of RemovalResult
250     list.next = remRes.list;
251     remRes.list = list;
252     return remRes;
253 }
254
255 /**
256     The remove method removes a given
257     element from linked list.
258 */
259
260 public boolean remove(String e)
261 {
262     RemovalResult remRes;
263     remRes = remove(e, first);
264
265     // Replace the list by the results of the removal
266     first = remRes.list;
267     if (remRes.node != null)
268         return true;
269     else
270         return false;
271 }
272
273 /**
274     The remove method recursively removes a
275     node containing a given element from
276     a specified list.
277     @param element The element to remove.
278     @param list The list to remove from.

```

```

279     @return the list containing
280     */
281
282     private RemovalResult remove(String e, Node list)
283     {
284         if (list == null)
285             return new RemovalResult(null, null);
286
287         // Is the first node on list the target of the removal?
288         if (list.value.equals(e))
289         {
290             RemovalResult remRes;
291             remRes = new RemovalResult(list, list.next);
292             list.next = null;
293             return remRes;
294         }
295
296         // Node to be removed is after the first node on list
297         // Recursively remove it from the tail
298         RemovalResult remRes;
299         remRes = remove(e, list.next);
300
301         list.next = remRes.list;
302         remRes.list = list;
303         return remRes;
304     }
305
306     public static void main(String [] args)
307     {
308         RLinkedList ll = new RLinkedList();
309         ll.add("Amy");
310         ll.add("Bob");
311         ll.add(0, "Al");
312         ll.add(2, "Beth");
313         ll.add(4, "Carol");
314         System.out.println("The members of the list are:");
315         System.out.print(ll);
316     }
317 }

```

The book's source code, available at www.pearson.com/gaddis, contains a file *RLinkedListDemo.java* that implements a graphical user interface for this class.

In the Spotlight:



Directed Graphs

A linked list defines a one-way relation among a set of nodes in such a way that from any of the nodes you can move to at most one other node. If you can move from a node *A* to another node *B*, we say that

A is *adjacent to B*, or that *B* is a *neighbor* of *A*. Using this terminology, we can see that a linked list is a collection of nodes in which every node may have at most one neighbor, and in which there is a path from a distinguished *head node* to every other node, as shown in [Figure 19-8](#).



Figure 19-8 The adjacency relation for nodes of a linked list

A directed graph is a data structure that is similar to a linked list, but with a much more general type of adjacency relation among its nodes. Specifically, a *directed graph* is a collection of nodes with an adjacency relation in which any node may have any number of neighbors. Directed graphs are often referred to as *digraphs*, and links between their nodes are called *edges*. An example of a directed graph is depicted in [Figure 19-9](#).




Figure 19-9 An example of a digraph

[Figure 19-9 Full Alternative Text](#)

A digraph does not have to have a head node from which you can reach every other node, nor does it have to be “connected.” Digraphs are very useful in modeling different types of systems. For example, a digraph can be used to model a computer network. In such a model, the nodes represent computers and the edges represent communication links.

The ideas we learned in this and preceding chapters can be used to create computer representations of digraphs. One such representation associates with each node *X* the list of its neighbors: that is, the list of nodes that can be reached from *X* by traversing a single edge. Such a list is called an *adjacency list* for *X*. For example, the digraph of [Figure 19-9](#) can be represented as shown in [Table 19-1](#).

Table 19-1 Representing a digraph by

adjacency lists

Node Neighbors

A B, C

B C, D, E

C A, D

D E

E

F G

G

The adjacency list representation can be realized by using a map that associates nodes with their adjacency lists. [Code Listing 19-6](#) initializes an adjacency list for the digraph of [Figure 19-9](#) and then prints it out. The program uses a `Map of String to List<String>`, which is defined in lines 12 and 13. In lines 22 through 27, the program acquires an entry set for the `Map` object and then uses the enhanced `for` loop to print out the adjacency list information. Entry sets for maps were discussed in [Section 18.4](#).

Code Listing 19-6 (AdjacencyList.java)

```
1 import java.util.*;
2 /**
3   This program demonstrates how to initialize an adjacency list
4   structure whose nodes are strings. It uses a map of strings (nodes)
5   to Lists of strings (nodes)
6 */
7
8 public class AdjacencyList
9 {
```

```

10 public static void main(String[] args)
11 {
12     Map<String, List<String>>
13         adjList = new HashMap<String, List<String>>>();
14     adjList.put("A", Arrays.asList("B", "C"));
15     adjList.put("B", Arrays.asList("C", "D", "E"));
16     adjList.put("C", Arrays.asList("A", "D"));
17     adjList.put("D", Arrays.asList("E"));
18     adjList.put("E", new ArrayList<String>());
19     adjList.put("F", Arrays.asList("G"));
20     adjList.put("G", new ArrayList<String>());
21
22     Set< Map.Entry<String, List<String>>>
23         entries = adjList.entrySet();
24     for (Map.Entry<String, List<String>> e : entries)
25     {
26         System.out.println(e.getKey() + ": " + e.getValue()) ;
27     }
28 }
29 }

```

Program Output

```

D: [E]
E: []
F: [G]
G: []
A: [B, C]
B: [C, D, E]
C: [A, D]

```

In practice, we do not build adjacency lists of strings, as shown here. Instead, we establish a mapping between the set of names of the nodes and a set of nonnegative integers $0..N-1$, where N is the number of nodes. In our example, if we number the nodes A through G in order, the adjacency list structure might be as shown in [Table 19-2](#).

Table 19-2 An adjacency list using integers to represent digraph nodes

Node Neighbors

```

0    1, 2

1    2, 3, 4

```

2 0, 3

3 4

4

5 6

6

The reason for the change is that, in practice, algorithms that work with adjacency lists often need to compare different nodes of the digraph for equality. It is more efficient to compare integers than it is to compare strings.

Data describing adjacency lists is usually stored in text files. Programs that work with digraphs read the data from such files and construct an in-memory representation of the adjacency list. The data stored in the file should follow a format that makes it easy for programs to *parse* the data. *Parsing* is the process of recognizing meaningful structure from syntax. In this case, we need to parse the file data to recognize which parts of the data represent nodes and which parts represent the adjacency lists for the nodes. [Table 19-3](#) uses the example digraph of [Figure 19-9](#) to illustrate a format that is relatively easy to parse.

Table 19-3 Storing digraph information in a file to simplify parsing

7

A 0

B 1

C 2

D 3

E 4

F 5

G 6

0 2 1 2

1 3 2 3 4

2 2 0 3

3 1 4

4 0

5 1 6

6 0

The file starts with the number N of nodes in the graph, in this case 7. This number is followed by N lines, with each line showing the name of a node and the number that will be used to represent the node in the adjacency list. Following this is another set of N lines, with each line specifying the neighbors reachable from a single node. Each line has the form

(node) (number of reachable neighbors) (list of reachable neighbors)

For example, the line for node 1 is

1 3 2 3 4

This line shows that the node with number 1 has 3 reachable neighbors and that these neighbors form the list 2 3 4.

19.5 Common Errors to Avoid

- Forgetting to check that a reference to a node is not null. Before you attempt to access the value stored in a node, or the `next` or `prev` references, make sure the reference is not null.
- Forgetting to check for special cases. Algorithms that operate on linked lists usually have to handle one or more special cases. Among the special cases to watch for are the following: an empty list when adding an element, a list of one element when removing an element, and the first or last node in the list when adding or removing a node.
- Incorrect handling of base cases in recursive linked list algorithms. Most recursive linked list algorithms do the recursion on the list parameter. In that case, the base case is usually when the list is empty. Algorithms that manipulate a list based on an index sometimes do the recursion on the index parameter: in that case the base case is when the index is zero.
- Using the head of the list to traverse the list. Do not use the head reference to walk down the list when you are traversing it, as it will destroy the structure of the list. Use an auxiliary reference initialized to the head.

Review Questions and Exercises

Multiple Choice and True/False

1. A list is a collection that .
 1. associates keys with elements
 2. assigns an index to each of its elements
 3. is implemented by the `JList` class
 4. none of the above
2. In Java, the first (head) element of a list .
 1. has a predecessor but no successor
 2. has both a predecessor and a successor
 3. has index 0
 4. has index 1
3. To add an element `e` just after a node referenced by `ref`, you should use the statement .
 1. `add(ref);`
 2. `ref.next = new Node(e);`
 3. `ref.next = new Node(e, ref.next);`
 4. `ref = new Node(e, ref.next);`
4. In general, adding an element `e` in *one step* to a singly linked list just before a node referenced by `ref` .
 1. cannot be done
 2. can be done if you use a loop
 3. can only be done when `ref` is the last node in the list
 4. can be done if you check for exceptions

5. To remove a node that is not the head of a singly linked list conveniently, you need to have a reference to .
1. the node that precedes it
 2. the node that comes after it
 3. the node itself
 4. none of the above
6. To easily move from any node to its successor, and from any node to its predecessor, and from the first to the last node and from the last to the first node, you need .
1. a singly linked list
 2. a singly and circularly linked list
 3. a doubly linked list
 4. a doubly and circularly linked list
7. A systematic method that starts at the beginning of the list and processes every node is called .
1. a walk down the list
 2. a list processing function
 3. a travesty of list processing
 4. a traversal
8. A method for removing the element at an index k in a list throws an exception unless .
1. k is negative
 2. k is nonnegative and less than the size of the list
 3. k is nonnegative and less than or equal to the size of the list
 4. k is positive and less than or equal to the size of the list
9. A method for adding an element at an index k in a list throws an exception unless .
1. k is negative
 2. k is nonnegative and less than the size of the list

3. k is nonnegative and less than or equal to the size of the list
 4. k is positive and less than or equal to the size of the list
10. A method that determines if a list is empty in a consistently written linked list class that uses a reference `first` to point to the first element, and a reference `last` to point to the last element, can return the value .
1. `first == null`
 2. `last == null;`
 3. `first == null || last == null`
 4. all of the above
11. **True or False:** A linked list is a recursive data structure.
12. **True or False:** A linked list is a looping data structure.
13. **True or False:** The head of a linked list is also a linked list.
14. **True or False:** The tail of a linked list is also a linked list.
15. **True or False:** The index of an element to be added to a list can be the same as the size of the list.
16. **True or False:** The index of an element to be removed from a list can be the same as the size of the list.
17. **True or False:** The statement
- ```
myList = new Node("Zoe", myList);
```
- adds "Zoe" to the end of the linked list `myList`;
18. **True or False:** There is a natural way to write a recursive method to check whether a linked list is empty.
19. **True or False:** The last item in a doubly linked list can sometimes have a successor.
20. **True or False:** The first item in a doubly linked list can sometimes have a successor.

## Find the Error

Find the error in each of the following code segments.

1. 

```
// Print the second element on
// a list myList of 3 elements
Node ref = myList;
ref ++ ;
System.out.print(ref.value);
```
2. 

```
// Print all elements in a list myList
Node ref = myList;
while (ref.next != null)
{
 System.out.print(ref.value + " ");
 ref = ref.next;
}
```
3. 

```
// Add a node to the beginning of
// a doubly linked list myList
myList = new Node("Abraham", myList, null);
```
4. 

```
// Remove the first node of a nonempty
// doubly linked list myList
myList = myList.next;
```
5. 

```
// The reference last points to the last
// node in a nonempty doubly linked list.
// Remove the last node from the list
last = last.prev;
```

## Algorithm Workbench

1. Write a recursive method `void print(Node ref)` that prints the values of all elements in the linked list whose first node is `ref`.
2. Write a nonrecursive method `reverse()` that reverses the order of the elements in a list. Assume that this method is to be added to the `LinkedList1` class in this chapter.
3. Write a method `reverse()` as in Algorithm Workbench 2, except implement the method to call a private recursive method `Node reverse(Node list)`. Implement this recursive method as well.
4. Add a method `String removeMin()` to the `LinkedList1` class in this chapter. The method removes and returns the minimum string (according to the usual dictionary order on strings) from the list. If the list is empty, the method returns `null`.

## Short Answer

1. Explain why recursive linked list methods should not be made public.

2. Explain how to modify the linked list classes in this chapter so that the size of the list can be computed more efficiently. Describe specifically the changes that have to be made to the fields and methods of the class.
3. What is the index of an element in a list?
4. Explain how contiguous allocation works to store elements of a list.
5. Explain how linked allocation works to store elements of a list.
6. Consult the online Java documentation and determine the differences between the `StringBuffer` and `StringBuilder` classes.

# Programming Challenges

## 1. Generic Linked Lists

Modify the doubly linked list class presented in this chapter so it works with generic types. Add the following methods drawn from the `java.util.List` interface:

- `void clear()`: remove all elements from the list.
- `E get(int index)`: return the element at position `index` in the list.
- `E set(int index, E element)`: replace the element at the specified position with the specified element and return the previous element.

Test your generic linked list class by processing a list of numbers of type `double`.

## 2. Linked List Sorting and Reversing

Modify the `LinkedList1` class presented in this chapter by adding `sort()` and `reverse()` methods. The `reverse` method reverses the order of the elements in the list, and the `sort` method rearranges the elements in the list so they are sorted in alphabetical order. Do not use recursion to implement either of these operations. Extend the graphical interface in the `LinkedList1Demo` class to support sort and reverse commands, and use it to test the new methods.

## 3. Recursive Linked List Operations



### Recursive Linked List Operations

Modify the `LinkedList1` class presented in this chapter by adding `sort()` and `reverse()` methods. The `reverse` method reverses the order of the elements in the list, and the `sort` method rearranges the elements in the list so they are sorted in alphabetical order. The class should use recursion to implement the sort and reverse operations. Extend the graphical interface in the `LinkedList1Demo` class to support sort and reverse commands, and use it to test the new methods.

## 4. Top Ten Gamers

Implement a class that maintains a list of the top 10 performers in a video game. An entry on the list consists of a name and score, and the list is kept sorted in descending order of scores. Here is an example of such a list when it has only four elements.

Spike        120

Whiz        105

G-Man      99

JediMaster 95

Use a class based on linked lists. The class should have a constructor that sets up an empty list, and a `void insert(String name, int score)` method that adds a name and a score pair to the list. The `insert` method puts the entry in the proper position so that the list stays sorted by score. The list should have a maximum size of 10. After the list has 10 elements, an attempt to add a name with a score that is less than or equal to the minimum score on the list is ignored, and adding a score that is greater than the minimum score causes an entry with the minimum score to be dropped from the list.

Test the score with a graphical user interface similar to `LinkedList1Demo.java`. The graphical interface should support a single command of the form

*insert name score*

An example of such a command is “insert Whiz 105.”

## 5. Wedding Lottery

A beautiful princess is to select a husband from  $n$  suitors. The suitors are assigned numbers  $0, \dots, n-1$  and seated at a round table, and a random sequence of  $n-1$  numbers  $x_0, x_1, \dots, x_{n-2}$ , all of them in the range  $0, \dots, n-1$  is selected. The princess starts at suitor 0, and moving in a clockwise direction, she counts off  $x_0$  positions and eliminates the suitor at that position. From there, she counts off another  $x_1$  positions and eliminates the suitor at that position. She continues doing this until one lucky guy is left: she marries him and they live happily ever after.

Assume that the princess’s position is always between two suitors. As an example, assume the suitors are  $A, B, C$ , and  $D$ , with  $A$  being numbered 0 and  $D$  being numbered 3. Let the random sequence be 0, 2, 1. Then the suitors will be eliminated in the order  $A, D$ , and  $C$ , and the lucky suitor will be  $B$ .

Write a class that uses linked lists to help the princess make her decision. The constructor for the



class will be passed the number  $n$  of suitors (say 4), followed by a list of  $n$  suitors' names (say  $A$ ,  $B$ ,  $C$ ,  $D$ ), followed by a list of  $n-1$  numbers used to “rotate” among the suitors (0, 2, 1). The current position starts at the beginning of the list. The class should have a method `rotate(int x)` that counts off  $x$  steps in the clockwise direction beginning at the current position. The class should have a method

```
List<String> getEliminated()
```

that returns a list, in order of elimination, of the names of the suitors already eliminated. The class should also have a method

```
List <String> getHopeful()
```

that returns a list of suitors not yet eliminated.

Write a driver program that obtains the input  $n$ , a list of  $n$  suitors' names, and a list of  $n-1$  integers to be used as “rotate steps.” Using the class you have written, the program prints the names of the rejected suitors (in the order of rejection), and then prints the name of the lucky guy.

## 6. Wedding Lottery II

Write an appropriate graphical user interface for the Wedding Lottery problem. The interface should allow the princess to specify the list of suitors using commands of the form *add name*. Once the princess has entered all the names, she presses a button labeled “Choose” to move from the input phase to the selection phase. In the selection phase, the user interface displays a list of suitors already rejected and another list of suitors still under consideration. The interface will have a text field in which the princess can enter a number for the next rotate step. As long as there is more than one suitor left, the program eliminates a suitor each time the princess types a number into the rotate step text field.

## 7. Adjacency List Input

Create a file containing the adjacency list data given at the end of this chapter, and then modify the program in [Code Listing 19-6](#) so that it reads the file and uses the data found in the file to build the adjacency list structure. Use an array of strings to keep track of the association between node names and node numbers that is given at the beginning of the file, and then use an array or an `ArrayList` of lists of numbers for the adjacency list structure. Print the adjacency list twice: once using numbers to name the nodes (as in [Table 19-3](#)) and again using the string names of the nodes, as in [Table 19-3](#).

# Chapter 20 Stacks and Queues

## TOPICS

1. [20.1 Stacks and Their Applications](#)
2. [20.2 Array Implementation of Stacks](#)
3. [20.3 Linked Implementation of Stacks](#)
4. [20.4 Queues and Their Applications](#)
5. [20.5 Array Implementation of Queues](#)
6. [20.6 Linked List Implementation of Queues](#)
7. [20.7 Generic Implementation of Stacks and Queues](#)
8. [20.8 Queues and Breadth-First Search](#)
9. [20.9 Common Errors to Avoid](#)

# 20.1 Stacks and Their Applications

CONCEPT: A stack is a collection of items that allows addition and removal of items in a last-in-first-out manner.

In [Chapter 18](#) you learned about linked lists, array lists, and other types of collection classes that are part of the Java Collection Framework. The `List` classes you learned about in that chapter are very general, and allow elements to be added and removed at any position in the list. There are many applications that use lists of objects in more restricted ways. For example, an application may need to maintain a list of objects that requires additions and removals only at the ends, or even at only one end of the list. In a list that requires additions and removals to take place at only one end, the only item available for removal is the item that was last added. Such a list is called a *stack*, and the end of the list at which additions and removals take place is called the *top of the stack*.

## Examples and Applications of Stacks

There are many examples of stacks in real life. For example, when a number of cars are parked single file in a narrow driveway, the last car parked must be the first one to leave. Another example is a stack of plates in a cafeteria. Workers add plates to the top of the stack, and each patron removes a plate from the top of the stack.

There are many problems, both in computer science and out of it, that can be solved with the use of stacks. For example, all programming languages allow programmers to define methods and call them from different places within the program. The computer must keep track of the *return address* for each method call, that is, the place in the program from which the method was called. The computer then uses this return address when the method executes a return statement. Whenever a method is called, the computer adds its return address to a stack. Later, when the method executes a return, the computer removes an address from the stack and uses it as the target address for the return statement. The stack is the right data structure to use because the last method called is the first to return.

## Stack Operations

A stack can be viewed as an abstract data type that supports three main operations, traditionally called push, peek, and pop. The *push* operation takes a single item and adds it to the top of the stack. The *peek* operation returns the item currently at the top of the stack, but does not remove it. The *pop* operation removes (and returns) the item currently stored at the top of the stack.

The Java Collection Framework provides a `Stack` class that supports generic types. The class is in the `java.util` package. [Table 20-1](#) lists most of its methods.

# Table 20-1 Stack class methods

**Stack<E>**  
**()**

**Constructor.**

**E push(E item)** Adds `item` to the top of the stack and returns the item added.

**E pop()** Removes and returns the item at the top of the stack. If the stack is empty, `pop` throws `EmptyStackException`.

**E peek()** Returns the item that is currently at the top of the stack, but does not remove it from the stack. If the stack is empty, `peek` throws `EmptyStackException`.

**boolean empty()** Returns `true` if this stack is empty and `false` otherwise.

[Code Listing 20-1](#) is a simple demonstration of the use of the `Stack` class. The program creates a stack and pushes `String` objects onto it. Then it keeps popping the stack and printing the values popped until the stack is empty. Notice that the last-in-first-out nature of the stack causes the strings to be popped in the reverse of the order in which they were pushed.

## Code Listing 20-1 (StackDemo1.java)

```
1 import java.util.*;
2
3 /**
4 * This program demonstrates the java.util.Stack class.
5 */
6
7 public class StackDemo1
8 {
9 public static void main(String [] args)
10 {
11 // Create a stack of strings and add some names
12 Stack<String> stack = new Stack<>();
13 String [] names = {"Al", "Bob", "Carol"};
14 System.out.println("Pushing onto the stack the names:");
15 System.out.println("Al Bob Carol");
16 for (String s : names)
17 stack.push(s);
18 }
```

```

19 // Now pop and print everything on the stack
20 String message = "Popping and printing all stack values:";
21 System.out.println(message);
22 while (!stack.empty())
23 System.out.print(stack.pop() + " ");
24 }
25 }

```

# Program Output

```

Pushing onto the stack the names:
Al Bob Carol
Popping and printing all stack values:
Carol Bob Al

```

# Stacks of Primitive Types

The `Stack` class provided by the Java Collections Framework does not directly accept values of primitive type. To use it with a primitive type such as `int` or `double`, you must use the corresponding wrapper class. For example, to have a stack of `int`, you must create a stack of `Integer`:

```
Stack<Integer> intStack = new Stack<>();
```

Then you can pass values of the appropriate primitive type to the `push` method, and the value will automatically be boxed. Likewise, popping a value from the stack and assigning to a variable of the appropriate primitive type will cause it to be unboxed. The use of stacks of primitive types is illustrated in the program of [Code Listing 20-2](#). The program pushes some numbers onto a stack, then pops and prints them.

# Code Listing 20-2 (StackDemo2.java)

```

1 import java.util.*;
2
3 /**
4 * This program demonstrates the use
5 * of stacks with primitive types.
6 */
7
8 public class StackDemo2
9 {
10 public static void main(String [] args)
11 {
12 Stack<Integer> intStack = new Stack<>();
13
14 // Push some numbers onto the stack
15 for (int k = 1; k < 10; k++)
16 intStack.push(k*k);

```

```
17
18 // Pop and print all numbers
19 while (!intStack.empty())
20 {
21 int x = intStack.pop();
22 System.out.print(x + " ");
23 }
24 }
25 }
```

## Program Output

81 64 49 36 25 16 9 4 1

# 20.2 Array Implementation of Stacks

CONCEPT: A stack can be implemented by using an array to hold the stack items and an integer to keep track of the top of the stack.

A stack is just a list that enforces last-in-first-out access. In this section and the next, we look at how arrays and linked lists can be used to implement stacks.

The idea is simple. You create an array large enough to hold the largest number of items you will need to have in the stack at any given time, say, four (the capacity of a stack in a real application would likely be much larger), and you use an integer index, `top`, to point to the next slot in the array that is available to receive an item. Thus, for a stack of integers, you would have:

```
int [] s = new int [4]; // Array holds stack elements
int top = 0; // Pointer to top of stack
```


The variable `top`, when used in this manner, is traditionally called the *stack top pointer*, or simply, the *stack pointer*.

## The Stack Push Operation

To push a new item `x`, first we check to see if there is room in the stack. If so, we place it in the slot pointed to by `top` and then increment `top` to point to the next slot. If the stack has no room, we throw an exception as follows:

```
if (top == s.length)
 throw new StackOverflowException();
else
{
 // Add the new item to the stack and update top
 s[top] = x;
 top++;
}
```

`StackOverflowException` is not a Java exception: it is one that we have to define. [Figure 20-1](#) illustrates a sequence of events that starts with an empty stack and adds two items.



# Figure 20-1 Pushing items onto a stack

[Figure 20-1 Full Alternative Text](#)

## The Stack `empty` Method

The boolean expression `top == 0` is true when the stack is empty and false when it is not, so the `empty` method simply returns the value of that expression:

```
return top == 0;
```

## The Stack `peek` and `pop` Methods

The two stack methods `peek` and `pop` are similar: both check to see if the stack is empty before attempting to access the item at the top of the stack, and throw `EmptyStackException` if the stack is empty. The `peek` method returns the item at the top of the stack as follows:

```
return s[top-1];
```

The `pop` method needs to return the same value as `peek`, but in addition, it needs to decrement the stack pointer to indicate that the stack slot currently holding the item at the top of the stack is becoming available. This is accomplished as follows:

```
top --;
return s[top];
```

The code we have outlined here will correctly implement a stack of integers. The stack class makes use of a separate file, *StackExceptions.java*, containing definitions of the following subclasses of the `RuntimeException` class. These two exception classes will be thrown by our stack methods.

```
class StackOverflowException extends RuntimeException
{
}
class EmptyStackException extends RuntimeException
{
}
```

The code for the stack class is given next, as [Code Listing 20-3](#).

## Code Listing 20-3 (ArrayStack.java)

```
1 /**
2 Array Implementation of a stack.
```



```

3 */
4
5 public class ArrayStack
6 {
7 private int [] s; // Holds stack elements
8 private int top; // Stack top pointer
9
10 /**
11 * Constructor.
12 * @param capacity The capacity of the stack.
13 */
14
15 public ArrayStack (int capacity)
16 {
17 s = new int[capacity];
18 top = 0;
19 }
20
21 /**
22 * The empty method checks for an empty stack.
23 * @return true if stack is empty.
24 */
25
26 public boolean empty()
27 {
28 return top == 0;
29 }
30
31 /**
32 * The push method pushes a value onto the stack.
33 * @param x The value to push onto the stack.
34 * @exception StackOverflowException When the
35 * stack is full.
36 */
37
38 public void push(int x)
39 {
40 if (top == s.length)
41 throw new StackOverFlowException();
42 else
43 {
44 s[top] = x;
45 top ++;
46 }
47 }
48
49 /**
50 * The pop method pops a value off the stack.
51 * @return The value popped.
52 * @exception EmptyStackException When the
53 * stack is empty.
54 */
55
56 public int pop()
57 {
58 if (empty())
59 throw new EmptyStackException();
60 else

```

```

61 {
62 top--;
63 return s[top];
64 }
65 }
66
67 /**
68 The peek method returns the value at the
69 top of the stack.
70 @return value at top of the stack.
71 @exception EmptyStackException When the
72 stack is empty.
73 */
74
75 int peek()
76 {
77 if (empty())
78 throw new EmptyStackException();
79 else
80 {
81 return s[top-1];
82 }
83 }
84 }

```

The `ArrayStackDemo` class, shown in [Code Listing 20-4](#), demonstrates the use of this stack. It has a `main` method that creates a stack with capacity 5, stores some integers on it, and then pops and prints them.

## Code Listing 20-4 (`ArrayStackDemo.java`)

```

1 /**
2 This class demonstrates the
3 use of the ArrayStack class.
4 */
5
6 public class ArrayStackDemo
7 {
8 public static void main(String [] arg)
9 {
10 String str; // Use for output
11 ArrayStack st = new ArrayStack(5);
12 str = "Pushing 10 20 onto the stack.";
13 System.out.println(str);
14 st.push(10);
15 st.push(20);
16 str = "Value at top of the stack is ";
17 System.out.println(str + st.peek());
18 str = "Popping and printing all values:";
19 System.out.println(str);
20 while (!st.empty())
21 System.out.print(st.pop() + " ");
22 }
23 }

```

# Program Output

```
Pushing 10 20 onto the stack.
Value at top of the stack is 20
Popping and printing all values:
20 10
```

## Stacks of Objects

The stack in [Code Listing 20-3](#) is a stack of integers. Writing a stack that works with other types is similar. For example, to write a stack that works with `String` objects, we would modify the `ArrayStack` class by changing all occurrences of `int` that refer to the type of the items being stored in the stack to `String`. In particular, the parameter and return types of `peek`, `pop`, and `push` would change as follows:

|                               |                |                                  |
|-------------------------------|----------------|----------------------------------|
| <code>int peek()</code>       | <i>becomes</i> | <code>String peek()</code>       |
| <code>int pop()</code>        | <i>becomes</i> | <code>String pop()</code>        |
| <code>void push(int x)</code> | <i>becomes</i> | <code>void push(String x)</code> |

The type of array that holds the stack items also needs to change as follows:

|                                    |                |                                        |
|------------------------------------|----------------|----------------------------------------|
| <code>int [ ] s;</code>            | <i>becomes</i> | <code>String [ ] s;</code>             |
| <code>s = new int[capacity]</code> | <i>becomes</i> | <code>s = new String[capacity];</code> |

One other change must be made to the `pop` method when the stack is being used to store object types: the array entry referring to the stack item being removed must be set to `null` to facilitate garbage collection of the item being removed from the stack as follows:

```
s[top-1] = null;
```

For example, the `pop` method modified to work with string objects would look like this:

```
public String pop()
{
 if (empty())
 throw new EmptyStackException();
 else
 {
 String retVal = s[top-1];
 s[top-1] = null; // Facilitate garbage collection
 top--;
 return retVal;
 }
}
```

Later in this chapter we will look at how we can use generic types to write a single stack class that works with values of all types.

A disadvantage of using an array implementation is the need to fix the size of the array at the time the stack is created. This causes a problem if the application using the stack needs to store more items than the size of the stack will allow. Basing the implementation of an array on a linked list avoids this problem.

## 20.3 Linked Implementation of Stacks

CONCEPT: A stack can be implemented by using a singly linked list to hold the stack items, and having the head of the list serve as the top of the stack.

A stack can be implemented by using any type of list object, including the `LinkedList` class in the `java.util` package, to hold the stack items. The `LinkedList` class, however, supports many operations that are not needed when all we need is a stack: for example, most applications of stacks do not need iterators, or the ability to add and remove items from the middle of the list. We can avoid this additional overhead by basing a stack on a simple linked list that we code ourselves.

The linked list will be based on the following `Node` class, declared private inside a `LinkedStack` class:

```
private class Node
{
 String value;
 Node next;
 Node(String val, Node n)
 {
 value = val;
 next = n;
 }
}
```

Inside the `LinkedStack` class, we will have a reference to a linked list of items. A reference `top` will point to the node at the front of the list: this node will be the one removed by the next call to `pop`. For example, starting with an empty stack and pushing the strings "Alice" and "Bob" in succession will result in the scenario shown in [Figure 20-2](#).




**Figure 20-2 A linked stack after pushing two items**

A further operation of pushing "Carol" on the stack would give the result shown in [Figure 20-3](#):



**Figure 20-3 The linked stack of [Figure 20-](#)**

## 2 after push of the string "Carol"

At this point, executing a pop operation would remove "Carol" from the stack and return us to the scenario shown in [Figure 20-2](#).

# Implementation of Stack Methods

The `LinkedStack` class will have a field `top` that points to the item at the top of the stack and simultaneously serves as the head of the linked list. The stack will be empty whenever the expression

```
top == null
```

is true, so the `empty` method simply returns the value of that expression:

```
public boolean empty()
{
 return top == null;
}
```

The `push` operation adds a new item `s` to the front of the list:

```
top = new Node(s, top);
```

The `peek` method returns the value field of the node at the front of the list:

```
return top.value;
```

The `pop` method saves the value in the node at the top, removes the top node from the list by moving the head pointer past it, and returns the saved value:

```
String retValue = top.value;
top = top.next;
return retValue;
```

Both `peek` and `pop` throw an exception if they are called on an empty stack. Finally, the `LinkedStack` class is equipped with a `toString` method that uses a `StringBuilder` object to create a string representation of the class. Having this method allows the contents of the entire stack `st` to be displayed by passing it to the `System.out.print` method:

```
System.out.println(st);
```

The `LinkedStack` class can be seen in [Code Listing 20-5](#).

## Code Listing 20-5 (LinkedStack.java)

```

1 /**
2 This program demonstrates how to write and
3 use a stack class based on linked lists.
4 */
5
6 class LinkedStack
7 {
8 /**
9 The Node class is used to implement the
10 linked list.
11 */
12
13 private class Node
14 {
15 String value;
16 Node next;
17 Node(String val, Node n)
18 {
19 value = val;
20 next = n;
21 }
22 }
23
24 private Node top = null; // Top of the stack
25
26 /**
27 The empty method checks for an empty stack.
28 @return true if stack is empty, false otherwise.
29 */
30
31 public boolean empty()
32 {
33 return top == null;
34 }
35
36 /**
37 The push method adds a new item to the stack.
38 @param s The item to be pushed onto the stack.
39 */
40
41 public void push(String s)
42 {
43 top = new Node(s, top);
44 }
45
46 /**
47 The Pop method removes the value at the
48 top of the stack.
49 @return The value at the top of the stack.
50 @exception EmptyStackException When the
51 stack is empty.
52 */
53
54 public String pop()
55 {
56 if (empty())
57 throw new EmptyStackException();
58 else

```

```

59 {
60 String retValue = top.value;
61 top = top.next;
62 return retValue;
63 }
64 }
65
66 /**
67 The peek method returns the top value
68 on the stack.
69 @return The value at the top of the stack.
70 @exception EmptyStackException When the
71 stack is empty.
72 */
73
74 public String peek()
75 {
76 if (empty())
77 throw new EmptyStackException();
78 else
79 return top.value;
80 }
81
82 /**
83 The toString method computes a string
84 representation of the contents of the stack.
85 @return The string representation of the
86 stack contents.
87 */
88
89 public String toString()
90 {
91 StringBuilder sBuilder = new StringBuilder();
92 Node p = top;
93 while (p != null)
94 {
95 sBuilder.append(p.value);
96 p = p.next;
97 if (p != null)
98 sBuilder.append("\n");
99 }
100 return sBuilder.toString();
101 }
102 }

```

[Code Listing 20-6](#) gives the `LinkedStackDemo` class, which demonstrates the use of the `LinkedStack` class.

## Code Listing 20-6 (LinkedStackDemo.java)

```

1 /**
2 This class demonstrates the use of the
3 LinkedStack class.
4 */

```



```

5
6 public class LinkedStackDemo
7 {
8 public static void main(String [] args)
9 {
10 LinkedStack st = new LinkedStack();
11 System.out.println("Pushing: Amy Bob Chuck");
12 System.out.println("Contents of Stack:");
13 st.push("Amy");
14 st.push("Bob");
15 st.push("Chuck");
16 System.out.println(st);
17 String name = st.pop();
18 System.out.println("Popped: " + name);
19 System.out.println("Contents of Stack:");
20 System.out.println(st);
21 }
22 }

```

## Program Output

```

Pushing: Amy Bob Chuck
Contents of Stack:
Chuck
Bob
Amy
Popped: Chuck
Contents of Stack:
Bob
Amy

```



## Checkpoint

1. 20.1 What is the common name for an operation that adds an element to a stack?
2. 20.2 What is the common name for an operation that removes an element from a stack?
3. 20.3 What should a stack method for removing an item do when the stack is empty?
4. 20.4 In what order are elements added and removed from a stack?
5. 20.5 Cite an example from everyday life where a collection of items behaves like a stack.

# 20.4 Queues and Their Applications

CONCEPT: A queue is a collection of items that is accessed in a first-in-first-out fashion.

You saw in the last section that a stack is a collection of items that is accessed in a last-in-first-out fashion. A queue is like a stack, except its items are added and removed in a first-in-first-out fashion. Queues, like stacks, have many applications in the real world. For example, at a grocery store checkout line, the first customer in the line is the first to be served.

Queues find many applications in computer science. Consider, for example, a network of computers that must share one printer. The printer server associates a queue with the printer, and print jobs that arrive from the various computers are added to this queue. Print jobs are removed from this queue and are serviced in a first-come-first-served fashion.

A queue can be viewed as a list where one end is designated as the *front* of the queue, and the other end is designated as the *rear*. We add items at the rear of the queue and remove items from the front. A queue is required to support the following operations:

- *enqueue*( $x$ ): add a new item  $x$  to the rear of the queue
- *dequeue*( ): remove and return the item at the front of the queue
- *empty*( ): check if the queue is empty
- *peek*( ): return, but do not remove, the item at the front of the queue

# 20.5 Array Implementation of Queues

CONCEPT: Fixed size arrays can be used to implement queues.

To implement a queue using an array, we need an array `q`, and two integer indices `front` and `rear`. Assuming we want to implement a queue that will hold strings, we declare these variables as follows:

```
String [] q;
int front, rear;
```




## Array Implementation of Queues

It is useful to establish a condition that specifies how the variables `front` and `rear` are being used, and require that all queue methods that manipulate `front` and `rear` preserve this condition. This condition, which we will refer to as the *queue invariant*, is as follows:

- `front` is the index of the slot in `q` that holds the next item that will be dequeued. This slot will normally be *filled*.
- `rear` is the index of the slot in `q` that will hold the next item that will be enqueued. This slot will normally be *empty*.

Initially, the queue is empty, and both `front` and `rear` are initialized to 0 as required by the queue invariant (the first element stored will be placed at 0, and the first element removed will come from 0). The queue invariant is illustrated in [Figure 20-4](#).



## Figure 20-4 Use of an array-based queue

### [Figure 20-4 Full Alternative Text](#)

Now suppose we want to add an item `s` to the queue. The queue invariant says that we should store `s` in the slot `q[rear]` and then increment `rear` so it points to the next empty slot as follows:

```
q[rear] = s;
rear ++;
```


Thus, after adding one item, `rear` will be 1, and after adding a second, `rear` will be 2. At this point `front` will still be 0. This is illustrated in the first two rows of [Figure 20-4](#).

Now consider the operation of removing an item. The item to be removed is `q[front]`. We save the value at that position so it can be returned to the caller, set the queue entry to `null`, and increment `front` so that it points to the next item that will be dequeued.

```
String value = q[front];
q[front] = null;
front ++;
```

This is illustrated in the third and fourth rows of [Figure 20-4](#). Notice that as long as we have not reached the end of the array, `rear` will increase as items are added, and `front` will increase as items are removed, and `rear` and `front` will be equal precisely when the queue is empty.

Now suppose we start with the queue shown in the last row of [Figure 20-4](#) and add two more items. At that point `front` will still be 2, and `rear` will be 0 because that is the next available empty slot in the array. That situation is shown in the top row of [Figure 20-5](#). We see that to make the best use of available space in the array, the indices `front` and `rear` need to wrap around to 0 whenever they get to the end of the array. Accordingly, the correct code for adding an item to the queue is



## Figure 20-5 Use of an array as a circular buffer

### [Figure 20-5 Full Alternative Text](#)

```
q[rear] = s;
rear++;
if (rear == q.length) rear = 0;
```

and the correct code for removing an item is

```
String value = q[front];
q[front] = null;
front++;
if (front == q.length) front = 0;
```

When array indices wrap around to the beginning of the array in this fashion, we say that we are using the array as a *circular buffer*.

Notice in [Figure 20-5](#) that `front` is also equal to `rear` when the array is completely full. This makes it impossible to distinguish a queue that is full from one that is empty just by looking at `front` and `rear`. One way to work around this is to have the queue keep a count of the number of items it contains: an enqueue operation increments this counter, and a dequeue operation decrements it. Then

we can tell that the queue is empty when this count is 0, and that the queue is full when the count is equal to the length of the array.

A complete implementation of a queue is given in [Code Listing 20-7](#). It uses a variable named `size` to keep track of the number of items in the queue, and has a method named `capacity` that returns the length of the array used inside the queue. Notice that the `enqueue` and `dequeue` methods check for an empty or full queue, and throw appropriately named exceptions when they cannot accomplish their designated task. The exceptions thrown are in a separate file named *QueueExceptions.java*:

```
class QueueOverflowException extends RuntimeException
{
}

class EmptyQueueException extends RuntimeException
{
}
```

The `ArrayQueue` class in [Code Listing 20-7](#) overrides the `toString` method to return a readable string representation of the contents of the queue. The string returned from this method encodes the current value of the `front` and `rear` indices, together with the value and position of each item in the queue. You can see the format of the string returned by the `toString` method, and how it encodes the state of the queue, in the Program Output following [Code Listing 20-8](#).

## Code Listing 20-7 (ArrayQueue.java)

```
1 import java.util.*;
2
3 /**
4 * The ArrayQueue class uses an
5 * array to implement a queue.
6 */
7
8 class ArrayQueue
9 {
10 private String [] q; // Holds queue elements
11 private int front; // Next item to be removed
12 private int rear; // Next slot to fill
13 private int size; // Number of items in queue
14
15 /**
16 * Constructor.
17 * @param capacity The capacity of the queue.
18 */
19
20 ArrayQueue(int capacity)
21 {
22 q = new String[capacity];
23 front = 0;
24 rear = 0;
25 size = 0;
26 }
```

```

27
28 /**
29 The capacity method returns the length of
30 the array used to implement the queue.
31 @return The capacity of the queue.
32 */
33
34 public int capacity()
35 {
36 return q.length;
37 }
38
39 /**
40 The enqueue method adds an element to the queue.
41 @param s The element to be added to the queue.
42 @exception QueueOverflowException When there
43 is no more room in the queue.
44 */
45
46 public void enqueue(String s)
47 {
48 if (size == q.length)
49 throw new QueueOverflowException();
50 else
51 {
52 // Add to rear
53 size++;
54 q[rear] = s;
55 rear++;
56 if (rear == q.length) rear = 0;
57 }
58 }
59
60 /**
61 The peek method returns the item
62 at the front of the queue.
63 @return element at front of queue.
64 @exception EmptyQueueException When
65 the queue is empty.
66 */
67
68 public String peek()
69 {
70 if (empty())
71 throw new EmptyQueueException();
72 else
73 return q[front];
74 }
75
76 /**
77 The dequeue method removes and returns
78 the element at the front of the queue.
79 @return The element at the front of the queue.
80 @exception EmptyQueueException When
81 the queue is empty.
82 */
83
84 public String dequeue()

```

```

85 {
86 if (empty())
87 throw new EmptyQueueException();
88 else
89 {
90 size--;
91 // Remove from front
92 String value = q[front];
93
94 // Facilitate garbage collection
95 q[front] = null;
96
97 // Update front
98 front++;
99 if (front == q.length) front = 0;
100
101 return value;
102 }
103 }
104
105 /**
106 The empty method checks to see if
107 this queue is empty.
108 @return true if the queue is empty and
109 false otherwise.
110 */
111
112 public boolean empty()
113 {
114 return size == 0;
115 }
116
117 /**
118 The toString method returns a
119 readable representation of the
120 contents of the queue.
121 @return The string representation
122 of the contents of the queue.
123 */
124
125 public String toString()
126 {
127 StringBuilder sBuilder = new StringBuilder();
128 sBuilder.append("front = " + front + "; ");
129 sBuilder.append("rear = " + rear + "\n");
130 for (int k = 0; k < q.length; k++)
131 {
132 if (q[k] != null)
133 sBuilder.append(k + " " + q[k]);
134 else
135 sBuilder.append(k + " ?");
136 if (k != q.length - 1)
137 sBuilder.append("\n");
138 }
139 return sBuilder.toString();
140 }
141 }

```



The program shown in [Code Listing 20-8](#) demonstrates the use of the `ArrayQueue` class.

## Code Listing 20-8 (`ArrayQueueDemo.java`)

```
1 /**
2 * The ArrayQueueDemo demonstrates the use
3 * of the ArrayQueue class.
4 */
5
6 public class ArrayQueueDemo
7 {
8 public static void main(String [] args)
9 {
10 String str; // Holds various string values
11
12 ArrayQueue queue = new ArrayQueue(4);
13 str = "Queue has capacity ";
14 System.out.println(str + queue.capacity());
15
16 // Add 4 names
17 String [] names
18 = {"Alfonso", "Bob", "Carol", "Deborah"};
19 System.out.println("Adding names: ");
20 for (String s : names)
21 {
22 System.out.print(s + " ");
23 queue.enqueue(s);
24 }
25 System.out.println("\nState of queue is: ");
26 System.out.println(queue);
27
28 // Remove 2 names
29 System.out.println("Removing 2 names.");
30 queue.dequeue(); queue.dequeue();
31 System.out.println("State of queue is: ");
32 System.out.println(queue);
33
34 // Add a name
35 System.out.println("Adding the name Elaine:");
36 queue.enqueue("Elaine");
37 System.out.println("State of queue is: ");
38 System.out.println(queue);
39 }
40 }
```

## Program Output

```
Queue has capacity 4
Adding names:
Alfonso Bob Carol Deborah
State of queue is:
front = 0; rear = 0
```

```

0 Alfonso
1 Bob
2 Carol
3 Deborah
Removing 2 names.
State of queue is:
front = 2; rear = 0
0 ?
1 ?
2 Carol
3 Deborah
Adding the name Elaine:
State of queue is:
front = 2; rear = 1
0 Elaine
1 ?
2 Carol
3 Deborah

```

# GUI Front End for the Queue Demo Program

Now we describe a graphical user interface to the queue implemented in [Code Listing 20-7](#). A screenshot of the user interface after the user starts with a queue of size four, adds three names, and then removes a name, is shown in [Figure 20-6](#).




**Figure 20-6 Screenshot of GUI front end for the queue program**

(Mark Schrier. Copyright Oracle, Inc.)

The user interface consists of a window with a pane at the top that has an array of text fields. At the bottom of the window, there is a text field used by the user to enter commands. The program recognizes the following commands:

- enqueue *x*: add *x* to the queue
- add *x*: add *x* to the queue (same as enqueue *x*)

- `dequeue`: remove an item from the queue
- `remove`: remove an item from the queue (same as `dequeue`)

The program sets an event handler on the command entry text field that is called whenever the user types a command. The handler retrieves the text of the command, parses it by using a `Scanner` object, identifies the command and its argument, and calls the appropriate queue method. The handler then calls the `refresh` method, passing it a string that represents the current state of the queue. The `refresh` method examines its argument to determine the position and value of each queue element, and then updates the queue view text fields at the top of the frame. [Code Listing 20-9](#) gives the rest of the details.

## Code Listing 20-9 (QueueDemo.java)

```

1 import java.util.Scanner;
2 import javafx.application.Application;
3 import javafx.event.*;
4 import javafx.geometry.*;
5 import javafx.scene.Scene;
6 import javafx.scene.control.*;
7 import javafx.scene.layout.*;
8 import javafx.stage.Stage;
9
10 /**
11 This program demonstrates the graphical display of a queue.
12 */
13
14 public class QueueDemo extends Application
15 {
16 public static void main(String [] args)
17 {
18 launch();
19 }
20
21 @Override
22 public void start(Stage stage) throws Exception
23 {
24 // Create a queue.
25 ArrayQueue queue = new ArrayQueue(4);
26
27 // Create the GUI to display the queue.
28 TextField[] qViewTextField = new TextField[4];
29 TextField commandEntryTextField = new TextField();
30
31 // Put the queue view text fields in a horizontal box.
32 HBox topHBox = new HBox(10);
33
34 for (int k = 0; k < qViewTextField.length; k++)
35 {
36 TextField tF = new TextField();
37 tF.setPrefColumnCount(5);
38 tF.setEditable(false);

```

```

39 topHBox.getChildren().add(tF);
40 qViewTextField[k] = tF;
41 }
42
43 // Put the label and commandEntryTextField in
44 // a horizontal box.
45 HBox hBox = new HBox(10);
46 hBox.setAlignment(Pos.CENTER);
47 hBox.getChildren().addAll(new Label("Command"),
48 commandEntryTextField);
49
50 // Use an outer VBox to hold all the GUI components
51 VBox outerPane = new VBox(10, topHBox, hBox);
52 outerPane.setPadding(new Insets(10));
53
54 // Set the event handler on the text field
55 EventHandler<ActionEvent> handler
56 = new CmdTextHandler(queue, qViewTextField);
57 commandEntryTextField.setOnAction(handler);
58
59 stage.setScene(new Scene(outerPane));
60 stage.setTitle("Queue Demo");
61 stage.show();
62 }
63 }
64
65 /**
66 * The handler for the command entry text field uses
67 * constructor parameters to store reference to
68 * the queue and the queue view text fields.
69 */
70 class CmdTextHandler implements EventHandler<ActionEvent>
71 {
72 private ArrayQueue queue;
73 private TextField[] qViewTextField;
74
75 public CmdTextHandler(ArrayQueue queue1, TextField[] qViewTfs)
76 {
77 queue = queue1;
78 qViewTextField = qViewTfs;
79 }
80
81 public void handle(ActionEvent evt)
82 {
83 TextField cmdTextField = (TextField) evt.getTarget();
84 String cmdText = cmdTextField.getText();
85 Scanner sc = new Scanner(cmdText);
86 if (!sc.hasNext())
87 {
88 return;
89 }
90 String cmd = sc.next();
91 if (cmd.equals("add") || cmd.equals("enqueue"))
92 {
93 String item = sc.next();
94 queue.enqueue(item);
95 refresh(queue.toString());
96 return;

```

```

97 }
98 if (cmd.equals("remove") || cmd.equals("dequeue"))
99 {
100 queue.dequeue();
101 refresh(queue.toString());
102 return;
103 }
104 }
105
106 /**
107 * The refresh method stores the current queue
108 * entries into the corresponding text fields of
109 * the queue view.
110
111 * @param qStr encodes the current contents of the queue.
112 */
113 private void refresh(String qStr)
114 {
115 Scanner sc = new Scanner(qStr);
116 sc.nextLine(); // Skip first, rear info
117 while (sc.hasNext())
118 {
119 int k = sc.nextInt();
120 String qEntry = sc.next();
121 qViewTextField[k].setText(qEntry);
122 }
123 }
124 }

```

# 20.6 Linked List Implementation of Queues

CONCEPT: Singly linked lists can be used to implement queues.

A simple linked list can be used to implement a queue. A reference *front* is used as the head of the linked list, and marks the end of the list from which items will be removed. A second reference, *rear*, is used to mark the end at which additions will take place. [Figure 20-7](#) shows a queue with three elements.




**Figure 20-7 A linked list representing a queue with three elements**

Notice that *front* and *rear* will be pointing to the same node whenever there is only one node in the list, and will both be set to `null` whenever the list is empty.

## Queue Initialization and Enqueuing of Items

Using the same `Node` class as in the `LinkedStack` program, we can define a class `LinkedListQueue` that has two fields, `front` and `rear`, to keep track of the list of queue items:

```
Node front = null;
Node rear = null;
```

To add an item `s` to an empty queue, we create a node containing `s` and set both `rear` and `front` to point to it as follows:

```
rear = new Node(s, null);
front = rear;
```

To add an item `s` to a nonempty queue, we create a node containing `s`, set `rear.next` to point to that node, and then move `rear` so that it points to the newly added node as follows:

```
rear.next = new Node(s, null);
rear = rear.next;
```

## Dequeuing Items

To remove and return a value from the queue, first we store the value to be returned. Then we move `front` down past the node being removed and return the stored value as follows:

```
String value = front.value;
front = front.next;
return value;
```

This code, however, will leave `rear` pointing to an item that has been removed if the item being removed is the last one in the list. In this case, `rear` must be set to `null`. The correct code for removing is therefore

```
String value = front.value;
front = front.next;
if (front == null)
 rear = null;
return value;
```

The code that implements the `LinkedList` class, with appropriate driver code is given in [Code Listing 20-10](#).

## Code Listing 20-10 (`LinkedList.java`)

```
1 /**
2 This class implements a queue based
3 on linked lists.
4 */
5
6 public class LinkedList
7 {
8 private class Node
9 {
10 String value;
11 Node next;
12 Node(String val, Node n)
13 {
14 value = val;
15 next = n;
16 }
17 }
18
19 private Node front = null;
20 private Node rear = null;
21
22 /**
23 The method enqueue adds a value
24 to the queue.
25 @param s The value to be added
26 to the queue.
27 */
28
29 public void enqueue(String s)
30 {
```

```

31 if (rear != null)
32 {
33 rear.next = new Node(s, null);
34 rear = rear.next;
35 }
36 else
37 {
38 rear = new Node(s, null);
39 front = rear;
40 }
41 }
42
43 /**
44 The empty method checks to see if
45 the queue is empty.
46 @return true if and only if queue
47 is empty.
48 */
49
50 public boolean empty()
51 {
52 return front == null;
53 }
54
55 /**
56 The method peek returns value at the
57 front of the queue.
58 @return item at front of queue.
59 @exception EmptyQueueException When the
60 queue is empty.
61 */
62
63 public String peek()
64 {
65 if (empty())
66 throw new EmptyQueueException();
67 else
68 return front.value;
69 }
70
71 /**
72 The dequeue method removes and returns
73 the item at the front of the queue.
74 @return item at front of queue.
75 @exception EmptyQueueException When
76 the queue is empty.
77 */
78
79 public String dequeue()
80 {
81 if (empty())
82 throw new EmptyQueueException();
83 else
84 {
85 String value = front.value;
86 front = front.next;
87 if (front == null) rear = null;
88 return value;

```



```

89 }
90 }
91
92 /**
93 * The toString method concatenates all strings
94 * in the queue to give a string representation
95 * of the contents of the queue.
96 * @return string representation of this queue.
97 */
98
99 public String toString()
100 {
101 StringBuilder sBuilder = new StringBuilder();
102
103 // Walk down the list and append all values
104 Node p = front;
105 while (p != null)
106 {
107 sBuilder.append(p.value + " ");
108 p = p.next;
109 }
110 return sBuilder.toString();
111 }
112 }

```

[Code Listing 20-11](#) demonstrates the use of the `LinkedList` class.

## Code Listing 20-11 (`LinkedListDemo.java`)

```

1 /**
2 * The LinkedListDemo class demonstrates
3 * the use of the LinkedList class.
4 */
5
6 public class LinkedListDemo
7 {
8 public static void main(String [] args)
9 {
10 LinkedList queue = new LinkedList();
11
12 // Add 4 names
13 String [] names =
14 {"Alfonso", "Bob", "Carol", "Deborah"};
15 System.out.println("Adding names: ");
16 for (String s : names)
17 {
18 System.out.print(s + " ");
19 queue.enqueue(s);
20 }
21
22 System.out.println("\nState of queue is: ");
23 System.out.println(queue);
24
25 // Remove 2 names

```

```

26 System.out.println("Removing 2 names.");
27 queue.dequeue(); queue.dequeue();
28 System.out.println("State of queue is: ");
29 System.out.println(queue);
30
31 // Add another name
32 System.out.println("Adding the name Elaine:");
33 queue.enqueue("Elaine");
34 System.out.println("State of queue is: ");
35 System.out.println(queue);
36 }
37 }

```

# Program Output

```

Adding names:
Alfonso Bob Carol Deborah
State of queue is:
Alfonso Bob Carol Deborah
Removing 2 names.
State of queue is:
Carol Deborah
Adding the name Elaine:
State of queue is:
Carol Deborah Elaine

```



## Checkpoint

1. 20.6 What is the common name for an operation that adds an element to a queue?
2. 20.7 In what order are elements added and removed from a queue?
3. 20.8 Cite an example of a queue from everyday life.
4. 20.9 In an array implementation of a queue, why is it necessary to treat the array as a circular buffer?

# 20.7 Generic Implementation of Stacks and Queues

CONCEPT: Classes that implement collections should be written to use generic types.

A generic collection class, whether a stack or a queue, has a great advantage: it can be used to hold items of any type. Implementing a generic collection is not much different from implementing a type-specific one: you simply parameterize the name of the class with a type parameter, say  $T$ , and then use  $T$  wherever you would need to use the type of the items being stored in the collection. For example, a skeleton for an array implementation of a generic stack class looks like this:

```
class GenStack<T>
{
 private T [] s; // Holds stack elements
 private int top; // Stack top pointer
 public GenStack (int capacity)
 {

 }

 public boolean empty() { return top == 0; }
 public T push(T x)
 {

 }

 public T pop()
 {

 }

}
```

There is a complication, however, because Java does not allow arrays of generic type to be instantiated. For example, the following statement to instantiate an array to hold the stack items will not compile:

```
T[] s = new T[capacity];
```

There is a workaround: first instantiate an array of `Object`, and then cast it to the desired generic type as follows:

```
T[] s = (T[]) new Object[capacity];
```

This compiles, albeit with a compiler warning. [Code Listing 20-12](#) shows a complete implementation of an array-based generic stack. The code uses the exception classes declared in the *StackExceptions.java* encountered earlier.

# Code Listing 20-12 (GenStack.java)

```
1 /**
2 This class implements a generic array
3 based stack.
4 */
5
6 public class GenStack<T>
7 {
8 private T [] s; // Body of stack
9 private int top; // Stack top pointer
10
11 /**
12 Constructor.
13 @param capacity The capacity of the stack.
14 */
15
16 public GenStack (int capacity)
17 {
18 s = (T[]) new Object [capacity];
19 top = 0;
20 }
21 /**
22 The empty method checks to see if
23 the stack is empty.
24 @return true if and only if the
25 stack is empty.
26 */
27
28 public boolean empty() { return top == 0; }
29
30 /**
31 The push method adds x to the stack.
32 @param x the value to be pushed onto
33 the stack.
34 @return the value that was pushed
35 onto the stack.
36 @exception StackOverFlowException When
37 the stack is full.
38 */
39
40 public T push(T x)
41 {
42 if (top == s.length)
43 throw new StackOverFlowException();
44 else
45 {
46 s[top] = x;
47 top ++;
48 return x;
49 }
50 }
51
52 /**
53 The pop method removes and returns the
```

```

54 item at the top of the stack.
55 @return item at the top of the stack.
56 @exception EmptyStackException When the
57 stack is empty.
58 */
59
60 public T pop()
61 {
62 if (empty())
63 throw new EmptyStackException();
64 else
65 {
66 T retVal = s[top-1];
67 s[top-1] = null;
68 top--;
69 return retVal;
70 }
71 }
72 }

```

The program in [Code Listing 20-13](#) shows how to use the `GenericStack` class.

## Code Listing 20-13 (`GenericStackDemo.java`)

```

1 /**
2 * This class demonstrates the use of
3 * the generic stack class GenStack
4 */
5
6 public class GenericStackDemo
7 {
8 public static void main(String [] arg)
9 {
10 GenStack <String> st = new GenStack<String>(5);
11 st.push("George");
12 st.push("Washington");
13 System.out.println(st.pop());
14 System.out.println(st.pop());
15 }
16 }

```

## Program Output

```

Washington
George

```

# 20.8 Queues and Breadth-First Search

You do not need to implement your own queue if you need to use one in your program. The Java Collection Framework provides many different classes that implement the `Queue` interface. One such class is the familiar `LinkedList` class that we encountered in [Chapter 19](#). The `Queue` interface declares the methods shown in [Table 20-2](#).

## Table 20-2 Some methods in the `Queue` interface

`boolean add(E, e)` Adds an element to the rear of the queue (*enqueue*)

`E peek()` Returns, but does not remove, the element at the front of the queue

`E remove()` Removes and returns the element at the front of the queue (*dequeue*)

`boolean isEmpty()` Checks to see if the queue is empty

## In the Spotlight:



### Directory Searching Using Breadth-First Search

Queues are very useful in solving problems that involve *breadth-first search*, a type of search through the nodes of a digraph. (Digraphs were introduced in the Spotlight section of [Chapter 19](#).) Breadth-first search starts searching at one node and then continues the search through other nodes in order of increasing distance from the start node. Here we show how breadth-first search can be used to search a directory and all contained subdirectories for a specified file or subdirectory. It is not necessary to understand what a digraph is to follow this presentation.

Files and directories in a computer system are stored in a hierarchical fashion, in which a directory may contain files and other directories. When a directory *D* is contained in another directory *E*, we

say that *D* is a child directory of *E*. This creates a situation in which directories can have children, children of children, and, in general, *descendants*. The problem here is to search through some initial directory and all its descendants to find a specified file *F* and print the path to *F* if it is found.

Breadth-first search is based on using a queue to store directories that are descendants of the initial directory. A descendant directory is placed on the queue if it has been encountered, but has not yet been searched. Initially, the queue contains only one directory: the initial directory. We proceed in stages. At each stage, we remove the directory *X* at the front of the queue and search it. If we find the desired file, we stop the search. Otherwise, we add all the child directories of *X* to the (rear of the) queue.

Suppose that the initial directory is *X*. In the beginning, the queue contains only *X*, and *X* will be removed from the queue and searched. If the process does not stop there, then all the children of *X* will be added to the queue. The children of *X* are the first-generation descendants. At the second stage, one of the first-generation descendants, say *Y*, will be removed and searched. If the desired file is not found, then all the children of *Y* (which are second-generation descendants of *X*) will be added to the rear of the queue for later processing. Because nodes will be removed from the queue and searched in the order in which they are added, all directories will be searched in increasing order of their generation. You can see from this that breadth-first search searches nodes in order of increasing distance from the start node.

We need several methods from the `java.util.File` class, as shown in [Table 20-3](#).

**Table 20-3 Some methods of the `File` class**

|                                       |                                                                                                                                                                                                                                                                            |
|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>String [] list()</code>         | Returns an array of strings representing names of the files or - directories stored in this directory. Returns <code>null</code> if the <code>File</code> object is not a directory or if there is an IO error in trying to determine the directory listing.               |
| <code>File [] listFiles()</code>      | Returns an array of <code>File</code> objects representing the files or - subdirectories stored in this directory. Returns <code>null</code> if the <code>File</code> object is not a directory or if there is an IO error in trying to - determine the directory listing. |
| <code>boolean isDirectory()</code>    | Returns <code>true</code> if this <code>File</code> object is a directory and not a regular file.                                                                                                                                                                          |
| <code>String getAbsolutePath()</code> | Returns the full pathname of the file or directory.                                                                                                                                                                                                                        |

The program in [Code Listing 20-14](#) follows closely the preceding description of breadth-first search. For brevity, the test data is hard coded into the main method, and the program searches for a file named “Menu.tex” starting in the directory “C:/Users/gcm.” As shown by the program output, this file is found in a descendant subdirectory named “Backup\CSC531.”

## Code Listing 20-14 (DirSearch.java)

```
1 /**
2 This program shows how to use breadth-first search
3 based on a queue to do a directory search.
4 */
5 import java.io.File;
6 import java.util.*;
7
8 public class DirSearch
9 {
10 public static void main(String[] args)
11 {
12 File initDir = new File("C:/Users/gcm");
13 String filePath = search(initDir, "Menus.tex");
14 if (filePath == null)
15 System.out.println("Not found");
16 else
17 System.out.println(filePath);
18 }
19
20 /**
21 This method searches a given directory and all its
22 subdirectories looking for specified file or subdirectory.
23 @param initDir : the initial directory to search.
24 @param searchFileName : the name of the file or subdirectory
25 to search for.
26 @return the full path name of the searched for file or directory.
27 */
28 static String search(File initDir, String searchFileName)
29 {
30 Queue<File> directoriesToSearch = new LinkedList<File>();
31 directoriesToSearch.add(initDir);
32
33 while (!directoriesToSearch.isEmpty())
34 {
35 // Get next directory to search
36 File currDir = directoriesToSearch.remove();
37 // Get contents of current directory
38 String[] dirContents = currDir.list();
39 // Directory contents will be null if there is
40 // a problem listing the directory contents
41 if (dirContents == null) continue;
42 // Do the directory contents contain the desired file?
43 if (Arrays.asList(dirContents).contains(searchFileName))
44 {
45 return currDir.getAbsolutePath();
46 }
47 // Desired file not in current directory
```



```
48 // Add all the children of this directory to the queue
49 // of directories to be searched
50 File [] childDirectories = currDir.listFiles();
51 for (File f : childDirectories)
52 {
53 if (f.isDirectory())
54 directoriesToSearch.add(f);
55 }
56 }
57 // No more directories left to search, so not found
58 return null;
59 }
60 }
```

# Program Output

C:\Users\gcm\Backup\CSC531

## 20.9 Common Errors to Avoid

- Forgetting to check for an empty stack or queue. You should always check to see if a stack is empty before calling `peek` or `pop`, and you should always check to see if a queue is empty before attempting to access or dequeue an element from the queue. If you forget to do this, the exception thrown will terminate your program.
- Handling an empty stack or queue by catching the exception thrown. Exception handling should be reserved for conditions your program cannot anticipate and prevent. Use the methods for checking if a collection is empty instead.
- Not maintaining the stack or queue invariant. If you are writing a stack or queue, establish a convention (invariant) for how you use the fields of your stack or queue class, and make sure all methods preserve the invariant.

# Review Questions and Exercises

## Multiple Choice and True/False

1. A collection that is accessed in *first-in-first-out* fashion is called .
  1. a stack
  2. a queue
  3. a linked list
  4. an array-based collection
2. A collection that is accessed in *last-in-first-out* fashion is called .
  1. a stack
  2. a queue
  3. a linked list
  4. none of the above
3. The order in which cars go through a toll booth is best described as .
  1. a stack
  2. a queue
  3. a linked list
  4. none of the above
4. The concept of seniority, which some employers use to hire and fire workers is .
  1. a stack
  2. a queue
  3. a linked list
  4. none of the above

5. The stack method that returns an element from the stack without removing it is .
1. `pop`
  2. `push`
  3. `peek`
  4. `spy`
6. If the stack method `push` is called on an empty stack, .
1. it throws an `EmptyStackException`
  2. it adds its argument to the stack
  3. it calls the stack method `empty`
  4. none of the above
7. **True or False:** You can use the JCF stack to directly create a stack of `int`.
8. **True or False:** If `pop` is called on an empty stack, it will not return until the user puts something on the stack.
9. **True or False:** When using an array to implement a stack, the `push` method will wrap around to the beginning of the stack when it reaches the end.
10. **True or False:** In a linked implementation of a queue, the references `front` and `rear` can only be equal if the queue is empty.

# Find the Error

Find the error in each of the following code segments:

```
1. // An array implementation of a stack
int pop()
{
 if (top == 0)
 throw new EmptyStackException();
 else
 {
 return s[top-1];
 top --;
 }
}
```

2. // A linked implementation of a queue  

```
void enqueue(int x)
{
 if (empty())
 {
 rear = new Node(x);
 front = rear;
 }
 else
 {
 rear = new Node(x);
 rear = rear.next;
 }
}
```
3. //A linked implementation of a stack  

```
int pop()
{
 if(empty())
 throw new EmptyStackException();
 return top.value;
}
```
4. // A linked implementation of a queue  

```
int dequeue()
{
 if(empty())
 throw new EmptyQueueException();
 int value = front.value;
 front ++;
 return value;
}
```
5. // An array implementation of a queue  

```
int dequeue()
{
 if (empty())
 throw new EmptyQueueException();
 int val = q[rear];
 rear ++;
 return val;
}
```

## Algorithm Workbench

1. A palindrome is a word that reads the same backward as forward. For example, the words *madam*, *radar*, *dad*, and *kayak* are all palindromes. Write a method that takes a parameter *s* of type `String` and uses a stack to see if *s* is a palindrome.
2. Write a class that uses an array to implement a stack of integers. The stack should have the following fields:

```
int []s; // Holds stack elements
int top = -1; // Points to last item pushed
```

The field `top` will point to an item that is actually stored on the stack, instead of pointing to the next available slot. The stack will set `top` to `-1` when it is empty. Write the constructor for this class, and all the usual stack methods.

3. Suppose that you have two stacks but no queues. You have an application that needs to use a queue. Explain how to use the two stacks to simulate a single queue.

## Short Answer

1. What is the name for a last-in-first-out collection?
2. What is the name for a first-in-first-out collection?
3. What is the name for a first-in-last-out collection?
4. What does the `pop` method do when it is called on an empty stack?
5. Why does the array entry containing an item that is being popped have to be set to `null`?
6. What problem would you encounter if you tried to use the head of a linked list for the rear of a queue, and use the other end of the list for the front?

# Programming Challenges

## 1. Double-Ended Queue

A *deque* (pronounced “deck”) is a list-based collection that allows additions and removals to take place at both ends. A deque supports the operations *addFront(x)*, *removeFront()*, *addRear(x)*, *removeRear()*, *size()*, and *empty()*. Write a class that implements a deque that stores strings using a doubly linked list that you code yourself. Demonstrate your class with a graphical user interface that allows users to manipulate the deque by typing appropriate commands in a `JTextField` component, and see the current state of the deque displayed in a `JTextArea` component. Consult the documentation for the `JTextArea` class for methods you can use to display each item in the deque on its own line.

## 2. Array-Based Deque

Implement a deque as described in Programming Challenge 1, except base your implementation on an array. The constructor for the class should accept an integer parameter for the capacity of the deque and create an array of that size. Use a graphical user interface based on an array of text fields similar to what is shown in [Figure 20-6](#). Test your deque class by constructing a deque with capacity 10.



Array-Based Deque

## 3. Prefix Expressions

An expression is in *prefix form* when operators are written before their operands. Here are some examples of prefix expressions and the values they evaluate to:

| Expression | Value |
|------------|-------|
| 12         | 12    |
| + 2 51     | 53    |

\* 5 7                      35

\* + 16 4 + 3 1 80

An expression (such as 12) that begins with an integer is a prefix expression that evaluates to itself. Otherwise, an expression is a prefix expression if it begins with an operator and is followed by two prefix expressions. In this latter case, the value of the expression is recursively computed from the values of its constituent prefix sub-expressions.

Write a program that allows the user to enter prefix expressions in a text field. The program reads the expression, evaluates it, and displays the value in a suitable GUI component. Assume that the user enters expressions that use only positive integers and the two operators + and \*. Your program should use a stack to store values of sub-expressions as they are computed, and another stack to store operators that have not yet been applied.

## 4. Properly Nested Delimiters

A Java program can have the following type of delimiters: {, }, (, ), [, and ]. In a correct Java program, these delimiters must be properly nested. Think of each left delimiter {, (, and [ as opening a scope, and think of each right delimiter }, ), and ] as closing a scope opened by a corresponding left delimiter. A string of characters containing these delimiters *has proper nesting of delimiters* if each scope that is opened is eventually closed, and the scopes are opened and closed in a last-opened-first-closed fashion.

Write a program that reads a file containing Java source code and checks it for proper nesting of delimiters. Your program should read the source code from the file and print it to the screen. If the file is properly nested, all of it is printed to the screen and a message is printed that the file is properly nested. If the file is not properly nested, then copying of the file to the screen stops as soon as improper nesting is detected, and your program prints a message that the file has errors. To simplify your task, you may assume these delimiters do not appear inside of comments and string literals, and that they do not appear in the program as character literals.

## 5. Tracing Genealogies

A file has genealogy data for a collection of  $N$  people. The first line of the file contains the integer  $N$  followed by  $N$  additional lines of data. Each of these additional lines specifies a list of children for a single person. The line starts with the name of the person, followed by the number of that person's children, followed by the names of the children. Here is an example of a file specifying genealogy information for 10 people.



```
10
Al 3 Beth Carol Dino
Beth 1 Pablo
Carol 3 Ben Alex Mary
Dino 0
Pablo 2 Angela Miguel
Ben 0
Alex 0
Mary 0
Angela 0
Miguel 0
```

For example, Al has three children named Beth, Carol, and Dino; Beth has one child named Pablo; and Dino has no children. You may assume that all names are unique.

Write a program which reads a file of genealogy information and then allows the user to enter pairs of names  $X$  and  $Y$ . The program then determines whether  $Y$  is a descendant of  $X$ , and if so, prints a list of names beginning with  $X$  and ending with  $Y$ , such that each person in the chain is a child of person preceding them on the list. Otherwise, the program states that  $Y$  is not a descendant of  $X$ .

# Chapter 21 Binary Trees, AVL Trees, and Priority Queues

## TOPICS

1. [21.1 Binary Trees and Their Applications](#)
2. [21.2 Binary Search Trees](#)
3. [21.3 AVL Trees](#)
4. [21.4 Priority Queues](#)
5. [21.5 Common Errors to Avoid](#)

# 21.1 Binary Trees and Their Applications

CONCEPT: A binary tree is like a list in which every node can have up to two successors.

A *binary tree* is a collection of nodes in which each node is associated with up to two successor nodes, respectively called the *left* and *right child*. Not every node in a binary tree will have two children: the left child may be missing, or the right child may be missing, or both nodes may be missing. A node in a binary tree that has no children is called a *leaf*.

A node that has children is said to be the *parent* of its children. For a nonempty collection of nodes to qualify as a binary tree, every node must have at most one parent, and there must be exactly one node with no parent. The one node that has no parent is called the *root* of the binary tree. An empty collection of nodes is regarded as constituting an empty binary tree.

A binary tree is similar in some ways to a linked list: the root of the binary tree corresponds to the head of the list, a child of a node in a binary tree corresponds to the successor of a node in a list, and the parent of a node in a binary tree corresponds to the *predecessor* of a node in a list. And of course, the analogue of the empty list is the empty binary tree.

[Figure 21-1](#) shows an example of a binary tree. The node labeled *A* is the root, and each arrow points from a parent to a child.





**Figure 21-1 A binary tree**

## Binary Tree Concepts

Let  $X$  be a node in a binary tree  $T$ . A node  $Y$  is a *descendant* of  $X$  if  $Y$  is on a path from  $X$  to a leaf of  $T$ . Thus, the descendants of  $X$  are  $X$  itself, all the children of  $X$ , all the children of the children of  $X$ , and so on. In [Figure 21-1](#), the descendants of  $C$  are  $C$ ,  $E$ ,  $F$ , and  $G$ . Any node  $X$  together with all its descendants forms a binary tree with root  $X$ : such a binary tree is called a *subtree* of  $T$ . In particular, the set of all descendants of the left child of the root of  $T$  forms a binary tree called the *left subtree* of  $T$ . Likewise, the set of all descendants of the right child of the root of  $T$  comprises the *right subtree* of  $T$ . [Figure 21-2](#) illustrates the left and right subtrees of the binary tree of [Figure 21-1](#).



## Figure 21-2 Left and right subtrees

[Figure 21-2 Full Alternative Text](#)

Although it is an abuse of language, we will sometimes speak of a subtree rooted at a child of a node  $X$  as being a subtree of  $X$ . According to this,  $\{B, D\}$  is the left subtree of  $A$ , and  $\{F, G\}$  is the right subtree of  $C$ .

An *ancestor* of  $X$  is any node on the path from  $X$  to the root of  $T$ . The ancestors of  $F$  in the binary tree of [Figure 21-1](#) are  $F$ ,  $C$ , and  $A$ .

## Applications of Binary Trees

Binary trees and their generalizations have many applications. They can be used to organize information in ways that allow efficient search and retrieval. Other forms of trees closely related to binary trees are used in database management systems to represent data, and in programming language compilers and interpreters to represent expressions. [Figure 21-3](#) shows the use of binary trees to

represent arithmetic expressions.




**Figure 21-3 Binary tree representation of expressions**

[Figure 21-3 Full Alternative Text](#)

## Representing Nodes of Binary Trees

A node in a binary tree is an object that contains a field for storing a value and two other fields for storing references to the left and right children. Here is a class that can be used to create nodes for a binary tree that stores integer data:

```
class Node
{
 int value; // Value to be stored
 Node left, right; // Left and right children
 Node(int val)
 {
 value = val;
 left = null;
 right = null;
 }
 Node(int val, Node left1, Node right1)
 {
 value = val;
 left = left1;
 right = right1;
 }
}
```

```
}
}
```

Notice the two constructors in the `Node` class. The constructor with one parameter is convenient when we need to create leaf nodes.

A binary tree is represented by a reference to its root node. Assuming we have integer values A, B, C, and D, we can build the binary tree shown in [Figure 21-4](#) by executing the following sequence of statements:




**Figure 21-4 A binary tree constructed with the `Node` class**

[Figure 21-4 Full Alternative Text](#)

```
Node dNode = new Node(D);
Node cNode = new Node(C, dNode, null);
Node bNode = new Node(B);
Node aNode = new Node(A, bNode, cNode);
```

## Traversing a Binary Tree

A *binary tree traversal* is a systematic method of doing some processing at every node of a binary tree. A common reason for traversing a binary tree is to print the values stored in it. There are three standard techniques for traversing binary trees. Each of these is most easily implemented as a recursive method called on the root node of the binary tree. They are as follows:

- Preorder traversal: process the data at the root node, traverse the left subtree, and then traverse the right subtree
- Inorder traversal: traverse the left subtree, process the data at the root node, and then traverse the right subtree
- Postorder traversal: traverse the left subtree, traverse the right subtree, and then process the data at the root node

Processing the data at a node is often called *visiting* the node.

A traversal method takes as parameter a reference `btree` to the root of a subtree. The method returns immediately if `btree` is `null`, otherwise, it visits the root and recursively traverses the two subtrees in the order prescribed by the traversal technique being used. Here is an implementation of inorder traversal that prints out all values stored in binary tree:

```
void inorder(Node btree)
{
 // If the tree is empty do nothing
 if (btree != null)
 {
 // Traverse the left subtree
 inorder(btree.left);
 // Visit the node
 System.out.print(btree.value + " ");
 // Traverse the right subtree
 inorder(btree.right);
 }
}
```

[Code Listing 21-1](#) shows a `Node` class that can be used as a basis for building a binary tree that stores integer values.

## Code Listing 21-1 (Node.java)

```
1 /**
2 * Node class.
3 */
4 public class Node
5 {
6 int value;
7 Node left, right;
8
9 // Constructor for leaf nodes.
10 Node(int val)
11 {
12 value = val;
13 left = null;
14 right = null;
15 }
16 }
```

```

17 // Constructor for non-leaf nodes.
18 Node(int val, Node leftChild, Node rightChild)
19 {
20 value = val;
21 left = leftChild;
22 right = rightChild;
23 }
24 }

```

[Code Listing 21-2](#) shows a class called `NodeUtilities`. This class can be used as a repository for various methods for working with node-based structures.

## Code Listing 21-2 (`NodeUtilities.java`)

```

1 /**
2 This class has various utility methods for working with nodes.
3 */
4 public class NodeUtilities
5 {
6 /**
7 Inorder traversal of a binary tree rooted at a node.
8 @param btree : The root of the binary tree.
9 */
10 static public void inorder(Node btree)
11 {
12 if (btree != null)
13 {
14 inorder(btree.left);
15 System.out.print(btree.value + " ");
16 inorder(btree.right);
17 }
18 }
19 }

```

[Code Listing 21-3](#) uses these two classes. The program creates a collection of nodes, links them together to form a binary tree, and then traverses the binary tree in inorder.

## Code Listing 21-3 (`BinaryTreeNodes.java`)

```

1 /**
2 This program demonstrates construction and traversal of binary trees.
3 */
4
5 public class BinaryTreeNodes
6 {
7 // Creates a binary tree from nodes and traverses it in
8 // in inorder.
9 public static void main(String[] args)
10 {
11 Node root = null; // Will be root of the binary tree.
12

```



```

13 Node aNode = new Node(10);
14 aNode.left = new Node(20);
15 Node dNode = new Node(40);
16 Node cNode = new Node(30, dNode, new Node(50));
17 aNode.right = cNode;
18 root = aNode;
19
20 System.out.print("Inorder traversal is: ");
21 NodeUtilities.inorder(root);
22 System.out.println();
23 }
24 }

```

## Program Output

```
Inorder traversal is: 20 10 40 30 50
```



## Checkpoint

1. 21.1 Consider the binary tree  $T$  shown in [Figure 21-1](#).
  1. List the ancestors of each node of  $T$ .
  2. List the descendants of each node of  $T$ .
  3. List all the leaves of  $T$ .
  4. List all the non-leaf nodes of  $T$ .
  5. Show the nodes of  $T$  as output by a preorder traversal.
  6. Show the nodes of  $T$  as output by a postorder traversal.
  7. Show the nodes of  $T$  as output by an inorder traversal.
2. 21.2 Show the output from a postorder traversal of the two trees in [Figure 21-3](#).
3. 21.3 Show the output from a preorder traversal of the two trees in [Figure 21-3](#).
4. 21.4 Using the `Node` class defined in this section, draw the binary tree constructed by the following statements:

```


Node left = new Node(10, new Node(20), new Node(50));
Node right = new Node(60, new Node(180), null);
Node tree = new Node(160, left, right);

```

and list the nodes of the resulting binary tree in preorder.

# Graphical Display of Binary Trees

You can graphically display a binary tree with the same recursive logic that is used to traverse the tree. The main idea is to recursively partition a view into three sections as shown in [Figure 21-5](#).



## Figure 21-5 Partitioning a view to display a binary tree

The partitioning is done by nesting views within views. An empty binary tree is displayed using a blank view. If the tree to be displayed is not empty, the display view is partitioned as in [Figure 21-5](#), the root is displayed at the top, and the left and right subtrees are recursively displayed in subviews nested side by side in the lower part of the containing view.

[Code Listing 21-4](#) shows a `GUINodeUtilities` class that has a recursive method for displaying a graphical view of a binary tree rooted at a node. If the tree is empty, its graphical representation is a blank `HBox` pane. A leaf node is displayed as a `HBox` containing a single un-editable text field that displays the value in the leaf node. A nonleaf node with a value *v* is displayed using a `BorderPane`: the value *v* is displayed in the TOP region, while the LEFT and RIGHT regions recursively display the left and right subtrees.

## Code Listing 21-4 (`GUINodeUtilities.java`)

```
1 import javafx.geometry.Insets;
2 import javafx.geometry.Pos;
3 import javafx.scene.control.Label;
4 import javafx.scene.control.TextField;
5 import javafx.scene.layout.BorderPane;
6 import javafx.scene.layout.HBox;
7 import javafx.scene.layout.Pane;
8
9 /**
10 * This class has a utility method for the
11 * graphical display of a binary tree.
12 */
13
14 class GUINodeUtilities
15 {
16 public static Pane getView(Node root)
```

```

17 {
18 // Empty tree is displayed as an empty HBox
19 if (root == null) { return new HBox(); }
20
21 // Leaf node displayed as a single un-editable text field
22 // in a HBox.
23 if (root.left == null && root.right == null)
24 {
25 TextField tF = new TextField(String.valueOf(root.value));
26 tF.setPrefColumnCount(2);
27 tF.setEditable(false);
28 HBox hBox = new HBox();
29 hBox.setAlignment(Pos.CENTER);
30 hBox.getChildren().add(tF);
31 hBox.setPadding(new Insets(5));
32 return hBox;
33 }
34 // Non-leaf node uses a BorderPane
35 // that puts the value in the TOP region, and left and right
36 // subtrees in the LEFT and RIGHT regions.
37 BorderPane pane = new BorderPane();
38 pane.setStyle("-fx-border-color: gray");
39 pane.setPadding(new Insets(10));
40
41 // Value of a non-leaf node is displayed in a text field
42 // in a HBox in the TOP region of the BorderPane.
43 String value = String.valueOf(root.value);
44 TextField tF2 = new TextField(value);
45 tF2.setPrefColumnCount(2);
46 tF2.setEditable(false);
47 HBox valueHBox = new HBox();
48 valueHBox.setAlignment(Pos.CENTER);
49 valueHBox.getChildren().add(tF2);
50 valueHBox.setPadding(new Insets(5));
51 pane.setTop(valueHBox);
52
53 // Left and right subtrees are recursively displayed
54 // in the LEFT and RIGHT region of the BorderPane.
55 pane.setLeft(getView(root.left));
56 pane.setCenter(new Label(" ")); // For spacing.
57 pane.setRight(getView(root.right));
58
59 return pane;
60 }
61 }

```

[Code Listing 21-5](#) shows a program that uses this utility method to display a binary tree in a scene on a JavaFX stage. The result of running the program can be seen in [Figure 21-6](#).




**Figure 21-6 GUI display of a binary tree**

(Mark Schrier. Copyright Oracle, Inc.)

## Code Listing 21-5

### (GUIBtreeDisplayDemo.java)

```

1 import javafx.application.Application;
2 import javafx.geometry.Insets;
3 import javafx.scene.Scene;
4 import javafx.scene.layout.VBox;
5 import javafx.stage.Stage;
6
7 /**
8 * This program demonstrates the GUI display of a
9 * binary tree.
10 */
11
12 public class GUIBtreeDisplayDemo extends Application
13 {
14 public static void main(String [] args)
15 {
16 launch();
17 }
18
19 @Override
20 public void start(Stage stage) throws Exception
21 {
22 // Create the tree.
23 Node root;
24 Node aNode = new Node(10);
25 aNode.left = new Node(20);
26 Node dNode = new Node(40);
27 Node cNode = new Node(30, dNode, new Node(50));
28 aNode.right = cNode;
29 root = aNode;

```

```

30
31 // Display the binary tree in a scene on the stage.
32 VBox vbox = new VBox(GUINodeUtilities.getView(root));
33 vbox.setPadding(new Insets(10));
34 Scene scene = new Scene(vbox);
35 stage.setScene(scene);
36 stage.setTitle("Display of Binary Tree");
37 stage.show();
38 }
39 }

```

# Class Implementation of Binary Trees

As we have seen, a binary tree is implemented as a reference to a root node that is an object of a suitable `Node` class such as shown in [Code Listing 21-1](#). A simple class implementation might use a private nested `Node` class like this:

```

class BinaryTree
{
 private Node { /* Node fields and methods */ }
 private Node root;
}

```

The class used to represent nodes is usually an implementation detail that should not concern users of the class, so the above approach makes sense. The class could then proceed to add public methods to work with binary tree objects.

Even so, it will be convenient for us to take a different approach. We will not nest the `Node` class, and we will declare the `root` field as *protected* instead of *private*:

```

class Node
{
 /* Node fields and methods */
}
class BinaryTree
{
 protected Node root;
}

```

Making `root` protected will make it accessible to subclasses, allowing us to use inheritance to write different types of binary trees. Not nesting `Node` will allow us to have different binary tree classes that refer to the same `Node` class, thus avoiding unnecessary repetition of code.

# 21.2 Binary Search Trees

CONCEPT: Binary search trees organize information to facilitate efficient search.

In [Chapter 16](#), you saw how binary search can be used to search quickly for a given item in a sorted array. Binary search trees are special types of binary trees that can be used to implement a form of binary search that does not require arrays. Look at the binary tree in [Figure 21-7](#).




**Figure 21-7 A binary search tree**

[Figure 21-7 Full Alternative Text](#)

Notice that the values stored in the left subtree (7, 13, and 19) are all less than the value stored at the root (25), and that the values stored in the right subtree (30, 50, 51, and 65) are all greater than the value stored in the root. In fact, the tree has the following property: at every node  $N$ , all values stored in the left subtree of  $N$  are less than the value stored at  $N$ , and all values stored in the right subtree of  $N$  are greater than the value stored at  $N$ . A binary tree with this property is called a *binary search tree*.

Binary search trees have the potential to facilitate rapid searches. Suppose we want to determine if  $X$  is in the tree. Starting at the root, we make our way down some path of the binary tree, comparing  $X$  to each node that we come to. At each node  $N$ , we continue the search in the left subtree if  $X$  is less than  $N$ , or in the right subtree if  $X$  is greater. We stop when we find  $X$  at some node, or when we come to a leaf that is not  $X$ . Assuming that the tree is reasonably balanced, we see that we will discard approximately half of the remaining tree every time we make a comparison. Thus, searching a binary tree with  $n$  nodes will require approximately  $\log n$  comparisons. Naturally, binary search can only be

used with items that support a comparison operation with the property that for any two distinct elements  $x$  and  $y$ , either  $x$  “is less than”  $y$  or  $x$  “is greater than”  $y$ . Sets that support such a comparison operation are said to be *totally* or *linearly* ordered.

The following method searches a binary search tree for an element  $x$ . It returns `true` if the tree contains  $x$  and `false` otherwise.

```
private boolean contains(int x, Node bstree)
{
 if (bstree == null)
 return false;
 if (x == bstree.value)
 return true;
 if (x < bstree.value)
 {
 // Recursively look in left subtree
 return contains(x, bstree.left);
 }
 else
 {
 // Recursively look in right subtree
 return contains(x, bstree.right);
 }
}
```

# Adding a Value to a Binary Search Tree

Like most operations on binary trees, adding a value to a binary search tree has a natural recursive implementation. The method

```
Node add(int x, Node bstree)
```

adds  $x$  to the binary search tree rooted at `bstree` and returns a reference to the root of the resulting binary search tree. First, the method checks to see if `bstree` is empty: if so, it creates and returns a new node containing  $x$ . If `bstree` is not empty, the method compares  $x$  to the element stored in the root: if  $x$  is less, it adds  $x$  to the left subtree and makes the resulting tree the new left subtree. If  $x$  is greater, it adds  $x$  to the right subtree and makes the resulting tree the new right subtree. Then it returns a reference to the root:

```
private Node add(int x, Node bstree)
{
 if (bstree == null)
 return new Node(x);
 // bstree is not null.
 if (x < bstree.value)
 {
 // Add x to the left subtree and replace the
 // current left subtree with the result
 bstree.left = add(x, bstree.left);
 }
 else
```

```

{
 // Add x to the right subtree
 bstree.right = add(x, bstree.right);
}
return bstree;
}

```

# Removing a Value from a Binary Search Tree



## Removing a Value from a Binary Search Tree

To remove a value from a binary search tree, you must locate the node  $N$  that contains the value and remove it from the tree. There are three cases to consider:

- The node  $N$  has no children. Remove  $N$  by setting the reference to it in its parent node to `null`.
- The node  $N$  has one child. Remove  $N$  and promote the child of  $N$  to take the place of its parent in the tree, as shown in [Figure 21-8](#).




**Figure 21-8 Before and after removal of a node with one child**



- The node  $N$  has two children. Remove  $N$  and replace it with the largest node in its left subtree, as shown in [Figure 21-9](#). Notice that the largest node in the left subtree of  $N$  cannot have a right child, so removing it reduces it to the simpler case of removing a node with at most one child.




**Figure 21-9 Removing a node with two children**

[Figure 21-9 Full Alternative Text](#)

The tree that results from removing the node with 90 in the tree of [Figure 21-9](#) is shown in [Figure 21-10](#).




**Figure 21-10** The tree of [Figure 21-9](#) after removing 90

[Figure 21-10 Full Alternative Text](#)

Let us now consider the task of coding the removal of a value  $x$  from a binary search tree  $T$ . This operation yields two objects: a node  $N$  that contains  $x$ , and the tree that remains after  $N$  has been removed from  $T$ . The following class will be used to represent the result of such an operation:

```
private class RemovalResult
{
 Node node; // Removed node
 Node tree; // Remaining tree
 RemovalResult(Node node, Node tree)
 {
 this.node = node;
 this.tree = tree;
 }
}
```

This class will be private to the class that represents the binary search tree. The method that performs the actual removal,

```
RemovalResult remove(Node bTree, int x)
```

will take a reference `bTree` to the root of a binary search tree and an integer value  $x$ . The `remove` method returns `null` if  $x$  is not found in `bTree`.

The algorithm for removing  $x$  considers three cases:

1. The relation `x < bTree.value` holds. In this case `x` is recursively removed from the left subtree:

```
RemovalResult result = remove(bTree.left, x);
if (result == null)
 return null;
bTree.left = result.tree;
result.tree = bTree;
return result;
```

If the recursive call returns `null`, then `x` is not in the left subtree and therefore cannot be in the tree, so the original call also returns `null`. Otherwise, the remaining tree from the recursive call becomes the left subtree of the remaining tree from the original call, and the `result` object is accordingly modified and returned.

2. The relation `x > bTree.value` holds. In this case `x` is recursively removed from the right subtree. This case is a mirror image of the preceding.
3. The relation `x == bTree.value` holds, so that the root node `bTree` contains the value to be removed. In this case, `bTree` becomes the removed node and the remaining tree is assembled from the left and right subtrees of `bTree`. The assembly of the remaining tree differs according to whether `bTree` has zero, one, or two children, and proceeds as previously explained and shown in [Figures 21-7](#), [21-8](#), and [21-9](#). The programming details are as follows:

Consider first the case where `bTree` has zero children. The removed node is just `bTree`, and the remaining tree is empty:

```
if (bTree.right == null && bTree.left == null)
 return new RemovalResult(bTree, null);
```

Next, consider the case of `bTree` having one child. In this case `bTree` is the removed node, and the nonempty subtree of `bTree` becomes the remaining tree:

```
Node node = bTree; // Removed node
Node tree; // Remaining tree

// Remaining tree is the one nonempty subtree
if (bTree.left != null)
 tree = bTree.left;
else
 tree = bTree.right;
node.left = null;
node.right = null;
return new RemovalResult(node, tree);
```

Finally, consider the case of `bTree` having two children. Again, `bTree` becomes the removed node. The removal of `bTree` leaves two subtrees, and we have to find a way to combine these two subtrees into one binary search tree. This can be done by removing the largest node from the left subtree and using it to replace the removed root. The removal of the largest node from a binary search tree is accomplished by the following recursive method:

```
RemovalResult removeLargest(Node bTree)
```

To see the logic behind the algorithm used by this method, observe that the largest node in a nonempty binary search tree is the root if the root has no right subtree; otherwise, it is the largest node in the right subtree.

```
RemovalResult removeLargest(Node bTree)
{
 if (bTree == null)
 return null;
 if (bTree.right == null)
 {
 // Root is the largest node
 Node tree = bTree.left; // Remaining tree
 bTree.left = null;
 return new RemovalResult(bTree, tree);
 }
 else
 {
 // Remove the largest node from the right subtree
 RemovalResult remResult = removeLargest(bTree.right);
 bTree.right = remResult.tree;
 remResult.tree = bTree;
 return remResult;
 }
}
```

## Implementation of Binary Search Trees

We are now ready to implement a binary tree search class. We begin with a base binary tree class, shown in [Code Listing 21-6](#), that contains a single method for checking whether the tree is empty.

### Code Listing 21-6 (BinaryTree.java)

```
1 /**
2 Base binary tree class
3 */
4 public class BinaryTree
5 {
6 protected Node root;
7
8 // Check if the binary tree is empty.
9 public boolean isEmpty()
10 {
11 return root == null;
12 }
13 }
```

A binary search tree is a special type of binary tree, so it makes sense to make the `BinarySearchTree` class a subclass of `BinaryTree`. [Code Listing 21-7](#) shows the implementation of

that class, together with all the methods that we have already discussed.

## Code Listing 21-7 (BinarySearchTree.java)

```
1 /**
2 This class implements a binary search tree.
3 */
4
5 public class BinarySearchTree extends BinaryTree
6 {
7 /**
8 This class represents the result of removing
9 a node from a binary tree.
10 */
11 private class RemovalResult
12 {
13 Node node; // Removed node
14 Node tree; // Remaining tree
15
16 RemovalResult(Node node, Node tree)
17 {
18 this.node = node;
19 this.tree = tree;
20 }
21 }
22
23 /**
24 The public add method adds a value to the tree by
25 calling a private add method and passing it the
26 root of the tree.
27 @param x The value to add to the tree.
28 @return true.
29 */
30 public boolean add(int x)
31 {
32 root = add(x, root);
33 return true;
34 }
35
36 /**
37 The contains method checks to see if a value is
38 in the binary tree.
39 @param x The value to check for.
40 @return true if x is in the tree, false otherwise.
41 */
42 public boolean contains(int x)
43 {
44 // Call the private recursive method.
45 return contains(x, root);
46 }
47
48 /**
49 The add method adds a value to the search tree.
50 @x the value to add.
51 @param bstree The root of the binary search tree.
```

```

52 @return The root of the resulting binary search tree.
53 */
54 private Node add(int x, Node bstree)
55 {
56 if (bstree == null)
57 {
58 return new Node(x);
59 }
60 // bstree is not null.
61 if (x < bstree.value)
62 {
63 // Add x to the left subtree and replace
64 // the current left subtree with the result.
65 bstree.left = add(x, bstree.left);
66 } else
67 {
68 // Add x to the right subtree.
69 bstree.right = add(x, bstree.right);
70 }
71 return bstree;
72 }
73
74 /**
75 The method contains checks whether an item is in
76 a binary search tree.
77 @param x The item to check for.
78 @param bstree The binary tree to look in.
79 @return true if found, false otherwise.
80 */
81 private boolean contains(int x, Node bstree)
82 {
83 if (bstree == null)
84 {
85 return false;
86 }
87 if (x == bstree.value)
88 {
89 return true;
90 }
91 if (x < bstree.value)
92 {
93 // Recursively look in left subtree.
94 return contains(x, bstree.left);
95 } else
96 {
97 // Recursively look in right subtree.
98 return contains(x, bstree.right);
99 }
100 }
101
102 /**
103 The remove method removes a value from the
104 binary search tree.
105 @param x The value to remove.
106 @returns true if x was removed, false if x not found.
107 */
108 public boolean remove(int x)
109 {

```

```

110 RemovalResult result = remove(root, x);
111 if (result == null)
112 {
113 return false;
114 } else
115 {
116 root = result.tree;
117 return true;
118 }
119 }
120
121 /**
122 * This remove method removes a value a from a
123 * binary search tree and returns the removed node
124 * and the remaining tree wrapped in a RemovalResult object.
125 *
126 * @param bTree The binary search tree.
127 * @param x The value to be removed.
128 * @return null if x is not found in bTree.
129 */
130 private RemovalResult remove(Node bTree, int x)
131 {
132 if (bTree == null)
133 {
134 return null;
135 }
136 if (x < bTree.value)
137 {
138 // Remove x from the left subtree.
139 RemovalResult result = remove(bTree.left, x);
140 if (result == null)
141 {
142 return null;
143 }
144 bTree.left = result.tree;
145 result.tree = bTree;
146 return result;
147 }
148 if (x > bTree.value)
149 {
150 // Remove x from the right subtree.
151 RemovalResult result = remove(bTree.right, x);
152 if (result == null)
153 {
154 return null;
155 }
156 bTree.right = result.tree;
157 result.tree = bTree;
158 return result;
159 }
160 // x is in this root node.
161 // Is it a leaf?
162 if (bTree.right == null && bTree.left == null)
163 {
164 return new RemovalResult(bTree, null);
165 }
166

```

```

167 // Does the node have two children?
168 if (bTree.right != null && bTree.left != null)
169 {
170 // Remove largest node in left subtree and
171 // make it the root of the remaining tree.
172 RemovalResult remResult = removeLargest(bTree.left);
173 Node newRoot = remResult.node;
174 newRoot.left = remResult.tree;
175 newRoot.right = bTree.right;
176
177 // Prepare the result to be returned.
178 bTree.left = null;
179 bTree.right = null;
180 return new RemovalResult(bTree, newRoot);
181 }
182 // The node has one child
183 Node node = bTree;
184 Node tree;
185 if (bTree.left != null)
186 {
187 tree = bTree.left;
188 } else
189 {
190 tree = bTree.right;
191 }
192 node.left = null;
193 node.right = null;
194 return new RemovalResult(node, tree);
195 }
196
197 /**
198 The removeLargest method removes the largest node
199 from a binary search tree.
200 @param bTree: The binary search tree.
201 @return The result of removing the largest node.
202 */
203 private RemovalResult removeLargest(Node bTree)
204 {
205 if (bTree == null)
206 {
207 return null;
208 }
209 if (bTree.right == null)
210 {
211 // Root is the largest node
212 Node tree = bTree.left;
213 bTree.left = null;
214 return new RemovalResult(bTree, tree);
215 } else
216 {
217 // Remove the largest node from the right subtree.
218 RemovalResult remResult
219 = removeLargest(bTree.right);
220 bTree.right = remResult.tree;
221 remResult.tree = bTree;
222 return remResult;
223 }
224 }

```



# Graphical Display of Binary Search Trees

The graphical display of binary trees can be achieved by using the `getView()` method of the `NodeUtilities` classes of [Code Listing 21-4](#). The `getView()` method requires access to the root node of the binary tree to be displayed. Because the `root` field of `BinaryTree` is protected, the `getView()` method must be called from a subclass of `BinaryTree`.

What we need is a `DisplayableBtree` class, a subclass of `BinaryTree` that is able to wrap a pre-existing binary tree and produce a graphical view of the wrapped binary tree. The definition of this class starts out like this:

```
public class DisplayableBtree extends BinaryTree
{
 private BinaryTree tree;
 public DisplayableBtree(BinaryTree t)
 {
 tree = t;
 }
}
```

Notice that this class has only one constructor, so you cannot create an instance of `DisplayableBtree` unless you already have an existing binary tree that you want to display:

```
BinarySearchTree tree1 = new BinarySearchTree();
tree1.add(50);
DisplayableBtree dTree = new DisplayableBTree(tree1);
```

Notice also that `DisplayableBtree` has two fields: a protected field `root` that it inherits from `Binarytree`, and a private field `tree` that references another binary tree. So, `DisplayableBtree` is a binary tree that contains, or wraps, another binary tree. Moreover, the inherited `root` field will always be `null`, because unlike `BinarySearchTree`, our `BinaryTree` class has no `add` method that can be used to add elements. To make the inherited `isEmpty()` method useful, we override it so that it checks the emptiness of the wrapped binary tree:

```
public class DisplayableBtree extends BinaryTree
{
 private BinaryTree tree;
 public DisplayableBtree(BinaryTree t)
 {
 tree = t;
 }

 @Override
 public boolean isEmpty()
 {
 return tree.isEmpty();
 }
}
```

Finally, we complete the class by adding a `getView()` method to return a graphical view of the binary tree. The completed class is shown in [Code Listing 21-8](#).

## Code Listing 21-8 (DisplayableBtree.java)

```
1 import javafx.scene.layout.Pane;
2
3 public class DisplayableBtree extends BinaryTree
4 {
5 private final BinaryTree tree;
6 public DisplayableBtree(BinaryTree t)
7 {
8 tree = t;
9 }
10
11 @Override
12 public boolean isEmpty()
13 {
14 return tree.isEmpty();
15 }
16
17 // Return a view of the binary tree suitable for
18 // display
19 public Pane getView()
20 {
21 return GUINodeUtilities.getView(tree.root);
22 }
23 }
```

## A Program for Testing the Binary Search Tree Class

The program for testing the class of [Code Listing 21-8](#) will have a graphical user interface that allows users to type commands to add, remove, and search for items in the binary search tree. The user interface consists of a main window with a command entry text field at the bottom and an uneditable text field at the top used to display the results from executing user commands. The center of the window is used to display in a hierarchical tree form, the elements that are currently in the search tree. [Figure 21-11](#) shows the interface after the user has executed commands to add 50, 100, 75, 80, 12, 105, and 77 in succession.




**Figure 21-11 A binary tree after adding some values**

(Mark Schrier. Copyright Oracle, Inc.)

[Figure 21-11 Full Alternative Text](#)

Removing the node with 100 causes the program to replace it with 80, resulting in the display shown in [Figure 21-12](#).



# Figure 21-12 The binary tree of [Figure 21-11](#) after removing 100

(Mark Schrier. Copyright Oracle, Inc.)

## [Figure 21-12 Full Alternative Text](#)

The program allows the following commands:

- `add x`: Adds the number `x` to the binary search tree, for example: `add 12`.
- `remove x`: Removes `x` from the tree, for example: `remove 12`.
- `contains x`: Checks if `x` is in the tree, for example: `contains 12`. The program will display `true` in the command result label at the top of the frame if `x` is found in the tree, and `false` otherwise.
- `isempty`: Checks if there are any values in the tree, for example: `isempty`. The program will display `true` in the command result text field if the tree is empty, and `false` otherwise.

The program for testing the `BinarySearchTree` class follows. It creates a binary search tree on which the `add`, `remove`, and other operations will be performed, and then wraps the search tree with a `DisplayableBtree`:

```
BinarySearchTree binTree = new BinarySearchTree();
DisplayableBtree displayBinTree = new DisplayableBtree(binTree);
```

The program then builds the user interface and attaches an event handler on the command entry text field at the bottom of the window. Every time the user enters a command, the handler method reads the command from the text field and executes the command on the binary search tree `binTree`. The old view of the tree is then removed from the user interface and replaced with the new view. The details are in [Code Listing 21-9](#).

## Code Listing 21-9 (GUIBSearchTreeDemo.java)

```
1 import java.util.Scanner;
2 import javafx.application.Application;
3 import javafx.event.ActionEvent;
4 import javafx.event.EventHandler;
5 import javafx.geometry.Insets;
6 import javafx.geometry.Pos;
7 import javafx.scene.Scene;
8 import javafx.scene.control.Label;
9 import javafx.scene.control.TextField;
10 import javafx.scene.layout.HBox;
```

```

11 import javafx.scene.layout.VBox;
12 import javafx.stage.Stage;
13
14 /**
15 This class is a GUI interface for testing the
16 the binary search tree class.
17 */
18
19 public class GUIBSearchTreeDemo extends Application
20 {
21 @Override
22 public void start(Stage stage) throws Exception
23 {
24 // Set up the binary tree to manipulate
25 BinarySearchTree binTree = new BinarySearchTree();
26
27 // This binary tree wraps the real binary tree
28 // so it can be displayed.
29 DisplayableBtree
30 displayBinTree = new DisplayableBtree(binTree);
31
32 // Displays the result of the user command.
33 TextField cmdResultTextField = new TextField();
34 cmdResultTextField.setEditable(false);
35
36 // Used by the user to enter a command.
37 TextField cmdTextField = new TextField();
38
39 // HBox1 holds cmdResultTextField and its label.
40 HBox hBox1 = new HBox(10);
41 Label cmdResultLabel = new Label("Command Result");
42 hBox1.getChildren().addAll(cmdResultLabel, cmdResultTextField);
43
44 // Place holder for the binary tree view.
45 HBox bTreeView = new HBox();
46 bTreeView.setAlignment(Pos.CENTER);
47
48 // HBox2 holds the cmdTextField and its label
49 HBox hBox2 = new HBox(10);
50 Label cmdLabel = new Label("Command");
51 hBox2.getChildren().addAll(cmdLabel, cmdTextField);
52
53 // Put the three HBox objects in a VBox
54 VBox outerPane = new VBox(10);
55 outerPane.setPadding(new Insets(10));
56 outerPane.getChildren().addAll(hBox1, bTreeView, hBox2);
57
58 // Command handler needs to access to the bTreeView
59 // so it can set the binary tree view, and to the
60 // cmdResultTextField so it can set the command result.
61 // Also needs access to the binary tree and its displayable
62 // twin.
63 EventHandler<ActionEvent>
64 handler = new CmdHandler(bTreeView, cmdResultTextField,
65 binTree, displayBinTree);
66 cmdTextField.setOnAction(handler);
67

```

```

68 // Set up the stage
69 stage.setTitle("Binary Search Tree Operations");
70 stage.setScene(new Scene(outerPane));
71 stage.show();
72 }
73
74 public static void main(String [] args)
75 {
76 launch(args) ;
77 }
78
79 }
80
81 // Event Handler class
82 class CmdHandler implements EventHandler<ActionEvent>
83 {
84 // These class variables reference
85 // data in the calling environment.
86 private HBox view;
87 private TextField cmdResultTextField;
88 private BinarySearchTree binTree;
89 private DisplayableBtree displayBinTree;
90
91 // Constructor stores references to data in
92 // the calling environment.
93 public CmdHandler(HBox tView, TextField cmdResTf,
94 BinarySearchTree tree,
95 DisplayableBtree dTree)
96 {
97 view = tView;
98 cmdResultTextField = cmdResTf;
99 binTree = tree;
100 displayBinTree = dTree;
101 }
102
103 public void handle(ActionEvent evt)
104 {
105 // Get the command from the event target.
106 TextField cmdTf = (TextField)evt.getTarget();
107 String cmdStr = cmdTf.getText();
108
109 // Use a scanner to read the command and its parts.
110 Scanner sc = new Scanner(cmdStr);
111 String cmd = sc.next();
112 switch (cmd)
113 {
114 case "add":
115 {
116 // Perform the operation on the tree.
117 int value = sc.nextInt();
118 binTree.add(value);
119
120 // Replace the old view of the tree.
121 view.getChildren().clear();
122 view.getChildren().add(displayBinTree.getView());
123 cmdResultTextField.setText(" ");
124 }

```

```

125 return;
126 case "isempty":
127 {
128 String resultText
129 = "Tree empty? " + binTree.isEmpty();
130 cmdResultTextField.setText(resultText);
131 }
132 return;
133 case "contains":
134 {
135 int value = sc.nextInt();
136 String resultText = "Contains " + value + "?: "
137 + binTree.contains(value);
138 cmdResultTextField.setText(resultText);
139 }
140 return;
141 case "remove":
142 {
143 int value = sc.nextInt();
144 binTree.remove(value);
145
146 // Replace the old view of the tree.
147 view.getChildren().clear();
148 view.getChildren().add(displayBinTree.getView());
149 cmdResultTextField.setText(" ");
150 }
151 return;
152 }
153 }
154 }

```

## 21.3 AVL Trees

**CONCEPT:** Operations on binary search trees are more efficient when the tree is balanced. An AVL tree is a binary search tree that satisfies a balance condition at each one of its nodes.

Because most operations on binary search trees trace out a path from the root to some node, the amount of work they do is proportional to the length of the path. Therefore it is desirable to have binary search trees that are “shallow” rather than “deep.” A binary tree can be kept shallow by imposing conditions on the height of its subtrees. The *height* of a binary tree is the length of the longest path from the root to a leaf. (The length of a path is the number of edges, or links, in the path: it is one less than the number of nodes in the path.) A binary tree with one node has height 0, and by convention, the empty binary tree has height  $-1$ .

An *AVL tree* is a binary search tree in which the heights of the left and right subtree of every node differ by at most 1. [Figure 21-13](#) shows two examples of AVL trees.




**Figure 21-13 Examples of AVL trees**

The tree of [Figure 21-14](#) fails to be an AVL tree because the right subtree of the root has height 2 whereas the left subtree has height 0.







**Figure 21-14 An example of a non-AVL tree**

## Adding New Elements to AVL Trees

AVL trees support the same operations as regular binary search trees: addition of new elements, removal, and search. To add a new element to an AVL tree, we simply add the new element using the same algorithm used to add elements to a regular binary search tree. It is possible, however, that the addition may upset the AVL balance condition, so we must check for this and rebalance the tree as necessary. We only need to check the balance condition at nodes that lie on the path between the newly inserted node and the root.

There are four types of imbalance that can result from adding a new node to an existing AVL tree:


- An *LL imbalance* occurs at a node  $N$  with a left child  $K$  when  $N$  and  $K$  are both left-heavy: that is, the height of the left subtree of  $N$  is greater than the height of the right subtree of  $N$ ; and the height of the left subtree of  $K$  is greater than the height of the right subtree of  $K$ . An LL imbalance is fixed by performing a *single right rotation* at  $N$ : this is illustrated in [Figure 21-15](#).



## Figure 21-15 A single right rotation is performed at a node with an LL imbalance

### [Figure 21-15 Full Alternative Text](#)

- An *RR imbalance* occurs at a node  $N$  with a right child  $M$  when  $N$  and  $M$  are both right-heavy. This situation is the mirror image of an LL imbalance: it is fixed by performing a *single left rotation* at  $N$ .
- An *LR imbalance* occurs at a node  $N$  with a left child  $K$  when  $N$  is left-heavy and  $K$  is right-heavy. An LR imbalance is fixed by performing a *double LR rotation* at  $N$ . This is illustrated in [Figure 21-16](#).



## Figure 21-16 A double LR rotation is performed at a node with an LR imbalance

### [Figure 21-16 Full Alternative Text](#)

- An *RL imbalance* occurs at a node  $N$  with a right child  $M$  when  $N$  is right-heavy and  $M$  is left-heavy. This situation is a mirror image of an LR imbalance, and can be fixed by performing a *double RL rotation* at  $M$ .

Consider now a general strategy for removing an LL imbalance that occurs after a new node has been added to an AVL tree. Let  $N$  be the node with the LL imbalance, let  $K$  be the left child of  $N$ , and let  $J$  be the left child of  $K$ . As shown in [Figure 21-17](#), let  $p$ ,  $q$ ,  $r$ , and  $s$  be respectively the left subtree of  $J$ , the right subtree of  $J$ , the right subtree of  $K$ , and the right subtree of  $N$ . Letting  $h$  be the height of  $s$ , we deduce that the height of the left subtree of  $N$  must be  $h+2$ . We can further deduce that the height of  $r$  must be  $h$ , and that one of  $p$  and  $q$  must have height  $h$  while the other has height  $h-1$ . Let us assume that  $p$  has height  $h$  and  $q$  has height  $h-1$ .




**Figure 21-17 A node with an LL imbalance**

In general, this situation occurs when a new node has been added to the subtree  $p$ , and the imbalance is discovered at  $N$  as we are tracing our way back to the root. At this point, we fix the imbalance by performing a single right rotation at  $N$ , that is, we rearrange the parent-child relationships among the nodes and subtrees shown to yield the tree shown in [Figure 21-18](#).




**Figure 21-18 Effect of a single right rotation**

Notice that the rearranged tree is still a binary search tree, and moreover, it now obeys the AVL balance condition at every node. An RR imbalance can similarly be fixed using a single left rotation.

Consider now how to remove an LR imbalance at a node  $N$ . Any such node  $N$  will have a left child  $J$ , which will in turn have a right child  $K$ . Let  $p$ ,  $q$ ,  $r$ , and  $s$  represent the left and right subtrees of these nodes as shown in [Figure 21-19](#). Assuming that  $s$  has height  $h$ , we deduce that the subtree rooted at  $J$  must have height  $h+2$ . We can further deduce that the height of  $p$  is  $h$ . Because the subtree rooted at  $K$  has height  $h+1$ , at least one of the subtrees of  $K$ , say  $q$ , will have height  $h$ . The other subtree,  $r$ , must have height  $h-1$  ( $r$  cannot have a height less than  $h-1$  because then the imbalance would be found at  $K$  rather than at  $N$ ).




**Figure 21-19 A node with an LR imbalance**

The LR imbalance shown in [Figure 21-19](#) is removed by executing a double LR rotation, which rearranges the tree into the form shown in [Figure 21-20](#).




**Figure 21-20 Effect of a double LR rotation on the tree of [Figure 21-19](#)**

The reconfigured tree is now an AVL tree. An RL imbalance can similarly be removed by performing a double RL rotation.

## Implementation of AVL Trees

An AVL node is a binary tree node with a height field. Our implementation of an AVL tree is based on the `AVLNode` class, which is a subclass of `Node`. The `AVLNode` class inherits the

```

int value;
Node left;
Node right;

```

fields from its superclass and adds an integer field to keep track of the height of the node in the AVL tree. The `Node` references `left` and `right` work fine when being assigned to, or when being used to access the fields inherited from `Node`, but require casting to access the height field:

```

int h = ((AVLNode)left).height;

```

To simplify the use of these references, we include in the `AVLNode` class convenience methods that apply the cast. You can see these methods in Lines 12 and 13 of [Code Listing 21-10](#).

## Code Listing 21-10 (`AVLNode.java`)

```

1 /**
2 AVLNode extends Node from which it inherits
3 the value field, and also the left and right links.

```

```

4 */
5
6 class AVLNode extends Node
7 {
8 int height;
9
10 // These methods cast the inherited NODE links
11 // AVLNode links.
12 AVLNode getLeft(){ return (AVLNode) left; }
13 AVLNode getRight(){ return (AVLNode) right; }
14
15 public AVLNode(int value)
16 {
17 // Call the other (sister) constructor.
18 this(value, null, null);
19 }
20
21 public AVLNode(int val, AVLNode left1, AVLNode right1)
22 {
23 // Pass the parameters to the superclass constructor.
24 super(val, left1, right1);
25 height = 0;
26 }
27
28 /**
29 * The resetHeight methods recomputes height if the
30 * left or right subtrees have changed.
31 */
32 void resetHeight()
33 {
34 int leftHeight = -1;
35 int rightHeight = -1;
36 if (left != null)
37 {
38 leftHeight = getLeft().height;
39 }
40 if (right != null)
41 {
42 rightHeight = getRight().height;
43 }
44 height = 1 + Math.max(leftHeight, rightHeight);
45 }
46 }

```

The implementation of AVL trees is similar to that of binary search trees, except that the add and remove methods must be modified to ensure that the tree stays balanced. [Code Listing 21-11](#) shows an AVL class `AVLTree` with the add method implemented. The class extends the `BinaryTree` class of [Code Listing 21-6](#), from which it inherits the `root` field. Although declared as being of type `Node`, the `root` field will point to `AVLNode` objects. Accordingly, the `AVLTree` class has a `getRoot()` method that retrieves `root` and automatically applies a cast to `AVLNode`. The `getRoot()` method appears in Lines 8–11 of [Code Listing 21-11](#).

The `add(bTree, x)` method is a straightforward implementation of the strategy already discussed. If `bTree` is empty, add creates and returns a new AVL tree with a single node. Otherwise, the method recursively adds the new value to one of the subtrees, recomputes the height of the subtrees, and

checks the balance condition at `bTree`. If the `bTree` is out of balance, add calls the `balance` method on it.

The `balance` method determines the type of imbalance existing at the node passed to it and calls one of four methods to perform the appropriate rotation. The methods that perform the rotations are named `llBalance`, `lrBalance`, `rlBalance`, and `rrBalance` after the type of imbalance they correct.

## Code Listing 21-11 (AVLTree.java)

```
1 /**
2 This class implements AVLTrees.
3 */
4 public class AVLTree extends BinaryTree
5 {
6 // Convenience method casts the inherited root
7 // from Node to AVLNode.
8 private AVLNode getRoot()
9 {
10 return (AVLNode) root;
11 }
12
13 /**
14 The getHeight method computes the height of an AVL tree.
15 @param tree An AVL tree.
16 @return The height of the AVL tree.
17 */
18 static int getHeight(AVLNode tree)
19 {
20 if (tree == null)
21 return -1;
22 else
23 return tree.height;
24 }
25
26 /**
27 The add method adds a value to this AVL tree.
28 @param x The value to add.
29 @return true.
30 */
31 public boolean add(int x)
32 {
33 root = add((AVLNode) root, x);
34 return true;
35 }
36
37 /**
38 The add method adds a value to an AVL tree.
39 @return The root of the augmented AVL tree.
40 */
41 private AVLNode add(AVLNode bTree, int x)
42 {
43 if (bTree == null)
44 {
```



```

45 return new AVLNode(x);
46 }
47 if (x < bTree.value)
48 {
49 bTree.left = add(bTree.getLeft(), x);
50 } else
51 {
52 bTree.right = add(bTree.getRight(), x);
53 }
54
55 // Compute heights of the left and right subtrees
56 // and rebalance the tree if needed.
57 int leftHeight = getHeight(bTree.getLeft());
58 int rightHeight = getHeight(bTree.getRight());
59 if (Math.abs(leftHeight - rightHeight) == 2)
60 {
61 return balance(bTree);
62 } else
63 {
64 bTree.resetHeight();
65 return bTree;
66 }
67 }
68
69 /**
70 The balance method rebalances an AVL tree.
71 @param bTree The AVL tree needing to be balanced.
72 @return The balanced AVL tree.
73 */
74 private AVLNode balance(AVLNode bTree)
75 {
76 int rHeight = getHeight(bTree.getRight());
77 int lHeight = getHeight(bTree.getLeft());
78
79 if (rHeight > lHeight)
80 {
81 AVLNode rightChild = bTree.getRight();
82 int rrHeight = getHeight(rightChild.getRight());
83 int rlHeight = getHeight(rightChild.getLeft());
84 if (rrHeight > rlHeight)
85 {
86 return rrBalance(bTree);
87 } else
88 {
89 return rlBalance(bTree);
90 }
91 } else
92 {
93 AVLNode leftChild = bTree.getLeft();
94 int llHeight = getHeight(leftChild.getLeft());
95 int lrHeight = getHeight(leftChild.getRight());
96 if (llHeight > lrHeight)
97 {
98 return llBalance(bTree);
99 } else
100 {
101 return lrBalance(bTree);
102 }

```

```

103 }
104 }
105
106 /**
107 The rrBalance method corrects an RR imbalance.
108 @param bTree The AVL tree with an RR imbalance.
109 @return The balanced AVL tree.
110 */
111 private AVLNode rrBalance(AVLNode bTree)
112 {
113 AVLNode rightChild = bTree.getRight();
114 AVLNode rightLeftChild = rightChild.getLeft();
115 rightChild.left = bTree;
116 bTree.right = rightLeftChild;
117 bTree.resetHeight();
118 rightChild.resetHeight();
119 return rightChild;
120 }
121
122 /**
123 The rlBalance method corrects an RL imbalance.
124 @param bTree The AVL tree with an RL imbalance.
125 @return The balanced AVL tree.
126 */
127 private AVLNode rlBalance(AVLNode bTree)
128 {
129 AVLNode root = bTree;
130 AVLNode rNode = root.getRight();
131 AVLNode rlNode = rNode.getLeft();
132 AVLNode rlrTree = rlNode.getRight();
133 AVLNode rllTree = rlNode.getLeft();
134
135 // Build the restructured tree.
136 rNode.left = rlrTree;
137 root.right = rllTree;
138 rlNode.left = root;
139 rlNode.right = rNode;
140
141 // Adjust heights
142 rNode.resetHeight();
143 root.resetHeight();
144 rlNode.resetHeight();
145
146 return rlNode;
147 }
148
149 /**
150 The llBalance method corrects an LL imbalance.
151 @param bTree The AVL tree with an LL imbalance.
152 @return The balanced AVL tree.
153 */
154 private AVLNode llBalance(AVLNode bTree)
155 {
156 AVLNode leftChild = bTree.getLeft();
157 AVLNode lrTree = leftChild.getRight();
158 leftChild.right = bTree;
159 bTree.left = lrTree;

```

```

160 bTree.resetHeight();
161 leftChild.resetHeight();
162 return leftChild;
163 }
164
165 /**
166 * The lrBalance method corrects an LR imbalance.
167 * @param bTree The AVL tree with an LR imbalance.
168 * @return The balanced AVL tree.
169 */
170 private AVLNode lrBalance(AVLNode bTree)
171 {
172 AVLNode root = bTree;
173 AVLNode lNode = root.getLeft();
174 AVLNode lrNode = lNode.getRight();
175 AVLNode lrlTree = lrNode.getLeft();
176 AVLNode lrrTree = lrNode.getRight();
177
178 // Build the restructured tree.
179 lNode.right = lrlTree;
180 root.left = lrrTree;
181 lrNode.left = lNode;
182 lrNode.right = root;
183
184 // Adjust heights.
185 lNode.resetHeight();
186 root.resetHeight();
187 lrNode.resetHeight();
188
189 return lrNode;
190 }
191 }
192

```

[Code Listing 21-12](#) implements the graphical interface to the `AVLTree` class. The program is a simple modification of the one in [Code Listing 21-9](#).

## Code Listing 21-12 (GUIAVLTreeDemo.java)

```

1 import java.util.Scanner;
2 import javafx.application.Application;
3 import javafx.event.ActionEvent;
4 import javafx.event.EventHandler;
5 import javafx.geometry.Insets;
6 import javafx.geometry.Pos;
7 import javafx.scene.Scene;
8 import javafx.scene.control.Label;
9 import javafx.scene.control.TextField;
10 import javafx.scene.layout.HBox;
11 import javafx.scene.layout.VBox;
12 import javafx.stage.Stage;
13
14 /**
15 This class is a JavaFX interface for testing the

```

```

16 the binary search tree class.
17 */
18
19 public class GUIAVLTreeDemo extends Application
20 {
21 @Override
22 public void start(Stage stage) throws Exception
23 {
24 // Set up the binary tree to manipulate
25
26 AVLTree avlTree = new AVLTree();
27
28 // This binary tree wraps the real binary tree
29 // so it can be displayed.
30 DisplayableBtree
31 displayBinTree = new DisplayableBtree(avlTree);
32
33 // Displays the result of the user command.
34 TextField cmdResultTextField = new TextField();
35 cmdResultTextField.setEditable(false);
36
37 // Used by the user to enter a command.
38 TextField cmdTextField = new TextField();
39
40 // HBox1 holds cmdResultTextField and its label.
41 HBox hBox1 = new HBox(10);
42 Label cmdResultLabel = new Label("Command Result");
43 hBox1.getChildren().addAll(cmdResultLabel, cmdResultTextField);
44
45 // Place holder for the binary tree view.
46 HBox bTreeView = new HBox();
47 bTreeView.setAlignment(Pos.CENTER);
48
49 // HBox2 holds the cmdTextField and its label
50 HBox hBox2 = new HBox(10);
51 Label cmdLabel = new Label("Command");
52 hBox2.getChildren().addAll(cmdLabel, cmdTextField);
53
54 // Put the three HBox objects in a VBox
55 VBox outerPane = new VBox(10);
56 outerPane.setPadding(new Insets(10));
57 outerPane.getChildren().addAll(hBox1, bTreeView, hBox2);
58
59 // Command handler needs to access to the bTreeView
60 // so it can set the binary tree view, and to the
61 // cmdResultTextField so it can set the command result.
62 // Also needs access to the binary tree and its displayable
63 // twin.
64 EventHandler<ActionEvent>
65 handler = new CmdHandler(bTreeView, cmdResultTextField,
66 avlTree, displayBinTree);
67 cmdTextField.setOnAction(handler);
68
69 // Set up the stage
70 stage.setTitle("Binary Search Tree Operations");
71 stage.setScene(new Scene(outerPane));
72 stage.show();
73 }

```

```

74
75 public static void main(String [] args)
76 {
77 launch(args) ;
78 }
79
80 }
81
82 // Event Handler class
83 class CmdHandler implements EventHandler<ActionEvent>
84 {
85 // These class variables reference
86 // data in the calling environment.
87 private HBox view;
88 private TextField cmdResultTextField;
89 private AVLTree binTree;
90 private DisplayableBtree displayBinTree;
91
92 // Constructor stores references to data in
93 // the calling environment.
94 public CmdHandler(HBox tView, TextField cmdResTf,
95 AVLTree tree,
96 DisplayableBtree dTree)
97 {
98 view = tView;
99 cmdResultTextField = cmdResTf;
100 binTree = tree;
101 displayBinTree = dTree;
102 }
103
104 public void handle(ActionEvent evt)
105 {
106 // Get the command from the event target.
107 TextField cmdTf = (TextField)evt.getTarget();
108 String cmdStr = cmdTf.getText();
109
110 // Use a scanner to read the command and its parts.
111 Scanner sc = new Scanner(cmdStr);
112 String cmd = sc.next();
113 switch (cmd)
114 {
115 case "add":
116 {
117 // Perform the operation on the tree.
118 int value = sc.nextInt();
119 binTree.add(value);
120
121 // Replace the old view of the tree.
122 view.getChildren().clear();
123 view.getChildren().add(displayBinTree.getView());
124 cmdResultTextField.setText(" ");
125 }
126 return;
127 }
128 }
129 }

```

A screenshot of this user interface after the user has issued commands to add the first six multiples of 10 to an initially empty AVL tree is shown in [Figure 21-21](#).




**Figure 21-21 Sample interaction with AVL program**

(Mark Schrier. Copyright Oracle, Inc.)

# 21.4 Priority Queues

**CONCEPT:** A priority queue is a collection that stores elements according to some order. Priority queues support addition and removal of elements, and a removal always returns the least element in the collection.

An *order* on a set of values is a way of establishing, for any two elements  $x$  and  $y$ , whether  $x$  is “less than”  $y$  or  $y$  is “less than”  $x$ . Many sets have a *natural order*: for example, the set of integers is ordered by the usual “less than” relation  $<$ , that is, in the natural order for the integers,  $x$  is “less than”  $y$  if and only if  $x < y$ . Likewise, in the natural order for strings, a string  $s$  is “less than” string  $t$  if  $s$  comes before  $t$  in alphabetical order. In an ordered set, an element that is “less than” every other element is called the *least* element, or the *minimum*. All orders referred to in this chapter will be assumed to be *total* (or *linear*), meaning that for any two elements  $x$  and  $y$ , exactly one of the following conditions must hold:

1.  $x$  is “less than”  $y$
2.  $x$  is equal to  $y$
3.  $y$  is “less than”  $x$

As an example of an order that is not natural, we can stipulate, for any two numbers  $x$  and  $y$ , that  $x$  is “less than”  $y$  if  $x > y$ . In this case, for a set such as  $\{1, 2, 3, 4\}$ , the defined order turns the set into the sequence  $\{4, 3, 2, 1\}$ , and makes 4 the “least” element of the set.

The natural order of primitive types in Java is specified by the built-in comparison operator  $<$ . A class type such as `Integer` or `String` specifies its natural order by implementing the `Comparable` interface.

A *priority queue* is a collection that stores elements and provides access to those elements according to some order. A remove operation executed on a priority queue always returns a minimum element. The queue element that will be returned by the next call to remove is said to be at the *head of the queue*.

The main operations supported by a priority queue when regarded as an abstract data structure are as follows:

- *add(x)*: Adds  $x$  to the priority queue.
- *removeMin()*: Removes and returns the element at the head of the queue; this will be an element that is less than, or equal to, any other element in the queue.

The Java Collection Framework provides an implementation of a priority queue in a class called `PriorityQueue`. It provides more methods than the two we have listed, and some of the method

names are different. Most of its methods are listed in [Table 21-1](#).

## Table 21-1 PriorityQueue class methods

|                                             |                                                                                                                                                                                                     |
|---------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>PriorityQueue&lt;E&gt;<br/>( )</code> | Creates a queue that stores items of class type <code>E</code> and orders them according to their natural ordering (using the <code>Comparable</code> interface as implemented by <code>E</code> ). |
|---------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|                                                                                     |  |
|-------------------------------------------------------------------------------------|--|
| <code>PriorityQueue&lt;E&gt;(int capacity, Comparator&lt;? Super E&gt; comp)</code> |  |
|-------------------------------------------------------------------------------------|--|

|  |                                                                                                        |
|--|--------------------------------------------------------------------------------------------------------|
|  | Creates a queue that stores items of class <code>E</code> and orders them by the specified comparator. |
|--|--------------------------------------------------------------------------------------------------------|

|                                      |                                                                              |
|--------------------------------------|------------------------------------------------------------------------------|
| <code>boolean add(E<br/>item)</code> | Adds <code>item</code> to the priority queue and returns <code>true</code> . |
|--------------------------------------|------------------------------------------------------------------------------|

|                       |                                                                                                                             |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------|
| <code>E poll()</code> | Removes and returns a minimum element from the priority queue. This method returns <code>null</code> if the queue is empty. |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------|

|                       |                                                                                                                                                                          |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>E peek()</code> | Returns the item that is currently at the head of the queue, but does not remove it from the queue. If the queue is empty, <code>peek</code> returns <code>null</code> . |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

|                         |                                                                      |
|-------------------------|----------------------------------------------------------------------|
| <code>int size()</code> | Returns the number of items currently stored in this priority queue. |
|-------------------------|----------------------------------------------------------------------|

## Applications of Priority Queues

Priority queues are useful whenever items have to be processed according to some priority. For example, a personal information manager application might keep a list of tasks to be completed in a priority queue ordered by deadline, with the minimum task being the one with the closest deadline. Priority queues are also used by computer operating systems to schedule tasks for execution, and in computer networks to schedule data packets for transmission.

If you store a set of items in a priority queue and then retrieve them, they will come out in sorted order. This is the basis of a very efficient sorting algorithm named *Heapsort*. [Code Listing 21-13](#) illustrates the use of a priority queue to sort an array.



# Code Listing 21-13 (Heapsort.java)

```
1 import java.util.*;
2
3 /**
4 * This program demonstrates the Heapsort algorithm.
5 */
6
7 public class Heapsort
8 {
9 public static void main(String [] args)
10 {
11 // Create and display an array of random integers
12 Random randy = new Random();
13 int [] arr = new int[10];
14 System.out.println("Here is the array to be sorted:");
15 for (int k = 0; k < arr.length; k++)
16 {
17 arr[k] = randy.nextInt(100);
18 System.out.print(arr[k] + " ");
19 }
20
21 // Create a priority queue of integers
22 // and use it to sort the array
23 PriorityQueue<Integer> pQueue =
24 new PriorityQueue<Integer>();
25 for (int x : arr) pQueue.add(x);
26 for (int k = 0; k < arr.length; k++)
27 arr[k] = pQueue.poll();
28
29 // Print the array
30 System.out.println("\nHere is the sorted array:");
31 for (int x : arr) System.out.print(x + " ");
32 }
33 }
```

## Program Output

```
Here is the array to be sorted:
13 81 78 93 94 60 2 32 88 58
Here is the sorted array:
2 13 32 58 60 78 81 88 93 94
```

## Using Comparators with Priority Queues

The no-argument constructor of the `PriorityQueue` class used in [Code Listing 21-13](#) orders items according to their natural ordering. This is the order specified by the `<` operator if the type of items is primitive, or specified by the `Comparable` interface for class types that implement it. If we need to use a different order, we must use the second constructor shown in [Table 21-1](#). That constructor takes

two parameters: an integer specifying the capacity of the queue to be created, and a `Comparator` object. `Comparator` is a generic interface defined as follows:

```
interface Comparator<T>
{
 int compare (T x, T y);
 boolean equals(Object o);
}
```

A class implementing this interface provides a definition of `compare` that returns a negative integer if  $x$  should be regarded as less than  $y$ , a positive integer if  $x$  should be regarded as greater than  $y$ , and 0 if  $x$  is equal to  $y$ . The `equals` method of the `Comparator` interface exists to allow two `Comparator` objects to be checked for equality, and does not have to be implemented if the default implementation inherited from `Object` is adequate.

As an example, let us look at a version of Heapsort that sorts integers in alphabetical order of their string representation. This is done by creating a `PriorityQueue` object that uses a `Comparator` object defined from the following class:

```
class AlphaOrder implements Comparator<Integer>
{
 public int compare(Integer x, Integer y)
 {
 return x.toString().compareTo(y.toString());
 }
}
```

The `AlphaOrder` class is defined inside `Alphaheapsort` as a static inner class. `AlphaOrder` is made static because it will be instantiated by the `main` method, which is itself static. In Java, static methods cannot instantiate non-static inner classes.

The following code fragment creates a priority queue that uses `AlphaOrder` as a comparator, and the program in [Code Listing 21-14](#) demonstrates its use:

```
AlphaOrder comp = new AlphaOrder();
PriorityQueue<Integer>
 pQueue = new PriorityQueue<Integer>(arr.length, comp);
```

## Code Listing 21-14 (Alphaheapsort.java)

```
1 import java.util.*;
2
3 /**
4 * This program shows the use of a PriorityQueue with
5 * a custom Comparator object.
6 */
7
8 public class Alphaheapsort
9 {
10 /**
```

```

11 Specialized Comparator.
12 */
13
14 static class AlphaOrder implements Comparator<Integer>
15 {
16 public int compare(Integer x, Integer y)
17 {
18 return x.toString().compareTo(y.toString());
19 }
20 }
21
22 public static void main(String [] args)
23 {
24 // Create and display an array of random integers
25 Random randy = new Random();
26 int [] arr = new int[10];
27 String str = "Here is the original list of numbers:";
28 System.out.println(str);
29 for (int k = 0; k < arr.length; k++)
30 {
31 arr[k] = randy.nextInt(10000);
32 System.out.print(arr[k] + " ");
33 }
34
35 // Create a priority queue of integers
36 // and use it to sort the array
37 AlphaOrder c = new AlphaOrder();
38 PriorityQueue<Integer>
39 pQueue = new PriorityQueue<Integer>(arr.length, c);
40 for (int x : arr)
41 pQueue.add(x);
42 for (int k = 0; k < arr.length; k++)
43 arr[k] = pQueue.poll();
44
45 // Print the array
46 System.out.println("\nHere are the numbers sorted " +
47 "in alphabetical order:");
48 for (int x : arr)
49 System.out.print(x + " ");
50 }
51 }

```

## Program Output

```

Here is the original list of numbers:
1065 8433 743 4962 9954 31 1815 4135 489 3094
Here are the numbers sorted alphabetical in order:
1065 1815 3094 31 4135 489 4962 743 8433 9954

```

## Analysis of Heapsort

We can analyze the efficiency of Heapsort by counting the number of comparisons of array elements

that it causes the priority queue to perform. It turns out that a properly implemented priority queue holding  $n$  elements requires at most  $\log n$  comparisons to add or remove an item. Given an array of  $n$  elements, we can add all of them to a priority queue using  $n \log n$  comparisons, and then remove all of them using another  $n \log n$  comparisons, for a total of  $2n \log n$  comparisons. In the language of [Chapter 16](#), Heapsort is a  $O(n \log n)$  algorithm. In practice, an  $n \log n$  algorithm is a lot faster than a  $O(n^2)$  algorithm such as Selectionsort or Bubblesort. It can be shown that no algorithm that sorts by comparing items can sort an array of  $n$  items in less than  $n \log n$  comparisons. This means that, within a constant factor, the performance of Heapsort is the best possible.

## Implementation of Priority Queues

Our goal in implementing a priority queue is to find a data structure that allows each of the two main priority queue operations, *add* and *removeMin*, to be performed with no more than  $\log n$  comparisons when the priority queue has  $n$  items. Two possibilities immediately come to mind:

1. We can use an unsorted linked list to hold the  $n$  items. This will allow us to add a new element to the list quickly, but will force us to perform  $n-1$  comparisons to find and remove the minimum element.
2. We can use a linked list sorted in increasing order to hold the  $n$  items. This allows us to quickly find and remove the minimum, but will force us, in the worst case, to perform  $n$  comparisons while looking for the right place to insert the new item.

Neither of these achieves the desired  $O(\log n)$  complexity for both operations, so we must look for another approach.

## Binary Trees with the Heap Order Property

A third possibility is to use a binary tree to hold the elements of the queue. Consider a binary tree  $T$  that satisfies the following condition, which we will dub *the heap order property*:

At every node  $N$  of  $T$ , the value stored in  $N$  is greater than the value stored in the parent of  $N$ .

Intuitively, a binary tree with this property is a tree in which every path from the root to a leaf is sorted in increasing order. [Figure 21-22](#) shows an example of such a tree.




**Figure 21-22 A binary tree with the heap order property**

We need to impose a condition on such a tree that will keep it shallow; that is, we want the length of the longest path through the tree to be approximately  $\log n$  rather than  $n$ . This will ensure that an addition operation that has to traverse a path from the root to a leaf can be performed in  $\log n$  time. We need to introduce some terminology before we can proceed.

Let  $T$  be a binary tree. The *level* of a node  $N$  in  $T$  is the length of a path from the root to  $N$ . For example, the level of the root is 0, and the level of a child of a root is 1. The *depth* of  $T$  is the maximum level of a node in  $T$ ; that is, it is the length of the longest path from the root to a leaf. A binary tree with depth  $d$  has  $d+1$  levels, because it has a level for each value of  $l$  in the range  $0 \leq l \leq d$ . We are looking for a family of trees whose depth is at most  $\log n$  when the tree has  $n$  nodes.

## Complete Binary Trees and Heaps

A binary tree with depth  $d$  is *complete* if it satisfies two conditions:

- $T$  has  $2^l$  nodes for all levels  $0 \leq l \leq d-1$ ; that is, each level other than the last has the maximum number of nodes possible
- All leaf nodes at level  $d$  are as far to the left as possible

Basically, a complete binary tree is one that you get when you build a binary tree by starting with the root and adding nodes level by level, following these rules:

- R1. No node gets a child until each node at the previous level has two children.
- R2. No node gets a child unless its left sibling already has two children.

- R3. No node gets a right child unless it already has a left child.

The first rule ensures that there are no “holes” in the first  $d$  levels of the tree. The next two rules ensure that there are no “holes” as you move from left to right through the last level.

The tree of [Figure 21-22](#) is not complete because the node that contains 36 has a right child but no left child. [Figure 21-23](#) shows an example of a complete binary tree.




**Figure 21-23 A complete binary tree that is also a heap**

A complete binary tree that stores values from an ordered set is called a *heap* if it satisfies the heap order property.

## The Depth of Complete Binary Trees


A complete binary tree  $T$  with  $n$  nodes is shallow in the sense that it cannot have depth much greater than  $\log n$ . If we disregard the last level, what remains of  $T$  is perfectly balanced in the sense that at each node, there are just as many nodes in the left subtree as in the right subtree. This means that if we start at the root and descend along any path, we retain less than half of the remaining nodes each time we pass to a new (numerically higher) level. This shows that the length of any path can be at most  $\log n + 1$ . We state this as follows:

A complete binary tree with  $n$  nodes has depth  $d \leq \log n + 1$ .

## Storing a Complete Binary Tree in an Array

The fact that a complete binary tree has no “holes” has an interesting consequence. If you build the binary tree using rules R1 to R3 and store each node in an array  $A$  as it is added to the tree, you get

the pattern shown in [Figure 21-24](#).



## Figure 21-24 Storing a complete binary tree in an array

[Figure 21-24 Full Alternative Text](#)

This pattern has the following characteristics:


- The root of the tree is  $A[0]$
- The parent of a node  $A[k]$  that is not the root is  $A[(k-1)/2]$
- The left child and right children of a node  $A[k]$  are  $A[2k+1]$  and  $A[2k+2]$ , respectively
- The rightmost leaf in the last level (last node in the tree) is at  $A[n-1]$
- A node  $A[k]$  is a leaf if  $2k+1 \geq n$ , it has a left child if  $2k+1 < n$ , and has a right child if  $2k+2 < n$

These observations allow us to use either an array or an array list to store a complete binary tree. In particular, to build a heap of integers, we can use an array list such as the following:

```
ArrayList<Integer> arrayHeap;
```

## Adding an Item to a Heap

An addition of a new item to a heap must maintain the structure of a complete binary tree as well as the heap order property. We can preserve the structure of the tree by adding the new node according to rules R1, R2, and R3. This means that the new node will be added as a leaf at the last level of the tree, and will be as far to the left as possible within that level. For example, adding 12 to the tree of [Figure 21-23](#) yields the tree of [Figure 21-25](#).



## Figure 21-25 Adding a new value may violate the heap order property

In terms of our array list representation, this means that the new value is added to the end of the array list as follows:

```
arrayHeap.add(x);
```

Once we have added the new node, say  $X$ , we restore the heap order property by traversing the path from the point of insertion back to the root. As we go, we compare the value of  $X$  to that of its parent and swap the two nodes if the value in the parent is greater. Then we climb up the tree to the new position of  $X$  and repeat. We stop when the  $X$  becomes the root, or when it gets to a place where its value is greater than that of its parent. When we do this, we say we are *sifting up* the new node. The code for adding a value to our heap is embodied in the following method:

```
boolean add(int x)
{
 // Add x at the end of the array list
 arrayHeap.add(x);
 // Sift up
 siftUp();
 return true;
}
```

The sifting up procedure is performed by the following method:

```
private void siftUp()
{
 int p = arrayHeap.size()-1; // Position to sift up
 while (p != 0)
 {
 int parent = (p-1) / 2; // Index of parent
 if (valueAt(p) >= valueAt(parent))
 {
 // We are done
 return;
 }
 }
}
```



```

 }
 else
 {
 // Do a swap
 Integer temp = arrayHeap.get(parent);
 arrayHeap.set(parent, arrayHeap.get(p));
 arrayHeap.set(p, temp);
 // Move up
 p = parent;
 }
}
}

```

## Removing the Minimum Element

To remove an element from a heap, we delete the root node and replace it with the node in the last place of the queue, that is, the rightmost leaf at the deepest level of the tree. If we do this to the heap of [Figure 21-23](#), we get the complete binary tree shown in [Figure 21-26](#).




**Figure 21-26 The heap of [Figure 21-23](#) after replacing the root with the deepest rightmost leaf**

In general, the tree will violate the heap order property when the root is replaced in this manner, and we must perform a *sift down* operation to remake the heap. Say that  $X$  is the value that is currently stored in the root and needs to be sifted down. We start at the root and make our way down some branch of the tree. If we get to a place where  $X$  is smaller than its children, or is a leaf, we stop the procedure. Otherwise, we swap  $X$  with its smallest child, and then follow  $X$  down to the next level. In this way, we sift  $X$  down the tree until we get it to where it belongs. At that point, the heap order property is restored and we have remade the tree into a heap.

In terms of our array implementation, we must first save the value at position 0, and then remove the value at the end of the array list and store it at position 0. The `removeMin` method is as follows:

```
public int removeMin()
{
 if (isEmpty())
 throw new RuntimeException("Priority Queue is empty.");
 else
 {
 int val = arrayHeap.get(0);

 // Replace root by last leaf
 arrayHeap.set(0, arrayHeap.get(arrayHeap.size()-1));

 // Remove the last leaf
 arrayHeap.remove(arrayHeap.size()-1);
 siftDown();
 return val;
 }
}
```

You can see the `siftDown` method, together with the methods for `add` and `removeMin`, in [Code Listing 21-15](#).

## Code Listing 21-15 (Arrayheap.java)

```
1 import java.util.*;
2
3 /**
4 * This class is an array-based heap
5 */
6
7 public class ArrayHeap
8 {
9 private ArrayList<Integer> arrayHeap
10 = new ArrayList<>();
11
12 /**
13 * The siftUp method sifts up the element
14 * at arrayHeap[size()-1].
15 */
16 private void siftUp()
17 {
18 int p = arrayHeap.size()-1; // Position to sift up
19 while (p != 0)
20 {
21 int parent = (p-1) / 2;
22 if (valueAt(p) >= valueAt(parent))
23 {
24 // We are done
25 break;
26 }
27 else
28 {
```

```

29 // Do a swap
30 Integer temp = arrayHeap.get(parent);
31 arrayHeap.set(parent, arrayHeap.get(p));
32 arrayHeap.set(p, temp);
33
34 // Move up
35 p = parent;
36 }
37 }
38 }
39
40 /**
41 The siftDown method sifts down the
42 element at arrayHeap[0].
43 */
44
45 private void siftDown()
46 {
47 int p = 0; // Position to sift down
48 int size = arrayHeap.size();
49 while (2*p + 1 < size)
50 {
51 int leftChildPos = 2*p + 1;
52 int rightChildPos = leftChildPos + 1;
53 int minChildPos = leftChildPos;
54
55 // Is there a right child?
56 if (rightChildPos < size)
57 {
58 // Which child is smaller
59 if (valueAt(rightChildPos) < valueAt(leftChildPos))
60 {
61 minChildPos = rightChildPos;
62 }
63 }
64 // If less than children we are done,
65 // otherwise swap node with smaller child
66 if (valueAt(p) <= valueAt(minChildPos))
67 break;
68 else
69 {
70 // Do the swap
71 Integer temp = arrayHeap.get(p);
72 arrayHeap.set(p, arrayHeap.get(minChildPos));
73 arrayHeap.set(minChildPos, temp);
74 }
75 // Go down to the child position
76 p = minChildPos;
77 }
78 }
79
80 /**
81 The add method adds a value to the heap.
82 @param x The value to add.
83 @return true.
84 */
85
86 boolean add(int x)

```

```

87 {
88 // Add x at the end of the array list
89 arrayHeap.add(x);
90
91 // Sift up
92 siftUp();
93 return true;
94 }
95
96 /**
97 * The removeMin method removes an item from the heap.
98 * @return The minimum element in the heap.
99 * @exception RuntimeException When priority
100 queue is empty.
101 */
102
103 public int removeMin()
104 {
105 if (isEmpty())
106 {
107 String message = "Priority Queue is empty.";
108 throw new RuntimeException(message);
109 }
110 else
111 {
112 int val = arrayHeap.get(0);
113
114 // Replace root by last leaf
115 int lastPos = arrayHeap.size()-1;
116 arrayHeap.set(0, arrayHeap.get(lastPos));
117
118 // Remove the last leaf
119 arrayHeap.remove(arrayHeap.size()-1);
120 siftDown();
121 return val;
122 }
123 }
124
125 /**
126 * The isEmpty method checks if the heap is empty.
127 * @return true if the heap is empty
128 * and false otherwise.
129 */
130
131 public boolean isEmpty()
132 {
133 return arrayHeap.size() == 0;
134 }
135
136 /**
137 * The valueAt method is a convenience method,
138 * it makes the code more readable.
139 * @return value stored at given position in the heap.
140 */
141
142 private int valueAt(int pos)
143 {

```

```

144 return arrayHeap.get(pos);
145 }
146 }

```

The `ArrayHeapsort` program, given in [Code Listing 21-16](#) demonstrates the use of the `ArrayHeap` class to sort an array. No output is shown for this program because it is similar to that of other programs in this chapter.

## Code Listing 21-16 (`ArrayHeapsort.java`)

```

1 import java.util.*;
2
3 /**
4 * This program demonstrates the use of the
5 * ArrayHeap class.
6 */
7
8 public class ArrayHeapsort
9 {
10 public static void main(String [] args)
11 {
12 // Create and display an array of random integers
13 Random randy = new Random();
14 int [] arr = new int[10];
15 System.out.println("Here is the array to be sorted:");
16 for (int k = 0; k < arr.length; k++)
17 {
18 arr[k] = randy.nextInt(100);
19 System.out.print(arr[k] + " ");
20 }
21
22 // Create a heap of integers
23 // and use it to sort the array
24 ArrayHeap heap = new ArrayHeap();
25 for (int x : arr) heap.add(x);
26 for (int k = 0; k < arr.length; k++)
27 arr[k] = heap.removeMin();
28
29 // Print the array
30 System.out.println("\nHere is the sorted array:");
31 for (int x : arr)
32 System.out.print(x + " ");
33 }
34 }

```

**In the Spotlight:**



A rooted tree is a generalization of a binary tree in which nodes may have more than two children. A familiar example of a rooted tree is the hierarchical file and directory system used on almost all computers. [Figure 21-27](#) shows a simple example of a rooted tree that might be a part of a genealogy tree.




**Figure 21-27 A rooted tree**

Many programs need to display rooted trees as part of their user interfaces. JavaFX provides a generic class `TreeItem<T>` to represent a node of a tree. The type parameter `T` represents the type of data stored in the tree item. The following code creates a tree item that will serve as the root of the tree in [Figure 21-27](#):

```
TreeItem<String> root = new TreeItem<>("Al");
```

A tree item can have other tree items as children. The following lines of code create three items and make them children of the one already created:

```
TreeItem<String> bob = new TreeItem<>("Bob");
TreeItem<String> kate = new TreeItem<>("Kate");
TreeItem<String> sidney = new TreeItem<>("Sidney");
root.getChildren().addAll(bob, kate, sidney);
```

To form an in-memory representation of a rooted tree, first create a tree item to serve as the root of the tree, and then add new items as children of those already in the tree. To display the tree as part of the JavaFX scene graph, you must wrap its root item in a `TreeView` component like this:

```
TreeView<String> treeView = new TreeView<>(root);
```

The tree view component can then be added to a `Pane` component that is already part of the scene graph. Alternatively, the tree view can itself be the root of the scene graph. [Code Listing 21-17](#) illustrates the process of building and displaying the tree of [Figure 21-28](#).




**Figure 21-28 Screen shot of the TreeView example program**

(Mark Schrier. Copyright Oracle, Inc.)

[Figure 21-28 Full Alternative Text](#)

## Code Listing 21-17 (TreeViewProg.java)

```
1 import javafx.application.Application;
2 import javafx.scene.Scene;
3 import javafx.scene.control.TreeItem;
4 import javafx.scene.control.TreeView;
5 import javafx.stage.Stage;
6
7 /**
8 * This program demonstrates use of
9 * the TreeView component
10 */
11
12 public class TreeViewProg extends Application
13 {
14 @Override
15 public void start(Stage primaryStage)
16 {
17 // Create root node
18 TreeItem<String> root = new TreeItem<>("Al");
19 root.setExpanded(true);
20
21 // Create three nodes to be children of the root
22 TreeItem<String> bob = new TreeItem<>("Bob");
```

```

23 TreeItem<String> kate = new TreeItem<>("Kate");
24 TreeItem<String> sidney = new TreeItem<>("Sidney");
25 root.getChildren().addAll(bob, kate, sidney);
26
27 // Create Bob's children
28 TreeItem<String> amy = new TreeItem<>("Amy");
29 TreeItem<String> chuck = new TreeItem<>("Chuck");
30 bob.getChildren().addAll(amy, chuck);
31
32 // Create Sidney's child
33 TreeItem<String> joey = new TreeItem<>("Joey");
34 sidney.getChildren().add(joey);
35
36 TreeView<String> treeView = new TreeView<>(root);
37
38 Scene scene = new Scene(treeView, 300, 250);
39 primaryStage.setTitle("Tree Exhibition Program");
40 primaryStage.setScene(scene);
41 primaryStage.show();
42 }
43
44 /**
45 @param args the command-line arguments
46 */
47 public static void main(String[] args)
48 {
49 launch(args);
50 }
51 }

```

The user interface displayed when this program is run is shown in [Figure 21-28](#).



# 21.5 Common Errors to Avoid

- Forgetting to check for an empty tree. Make sure a reference to a node is not `null` before you try to access the value or children of the node.
- Forgetting to replace a subtree with a modified subtree. If you have to alter a tree, you will often have to recursively alter one or both subtrees. Make sure that you replace the current subtree with the altered version.
- Trying to use a loop when recursion is the best way to go. Most problems on binary trees have simple and elegant recursive solutions. Very often you don't need to use both iteration and recursion on the same problem.

# Review Questions and Exercises

## Multiple Choice and True/False

1. A binary tree is a collection of items in which each item .
  1. has no successor
  2. has one successor
  3. has exactly two successors
  4. has at most two successors
2. The three standard techniques for traversing a binary tree are .
  1. preorder, inorder, and postorder
  2. prefix, infix, and postfix
  3. left-to-right, top-down, and parent-to-child
  4. none of the above
3. The set of all descendants of a node in a binary tree form .
  1. a path from the node to the end of the tree
  2. a path from the node to the root
  3. a subtree rooted at that node
  4. a set that is best traversed using inorder traversal
4. A binary search tree is structured in such a way that .
  1. for each node, the left subtree holds values that are less than the node, and the right subtree holds values that are greater
  2. for each node, the left child holds a value that is less than the node, and the right child holds a value that is greater
  3. the value stored at each node is less than values stored in the children

4. the heights of the subtrees at any node differ by at most 1
5. An AVL tree is .
  1. a kind of priority queue
  2. a kind of heap
  3. a kind of binary search tree
  4. none of the above
6. When a node with two children is removed from a binary search tree, it can be replaced with .
  1. the rightmost leaf at the deepest level of the tree
  2. the smallest node in its right subtree
  3. the smallest node in its left subtree
  4. the smallest ancestor of the node
7. In a binary tree that is a heap, .
  1. every non-root node must have a value greater than that of its parent
  2. every left child must be less than its parent, and every right child must be greater
  3. the height of the left and right subtrees of each node must differ by at most 1
  4. leaf nodes must be greater or equal to every other node
8. In a complete binary tree, .
  1. every level  $l$  other than the last must have  $2^l$  nodes
  2. the depth  $d$  of the tree satisfies  $d \leq \log n + 1$ , where  $n$  is the number of nodes in the tree
  3. every node that has a right child also has a left child
  4. all of the above
9. The Heapsort method .
  1. is based on a complete binary search tree structure with the heap order property
  2. makes  $O(n \log n)$  comparisons when sorting an array of  $n$  entries

3. is a very good sorting method, but there are other general sorting methods that are much faster
  4. all of the above
10. To add a new value to a heap, you should .
1. add it as a leaf and then sift up
  2. sift the root down to make room for the new node
  3. replace the root with the deepest and rightmost leaf, and then sift down
  4. none of the above
11. **True or False:** A node in a binary tree can have two parents.
12. **True or False:** A node in a binary tree can have no parents.
13. **True or False:** The depth of a binary tree is always the same as the height of its root.
14. **True or False:** The level of a node is the height of the tree minus the height of the node.
15. **True or False:** Removing an item from a priority queue gives you the minimum element.
16. **True or False:** The height of an empty binary tree is not defined.
17. **True or False:** A heap is a complete binary tree that has the heap order property.
18. **True or False:** An ancestor of a node can never be in a subtree of the node.
19. **True or False:** A descendant of a node must belong to at least one subtree of the node.
20. **True or False:** The subtrees of a node in an AVL tree can be equal in height.
21. **True or False:** The subtrees of a node in a complete binary tree must be equal in height.

## Find the Error

Find the error in each of the following code segments.

```
1. // return number of nodes in a binary tree
 int NodeCount(Node tree)
 {
 int count = 0;
 if (tree != null)
```

```

{
 count++;
 NodeCount(tree.left);
 NodeCount(tree.right);
}
return count;
}

```

2. // Add a value to a binary search tree  
// Return root of resulting search tree

```

Node add(Node tree, int value)
{
 if (tree == null)
 return new Node(x);
 if (value < tree.value)
 return add(tree.left, value);
 else
 return add(tree.right, value);
}

```

3. // Add 1 to the value in each node of the tree

```

void incrementAll(Node tree)
{
 incrementAll(tree.left);
 incrementAll(tree.right);
 tree.Value++;
}

```

## Short Answer

1. Why is it good for binary search trees to be balanced?
2. A binary search tree has a right subtree but no left subtree. What node contains the least element in the tree?
3. What is the height of a binary tree?
4. Explain the concept of *order* as used with priority queues.
5. What is a priority queue?
6. Explain why the depth of a complete binary tree with  $n$  nodes is at most  $\log n + 1$ .
7. How does an AVL tree differ from an ordinary binary search tree?

## Algorithm Workbench

1. Write a method

```
void preorder(Node bTree)
```

that uses a preorder traversal to print all values stored in a binary tree.

2. Write a method

```
void postorder(Node bTree)
```

that uses a postorder traversal to print all values stored in a binary tree.

3. Write a method

```
int treeSize(Node bTree)
```

that computes and returns the number of nodes in a binary tree.

4. Write a method

```
int leafCount(Node btree)
```

that computes and returns the number of leaves in a binary tree.

5. Assume that data is stored in a binary tree, but that unlike the case of a binary search tree, no attempt is made to maintain any sort of order in the data stored. Give an algorithm for a method `contains` that searches a binary tree for a particular value `x` and returns `true` or `false` according to whether `x` is found in the tree.
6. Define a `Comparator` class `DecOrder` that orders integers in decreasing order, and give the code that shows how to create a `PriorityQueue` object that uses such a `Comparator` object.
7. Define an order on strings where a string `s` is less than `t` if the length of `s` is less than the length of `t`. If `s` and `t` have the same length, then `s` is less than `t` if `s` comes before `t` in alphabetic order. Write a `Comparator` class that corresponds to this order.

# Programming Challenges

## 1. AVL Trees

Add a `remove` method to the AVL tree class of this chapter. Use a graphical user interface that allows a user to work interactively with the AVL tree, adding and removing items as desired.

## 2. Family Tree

Write a program, with a graphical user interface, that allows users to create genealogy trees and work with them. Each person in this genealogy tree will be restricted to having at most two children. The nodes of the tree will store string values, which are interpreted as names of people. The program should support the following commands:

- `root name`: Adds `name` as the root of the tree. This must be the first command entered: it creates a tree with a single node, whose value field holds `name`. The created node becomes the root of the genealogy tree. The command is ignored if it is entered after the tree has been created.
- `left parent child`: Adds a new node `child` as the left child of the node with name `parent`. The command is ignored if there is no node with name `parent` in the tree, or if there is such a node, but it already has a left child.
- `right parent child`: Adds a new node `child` as the right child of the node with name `parent`. The command is ignored if there is no node with name `parent` in the tree, or if there is such a node, but it already has a right child.
- `descendants person`: Displays a list of all descendants of the node with name `person`. The command is ignored if there is no node in the tree with that name.
- `ancestors person`: Displays a list of all ancestors of the node with the name `person`. The command is ignored if there is no node in the tree with that name.

## 3. Balanced Search Tree

Write a program that can convert a sorted array into a balanced binary search tree. For this project, a balanced binary tree is one where the size of the left and right subtrees at each node differs by at most one. Your program should have a graphical user interface. The program allows the user to enter a number  $n$ , generates an array of  $n$  random integers, sorts the array, and then converts the sorted array into a balanced binary search tree. The program should display a graphical representation of the

## 4. Prefix Calculator



### Prefix Calculator

An arithmetic expression is in *prefix form* when operators are written before their operands. Here are some examples of prefix expressions and the values they evaluate to:

| Expression        | Value |
|-------------------|-------|
| 12                | 12    |
| + 2 51            | 53    |
| * 5 7             | 35    |
| * + 16 4 + 3 1 80 |       |

An expression (such as 12) that begins with an integer is a prefix expression that evaluates to itself. Otherwise, an expression is a prefix expression if it begins with an operator and is followed by two prefix expressions. In this latter case, the value of the expression is recursively computed from the values of its constituent prefix sub-expressions.

Write a program that allows the user to enter prefix expressions in a text field. The program reads the expression, evaluates it, and displays the value in a suitable GUI component. Assume that the user enters expressions that use only positive integers and the two operators + and \*. Your program should internally convert the prefix expression into a binary tree before evaluating it.

## 5. Levels in an AVL Tree

Extend the AVL tree class presented in this chapter by adding a method that takes a value, locates it in the AVL tree, and displays its level. Do this by extending the list of commands supported by the program to include commands of the form *level value*. For example, if the user types in



level 12

the program displays the level of a node containing 12. The program should display  $-1$  if the user types a value not in the tree. If there is more than one instance of the value in the tree, the program displays the level of any one of them.

## 6. Binary Tree Input and Display

You are given a file that contains data describing a binary tree whose nodes are strings. The file begins with an integer  $N$ , the number of nodes in the tree, on a line by itself. This first line is followed by  $N$  additional lines where each line specifies child information for a node in the tree. The line for a node  $X$  starts with  $X$  itself followed by a list of the children of  $X$ . Names of nodes are separated by whitespace. For example, the data set

```
5
Al Bob Carol
Bob Debby Elaine
Carol
Debby
Elaine
```

represents a binary tree in which Al has two children named Bob and Carol; Bob has two children named Debby and Elaine; and Carol, Debby, and Elaine have no children. Write a program that reads in data from such a file and displays it in a `JTree` component.

## 7. Breadth-First Traversal of Binary trees

Write a program that reads a file that describes a binary tree using the format of Programming Challenge 6. The program then lists all the nodes in the tree by their distance from the root: first the root, then all the children of the root, then all nodes that are children of a child of the root, and so on. For the example tree in Programming Challenge 6, the program might output the nodes in any of the following orders:

```
Al Bob Carol Debby Elaine
Al Carol Bob Elaine Debby
Al Bob Carol Debby Elaine
```

Such ordering are called breadth-first orders. Note that these three are not the only possible breadth-first orders.

# Chapter 22 Databases

## TOPICS


1. [22.1 Introduction to Database Management Systems](#)
2. [22.2 Tables, Rows, and Columns](#)
3. [22.3 Introduction to the SQL `SELECT` Statement](#)
4. [22.4 Inserting Rows](#)
5. [22.5 Updating and Deleting Existing Rows](#)
6. [22.6 Creating and Deleting Tables](#)
7. [22.7 Creating a New Database with JDBC](#)
8. [22.8 Scrollable Result Sets](#)
9. [22.9 Result Set Metadata](#)
10. [22.10 Relational Data](#)
11. [22.11 Advanced Topics](#)
12. [22.12 Common Errors to Avoid](#)

# 22.1 Introduction to Database Management Systems

CONCEPT: A database management system (DBMS) is software that manages large collections of data.

If an application needs to store only a small amount of data, text and binary files work well. These types of files, however, are not practical when a large amount of data must be stored and manipulated. Many businesses keep hundreds of thousands, or even millions, of data items in files. When a text or binary file contains this much data, simple operations such as searching, inserting, and deleting become cumbersome and inefficient.

When developing applications that work with an extensive amount of data, most developers prefer to use a database management system. A *database management system (DBMS)* is software specifically designed to store, retrieve, and manipulate large amounts of data in an organized and efficient manner. Once the data is stored using the database management system, applications may be written in Java or other languages to communicate with the DBMS. Rather than retrieving or manipulating the data directly, applications can send instructions to the DBMS. The DBMS carries out those instructions and sends the results back to the application. [Figure 22-1](#) illustrates this.



# Figure 22-1 A Java application interacts with a DBMS, which manipulates data

## [Figure 22-1 Full Alternative Text](#)

Although [Figure 22-1](#) is simplified, it illustrates the layered nature of an application that works with a database management system. The topmost layer of software, which in this case is written in Java, interacts with the user. It also sends instructions to the next layer of software, the DBMS. The DBMS works directly with the data, and sends the results of operations back to the application.

For example, suppose that a company keeps all of its product records in a database. The company has a Java application that allows the user to look up information on any product by entering its product ID number. The Java application instructs the DBMS to retrieve the record for the product with the specified product ID number. The DBMS retrieves the product record and sends the data back to the Java application. The Java application displays the data to the user.

The advantage of this layered approach to software development is that the Java programmer does not need to know about the physical structure of the data. He or she only needs to know how to write code that interacts with the DBMS. The DBMS handles the actual reading, writing, and searching of data.

## JDBC

[Figure 22-1](#) gives a simple illustration of a Java application communicating with a DBMS. The technology that makes this communication possible is known as *JDBC*, which stands for *Java Database Connectivity*. The Java API contains numerous JDBC classes that allow your Java applications to interact with a DBMS. This is illustrated in [Figure 22-2](#).




**Figure 22-2 A Java application uses the JDBC classes to interact with a DBMS**

[Figure 22-2 Full Alternative Text](#)

## SQL

*SQL*, which stands for *Structured Query Language*, is a standard language for working with database management systems. It was originally developed by IBM in the 1970s. Since then, SQL has been adopted by most database software vendors as the language of choice for interacting with its DBMS.

SQL consists of several key words. You use the key words to construct statements, which are also known as *queries*. These statements, or queries, are submitted to the DBMS, and are instructions for the DBMS to carry out operations on its data. When a Java application interacts with a DBMS, the Java application must construct SQL statements as strings, and then use an API method to pass those strings to the DBMS. In this chapter, you will learn how to construct simple SQL statements and then

pass them to a DBMS using an API method call.



Although SQL is a language, you don't use it to write applications. It is intended only as a standard means of interacting with a DBMS. You still need a general programming language, such as Java, to write an application for the ordinary user.

## Using a DBMS

In order to use JDBC to work with a database, you will need a DBMS installed on your system, or available to you in a school lab environment. There are many commercial DBMS packages available. Oracle, Microsoft SQL Server, DB2, and MySQL are just a few of the popular ones. In your school's lab, you may already have access to one of these, or perhaps another DBMS.

## Java DB and Apache Derby

If you do not have access to a DBMS in a school lab, you can use either Java DB or Apache Derby, depending on the version of the JDK that you are using. If you are using either Java 7 or Java 8, a DBMS named Java DB is automatically installed on your system when you install the JDK. Java DB is an open source distribution of Apache Derby, a pure Java DBMS that is designed specifically for Java applications.

Beginning with Java 9, however, Java DB is no longer bundled with the JDK. This is because Java DB is simply a rebranded version of Apache Derby. If you are using Java 9, you will need to separately download Apache Derby from the Apache Software Foundation's website and install it.

All of the examples in this chapter were created with Java DB and/or Apache Derby. If you wish to use either Java DB or Apache Derby, see Appendix I, available for download on this book's companion Web site at [www.pearson.com/gaddis](http://www.pearson.com/gaddis). Appendix I can help you configure Java DB, or download and install Apache Derby.

## Creating the `CoffeeDB` Database

In this chapter we will use a database named `CoffeeDB` as our example. The `CoffeeDB` database is used in the business operations of The Midnight Roastery, a small coffee roasting company. After you have installed either Java DB or Apache Derby, perform the following steps to create the `CoffeeDB` database:

1. Make sure you have downloaded student source code files from the book's companion Web site.
2. In this chapter's source code files, locate a program named `CreateCoffeeDB.java`.
3. Compile and execute the `CreateCoffeeDB.java` program. If either Java DB or Apache Derby is properly installed, this program will create the `CoffeeDB` database on your system.



If you are in a school lab environment using a DBMS other than Java DB or Apache Derby, consult with your instructor on how to modify the program to work with your specific DBMS.

## Connecting to the `CoffeeDB` Database

After installing either Java DB or Apache Derby and creating the `CoffeeDB` database, you should attempt to connect to the database with a Java program. A program can call the static JDBC method `DriverManager.getConnection` to get a connection to a database. There are overloaded versions of this method, but the simplest one has the following general format:

```
DriverManager.getConnection(DatabaseURL);
```

The method returns a reference to a `Connection` object, which we will discuss in a moment. In the general format, `DatabaseURL` is a string known as a *database URL*. URL stands for Uniform Resource Locator. A database URL lists the protocol that should be used to access the database, the name of the database, and potentially other items. A simple database URL has the following general format:

```
protocol:subprotocol:databaseName
```

In this very simple general format, three items are listed, separated by colons: *protocol*, *subprotocol*, and *databaseName*. Let's take a closer look at each one.

- *protocol* is the database protocol. When using JDBC, the protocol will always be `jdbc`.
- The value for *subprotocol* will be dependent upon the particular type of DBMS you are connecting to. If you are using either Java DB or Apache Derby, the subprotocol is `derby`.
- *databaseName* is the name of the database you are connecting to.

If we are using either Java DB or Apache Derby the URL for the `CoffeeDB` database is:

```
jdbc:derby:CoffeeDB
```

The `DriverManager.getConnection` method searches for and loads a JDBC driver that is

compatible with the database specified by the URL. A *JDBC driver* is a Java class that is designed to communicate with a specific DBMS. Each DBMS usually comes with its own JDBC driver. Typically, when you install a DBMS, you also update your system's `CLASSPATH` variable to include the JDBC driver's location. This will enable the JVM to find the driver class when you call the `DriverManager.getConnection` method.

When the `DriverManager.getConnection` method finds a compatible driver, it returns a `Connection` object. `Connection` is an interface in the `java.sql` package. You will need to use this `Connection` object to perform various tasks with the database, so save the reference in a variable. Here is an example of code that we can use in a Java application to get a connection to the `CoffeeDB` database using either Java DB or Apache Derby:

```
final String DB_URL = "jdbc:derby:CoffeeDB";
Connection conn = DriverManager.getConnection(DB_URL);
```

In the second statement shown here, we call the `DriverManager.getConnection` method, passing the URL for the `CoffeeDB` database. The method returns a reference to a `Connection` object, which we assign to the `conn` variable. If the `DriverManager.getConnection` method fails to load an appropriate driver for the specified database, it will throw an `SQLException`.

Before going any further, compile and execute the `TestConnection.java` program shown in [Code Listing 22-1](#). It demonstrates what we've covered so far. (This program assumes that either Java DB or Apache Derby has been installed, and that the `CoffeeDB` database has been created.)

## Code Listing 22-1 (TestConnection.java)

```
1 import java.sql.*; // Needed for JDBC classes
2
3 /**
4 * This program demonstrates how to connect to
5 * a Java DB or Apache Derby database using JDBC.
6 */
7
8 public class TestConnection
9 {
10 public static void main(String[] args)
11 {
12 // Create a named constant for the URL.
13 // NOTE: This value is specific for either Java DB or Apache Derby.
14 final String DB_URL = "jdbc:derby:CoffeeDB";
15
16 try
17 {
18 // Create a connection to the database.
19 Connection conn = DriverManager.getConnection(DB_URL);
20 System.out.println("Connection created to CoffeeDB.");
21
22 // Close the connection.
23 conn.close();
24 System.out.println("Connection closed.");
```



```

25 }
26 catch (Exception ex)
27 {
28 System.out.println("ERROR: " + ex.getMessage());
29 }
30 }
31 }

```

# Program Output

```

Connection created to CoffeeDB.
Connection closed.

```

Notice that line 1 imports all of the classes in the `java.sql` package. This package contains many of the necessary JDBC classes. Line 14 creates a string constant containing the URL for the `CoffeeDB` database.

JDBC methods throw an `SQLException` if they encounter a problem with a database. For that reason, we use a `try-catch` statement to handle any such exceptions. Let's take a closer look at the statements inside the `try` block:

- Line 19 does the following:
  - It declares a `Connection` variable named `conn`.
  - It calls the `DriverManager.getConnection` method to get a connection to the `CoffeeDB` database.
  - The `DriverManager.getConnection` method returns a reference to a `Connection` object. The reference is assigned to the `conn` variable.
- Line 20 displays a message indicating that a connection was created.
- Line 23 calls the `Connection` object's `close` method, which simply closes the database connection.
- Line 24 displays a message indicating that the connection is closed.

If a connection cannot be created in line 19, or the connection cannot be closed in line 23, an exception will be thrown. The `catch` clause in line 26 will handle the exception, and line 28 will display the exception object's default error message.

# Connecting to a Password-Protected Database

If the database that you are connecting to requires a user name and a password, you can use the

following form of the `DriverManager.getConnection` method:

```
DriverManager.getConnection(DatabaseURL, Username, Password);
```

In this general format, *Username* is a string containing a valid username, and *Password* is a string containing the password.



MyProgrammingLab™ [www.myprogramminglab.com](http://www.myprogramminglab.com)

1. 22.1 Why do most businesses use a DBMS to store their data instead of creating their own text files or binary files to hold the data?
2. 22.2 When a Java programmer uses a DBMS to store and manipulate data, why doesn't the programmer need to know specific details about the physical structure of the data?
3. 22.3 What is the technology that makes it possible for a Java application to communicate with a DBMS?
4. 22.4 What is the standard language for working with database management systems?
5. 22.5 What is a database URL?
6. 22.6 Suppose you have either Java DB or Apache Derby database on your system named `InventoryDB`. What database URL would you use in a Java program to get a connection to the database?
7. 22.7 What static JDBC method do you call to get a connection to a database?

# 22.2 Tables, Rows, and Columns

CONCEPT: Data that is stored in a database is organized into tables, rows, and columns.

A database management system stores data in a *database*. Your first step in learning to use a DBMS is to learn how data is organized inside a database. The data that is stored in a database is organized into one or more tables. Each *table* holds a collection of related data. The data that is stored in a table is then organized into rows and columns. A row is a complete set of information about a single item. The data that is stored in a row is divided into columns. Each column is an individual piece of information about the item.

The `CoffeeDB` database has a table named `Coffee`, which holds records for all of the different coffees sold by the company. [Table 22-1](#) shows the contents of the table.

Table 22-1 The `Coffee` database table

| Description             | ProdNum | Price |
|-------------------------|---------|-------|
| Bolivian Dark           | 14-001  | 8.95  |
| Bolivian Medium         | 14-002  | 8.95  |
| Brazilian Dark          | 15-001  | 7.95  |
| Brazilian Medium        | 15-002  | 7.95  |
| Brazilian Decaf         | 15-003  | 8.55  |
| Central American Dark   | 16-001  | 9.95  |
| Central American Medium | 16-002  | 9.95  |
| Sumatra Dark            | 17-001  | 7.95  |

|                      |        |       |
|----------------------|--------|-------|
| Sumatra Decaf        | 17-002 | 8.95  |
| Sumatra Medium       | 17-003 | 7.95  |
| Sumatra Organic Dark | 17-004 | 11.95 |
| Kona Medium          | 18-001 | 18.45 |
| Kona Dark            | 18-002 | 18.45 |
| French Roast Dark    | 19-001 | 9.65  |
| Galapagos Medium     | 20-001 | 6.85  |
| Guatemalan Dark      | 21-001 | 9.95  |
| Guatemalan Decaf     | 21-002 | 10.45 |
| Guatemalan Medium    | 21-003 | 9.95  |

As you can see, the table has 18 rows. Each row holds data about a type of coffee. The rows are divided into three columns. The first column is named `Description`, and it holds the description of a type of coffee. The second column is named `ProdNum`, and it holds a coffee's product number. The third column is named `Price`, and it holds a coffee's price per pound. As illustrated in [Figure 22-3](#), the third row in the table holds the following data:

- Description: Brazilian Dark
- Product Number: 15-001
- Price: 7.95

This row contains data about a single item.  
 Description: Brazilian Dark →  
 Product Number: 15-001  
 Price: 7.95

| Description             | ProdNum | Price |
|-------------------------|---------|-------|
| Bolivian Dark           | 14-001  | 8.95  |
| Bolivian Medium         | 14-002  | 8.95  |
| Brazilian Dark          | 15-001  | 7.95  |
| Brazilian Medium        | 15-002  | 7.95  |
| Brazilian Decaf         | 15-003  | 8.55  |
| Central American Dark   | 16-001  | 9.95  |
| Central American Medium | 16-002  | 9.95  |
| Sumatra Dark            | 17-001  | 7.95  |
| Sumatra Decaf           | 17-002  | 8.95  |
| Sumatra Medium          | 17-003  | 7.95  |
| Sumatra Organic Dark    | 17-004  | 11.95 |
| Kona Medium             | 18-001  | 18.45 |
| Kona Dark               | 18-002  | 18.45 |
| French Roast Dark       | 19-001  | 9.65  |
| Galapagos Medium        | 20-001  | 6.85  |
| Guatemalan Dark         | 21-001  | 9.95  |
| Guatemalan Decaf        | 21-002  | 10.45 |
| Guatemalan Medium       | 21-003  | 9.95  |

↑  
Description  
Column

↑  
ProdNum  
Column

↑  
Price  
Column

Figure 22-3 The Coffee database table

## Column Data Types

The columns in a database table are assigned a data type. Notice that the `Description` and `ProdNum` columns in the `Coffee` table hold strings, and the `Price` column holds floating-point numbers. The data types of the columns are not Java data types, however. Instead, they are SQL data types. [Table 22-2](#) lists a few of the standard SQL data types, and shows the Java data type that each is generally compatible with.

Table 22-2 A few of the SQL data types

| SQL Data Type | Description | Corresponding Java Data Type |
|---------------|-------------|------------------------------|
|---------------|-------------|------------------------------|

|                                              |                                                                                                  |                      |
|----------------------------------------------|--------------------------------------------------------------------------------------------------|----------------------|
| INTEGER or INT                               | An integer number                                                                                | int                  |
| CHARACTER( <i>n</i> )<br>or CHAR( <i>n</i> ) | A fixed-length string with a length of <i>n</i> characters                                       | String               |
| VARCHAR( <i>n</i> )                          | A variable-length string with a maximum length of <i>n</i> characters                            | String               |
| REAL                                         | A single-precision floating-point number                                                         | float                |
| DOUBLE                                       | A double-precision floating-point number                                                         | double               |
| DECIMAL( <i>t</i> , <i>d</i> )               | A decimal value with <i>t</i> total digits and <i>d</i> digits appearing after the decimal point | java.math.BigDecimal |
| DATE                                         | A date                                                                                           | java.sql.Date        |

There are many other standard data types in SQL. When the `Coffee` table was created, the following data types were used for the columns:

- The data type for the `Description` column is `CHAR(25)`. This means that each value in the `Description` column is a string with a fixed length of 25 characters, compatible with the `String` type in Java.
- The data type for the `ProdNum` column is `CHAR(10)`. This means that each value in the `ProdNum` column is a string with a fixed length of 10 characters, compatible with the `String` type in Java.
- The data type for the `Price` column is `DOUBLE`. This means that each value in the `Price` column is a double-precision floating-point number, compatible with the `double` data type in Java.

## Primary Keys

Most database tables have a *primary key*, which is a column that can be used to identify a specific row in a table. The column that is designated as the primary key holds a unique value for each row. If you try to store duplicate data in the primary key column, an error will occur.

In the `Coffee` table, the `ProdNum` column is the primary key because it holds a unique product number for each type of coffee. Here are some other examples:

- Suppose a table stores employee data, and one of the columns holds employee ID numbers. Because each employee's ID number is unique, this column can be used as the primary key.
- Suppose a table stores data about a cell phone company's inventory of phones, and one of the columns holds cell phone serial numbers. Because each phone's serial number is unique, this column can be used as the primary key.
- Suppose a table stores invoice data, and one of the columns holds invoice numbers. Each invoice has a unique invoice number, so this column can be used as a primary key.



It is possible for a table's primary key to be the combination of several columns in the table.



MyProgrammingLab™ [www.myprogramminglab.com](http://www.myprogramminglab.com)

1. 22.8 Describe how the data that is stored in a table is organized.
2. 22.9 What is a primary key?
3. 22.10 What Java data types correspond with the following SQL types?
  - INTEGER
  - INT
  - REAL
  - CHAR
  - CHARACTER
  - VARCHAR
  - DOUBLE

# 22.3 Introduction to the SQL `SELECT` Statement

CONCEPT: The `SELECT` statement is used in SQL to retrieve data from a database.

The first SQL statement we will discuss is the `SELECT` statement. You use the *`SELECT` statement* to retrieve the rows in a table. As its name implies, the `SELECT` statement allows you to select specific rows. We will start with a very simple form of the statement, as shown here:

```
SELECT Columns FROM Table
```

In the general form, *Columns* is one or more column names, and *Table* is a table name. Here is an example `SELECT` statement that we might execute on the `CoffeeDB` database:

```
SELECT Description FROM Coffee
```

This statement will retrieve the `Description` column for every row in the `Coffee` table. [Figure 22-4](#) shows an example of the results.

## Figure 22-4 Description column

| Description             |
|-------------------------|
| Bolivian Dark           |
| Bolivian Medium         |
| Brazilian Dark          |
| Brazilian Medium        |
| Brazilian Decaf         |
| Central American Dark   |
| Central American Medium |



Sumatra Dark

Sumatra Decaf

Sumatra Medium

Sumatra Organic Dark

Kona Medium

Kona Dark

French Roast Dark

Galapagos Medium

Guatemalan Dark

Guatemalan Decaf

Guatemalan Medium

[Figure 22-4](#) shows the results of a `SELECT` statement, but what happens to these results? In a Java program, the results of a `SELECT` statement are returned to the program in a `ResultSet` object. A *ResultSet object* is simply an object that contains the results of an SQL statement. The process of sending an SQL statement to a DBMS can be summarized in the following steps:

1. Get a connection to the database.
2. Pass a string containing an SQL statement to the DBMS. If the SQL statement has results to send back, a `ResultSet` object will be returned to the program.
3. Process the contents of the `ResultSet` object, if one has been returned to the program.

4. When finished working with the database, close the connection.

You previously saw, in [Code Listing 22-1](#), an example of how to perform step 1 (get a connection to the database) and step 4 (close the connection). Next we look at the details of how an SQL statement is sent to the DBMS and how its results are processed in steps 2 and 3.

## Passing an SQL Statement to the DBMS


Once you have gotten a connection to the database, you are ready to issue SQL statements to the DBMS. First, you must get a `Statement` object from the `Connection` object, using its `createStatement` method. Here is an example:

```
Statement stmt = conn.createStatement();
```

After this code executes, the `stmt` variable will reference a `Statement` object. `Statement` is an interface in the `java.sql` package. `Statement` objects have a variety of methods that can be used to execute SQL queries. To execute a `SELECT` query, you use the `executeQuery` method. The method returns a `ResultSet` object. Here is an example:

```
String sqlStatement = "SELECT Description FROM Coffee";
ResultSet result = stmt.executeQuery(sqlStatement);
```

The first statement creates a string containing an SQL query. The second statement passes this string as an argument to the `executeQuery` method. The method returns a reference to a `ResultSet` object containing the results of the query. The reference is assigned to the `result` variable. [Figure 22-5](#) illustrates how the `result` variable references the `ResultSet` object.



## Figure 22-5 A `ResultSet` object contains the results of an SQL query

A `ResultSet` object contains a set of rows and columns. The `ResultSet` object in [Figure 22-5](#) has 18 rows and one column. The rows in a `ResultSet` are numbered, with the first row being row 1, the second row being row 2, and so forth. The columns also are numbered, with the first column being column 1, the second column being column 2, and so forth. [Figure 22-6](#) shows the same `ResultSet` with the row and column numbers labeled.

|        |  | Column 1                |
|--------|--|-------------------------|
| Row 1  |  | Bolivian Dark           |
| Row 2  |  | Bolivian Medium         |
| Row 3  |  | Brazilian Dark          |
| Row 4  |  | Brazilian Medium        |
| Row 5  |  | Brazilian Decaf         |
| Row 6  |  | Central American Dark   |
| Row 7  |  | Central American Medium |
| Row 8  |  | Sumatra Dark            |
| Row 9  |  | Sumatra Decaf           |
| Row 10 |  | Sumatra Medium          |
| Row 11 |  | Sumatra Organic Dark    |
| Row 12 |  | Kona Medium             |
| Row 13 |  | Kona Dark               |
| Row 14 |  | French Roast Dark       |
| Row 15 |  | Galapagos Medium        |
| Row 16 |  | Guatemalan Dark         |
| Row 17 |  | Guatemalan Decaf        |
| Row 18 |  | Guatemalan Medium       |

## Figure 22-6 `ResultSet` rows and columns

### Getting a Row from the `ResultSet` Object

A `ResultSet` object has an internal *cursor* that points to a specific row in the `ResultSet`. The row that the cursor points to is considered the *current row*. The cursor can be moved from row to row, and this provides you with a way to examine all of the rows in the `ResultSet`.

At first, the cursor is not pointing to a row, but is positioned just before the first row. This is illustrated in [Figure 22-7](#).

Initially the cursor is positioned just before the first row in the `ResultSet`.

| Cursor → | Column 1                |
|----------|-------------------------|
| Row 1    | Bolivian Dark           |
| Row 2    | Bolivian Medium         |
| Row 3    | Brazilian Dark          |
| Row 4    | Brazilian Medium        |
| Row 5    | Brazilian Decaf         |
| Row 6    | Central American Dark   |
| Row 7    | Central American Medium |
| Row 8    | Sumatra Dark            |
| Row 9    | Sumatra Decaf           |
| Row 10   | Sumatra Medium          |
| Row 11   | Sumatra Organic Dark    |
| Row 12   | Kona Medium             |
| Row 13   | Kona Dark               |
| Row 14   | French Roast Dark       |
| Row 15   | Galapagos Medium        |
| Row 16   | Guatemalan Dark         |
| Row 17   | Guatemalan Decaf        |
| Row 18   | Guatemalan Medium       |

## Figure 22-7 The cursor is initially positioned before the first row

To move the cursor to the first row in a newly created `ResultSet`, you call the object's `next` method. Here is an example:

```
result.next();
```

Assuming that `result` references a newly created `ResultSet` object, this statement moves the cursor to the first row in the `ResultSet`. [Figure 22-8](#) shows how the cursor has moved to the first row in the `ResultSet` after the `next` method is called the first time.

After the `ResultSet` object's `next` method is called the first time, the cursor is positioned at the first row.

| Column 1 |        |                         |
|----------|--------|-------------------------|
| Cursor → | Row 1  | Bolivian Dark           |
|          | Row 2  | Bolivian Medium         |
|          | Row 3  | Brazilian Dark          |
|          | Row 4  | Brazilian Medium        |
|          | Row 5  | Brazilian Decaf         |
|          | Row 6  | Central American Dark   |
|          | Row 7  | Central American Medium |
|          | Row 8  | Sumatra Dark            |
|          | Row 9  | Sumatra Decaf           |
|          | Row 10 | Sumatra Medium          |
|          | Row 11 | Sumatra Organic Dark    |
|          | Row 12 | Kona Medium             |
|          | Row 13 | Kona Dark               |
|          | Row 14 | French Roast Dark       |
|          | Row 15 | Galapagos Medium        |
|          | Row 16 | Guatemalan Dark         |
|          | Row 17 | Guatemalan Decaf        |
|          | Row 18 | Guatemalan Medium       |

## Figure 22-8 The `next` method moves the cursor forward

[Figure 22-8 Full Alternative Text](#)

Each time you call the `next` method, it moves the cursor to the next row in the `ResultSet`.

The `next` method returns a `boolean` value. It returns `true` if the cursor successfully moved to the next row. If there are no more rows, it returns `false`. The following code shows how you can move the cursor through all of the rows of a newly created `ResultSet`.

```
while (result.next())
{
 // Process the current row.
}
```

There are other `ResultSet` methods for navigating the rows in a `ResultSet` object. We will look at some of them later in this chapter.

## Getting Columns in a `ResultSet` Row

You use one of the `ResultSet` object's "get" methods to retrieve the contents of a specific column in

the current row. When you call one of these methods, you can pass either the column number or the column name as an argument. There are numerous “get” methods defined in the `ResultSet` interface. [Table 22-3](#) lists a few of them.

Table 22-3 A few of the `ResultSet` methods

| ResultSet Method                                                            | Description                                                                                                                                                                                                                                                                        |
|-----------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>double getDouble(int colNumber) double getDouble(String colName)</pre> | Returns the <code>double</code> that is stored in the column - specified by <i>colNumber</i> or <i>colName</i> . The column must hold data that is compatible with the <code>double</code> data type in Java. If an error occurs, the method throws an <code>SQLException</code> . |
| <pre>int getInt(int colNumber) int getInt(String colName)</pre>             | Returns the <code>int</code> that is stored in the column - specified by <i>colNumber</i> or <i>colName</i> . The column must hold data that is compatible with the <code>int</code> data type in Java. If an error occurs, the method throws an <code>SQLException</code> .       |
| <pre>String getString(int colNumber) String getString(String colName)</pre> | Returns the <code>string</code> that is stored in the column - specified by <i>colNumber</i> or <i>colName</i> . The column must hold data that is compatible with the <code>String</code> type in Java. If an error occurs, the method throws an <code>SQLException</code> .      |

Recall that columns have an SQL data type. The SQL data types are not the same as the Java data types, but are compatible with them. To retrieve the contents of a column, you call the method that is designed to work with that column’s data type. For example, if the column contains a string, you would use the `getString` method to retrieve its value. If the column contains an integer, you would use the `getInt` method. Likewise, if the column contains a `double`, you would call the `getDouble` method.

When you call one of the “get” methods, you must tell it which column in the current row you want to retrieve. These methods accept either an integer argument, which is a column number, or a `String` holding the column name.

The `ResultSet` that we have been looking at in our example has only one column: the `Description` column. The `Description` column's data type is `CHAR(25)`, which means it is a fixed-length string of 25 characters. This is compatible with the `String` type in Java. To display the contents of the `Description` column in the current row, we could use the following statement:

```
System.out.println(result.getString("Description"));
```

The `Description` column holds values that are compatible with the `String` type, so we use the `getString` method to retrieve its contents. We could also use the column number to retrieve the column contents. Here is an example:

```
System.out.println(result.getString(1));
```

## NOTE:

Column names in a database table are not case sensitive. The column names `DESCRIPTION`, `description`, and `Description` are all the same.

Let's look at a complete program that demonstrates what we have covered so far. [Code Listing 22-2](#) displays the `Description` column from all of the rows in the `CoffeeDB` database.

## Code Listing 22-2 (`ShowCoffeeDescriptions.java`)

```
1 import java.sql.*; // Needed for JDBC classes
2
3 /**
4 * This program demonstrates how to issue an SQL
5 * SELECT statement to a database using JDBC.
6 */
7
8 public class ShowCoffeeDescriptions
9 {
10 public static void main(String[] args)
11 {
12 // Create a named constant for the URL.
13 // NOTE: This value is specific for either Java DB or Apache Derby.
14 final String DB_URL = "jdbc:derby:CoffeeDB";
15
16 try
17 {
18 // Create a connection to the database.
19 Connection conn = DriverManager.getConnection(DB_URL);
20
21 // Create a Statement object.
22 Statement stmt = conn.createStatement();
23
```

```

24 // Create a string with a SELECT statement.
25 String sqlStatement = "SELECT Description FROM Coffee";
26
27 // Send the statement to the DBMS.
28 ResultSet result = stmt.executeQuery(sqlStatement);
29
30 // Display a header for the listing.
31 System.out.println("Coffees Found in the Database");
32 System.out.println("-----");
33
34 // Display the contents of the result set.
35 // The result set will have three columns.
36 while (result.next())
37 {
38 System.out.println(result.getString("Description"));
39 }
40
41 // Close the connection.
42 conn.close();
43 }
44 catch(Exception ex)
45 {
46 System.out.println("ERROR: " + ex.getMessage());
47 }
48 }
49 }

```

## Program Output

```

Coffees Found in the Database

Bolivian Dark
Bolivian Medium
Brazilian Dark
Brazilian Medium
Brazilian Decaf
Central American Dark
Central American Medium
Sumatra Dark
Sumatra Decaf
Sumatra Medium
Sumatra Organic Dark
Kona Medium
Kona Dark
French Roast Dark
Galapagos Medium
Guatemalan Dark
Guatemalan Decaf
Guatemalan Medium

```

Let's take a closer look at the code. Line 14 declares a string constant, initialized with the URL for the `CoffeeDB` database. The statements that access the database are written inside the try block that appears in lines 18 through 42. Line 19 gets a connection to the database. After line 19 executes, the `conn` variable will reference a `Connection` object that can be used to access the database.



At this point in the program, we have a connection to the `CoffeeDB` database, but we are not ready to send a `SELECT` statement to the database. In order to send a `SELECT` statement to the database, we must have a `Statement` object. Line 22 calls the `Connection` object's `createStatement` method, which returns a reference to a `Statement` object. The reference is assigned to the `stmt` variable.

Line 25 declares a `String` variable named `sqlStatement`, initialized with the string `"SELECT Description FROM Coffee"`. This is the SQL statement that we want to submit to the database. Line 28 passes this string as an argument to the `Statement` object's `executeQuery` method, which executes the statement. The method returns a reference to a `ResultSet` object, which is assigned to the `result` variable. The `ResultSet` object contains the results of the `SELECT` statement.

The `while` loop that appears in lines 36 through 39 displays the contents of the `ResultSet` object. It works like this:

- The `while` statement in line 36 calls the `ResultSet` object's `next` method to advance the internal cursor. If the cursor is successfully advanced, the method returns `true` and the loop iterates. If the cursor is at the end of the `ResultSet` object's rows, the method returns `false` and the loop terminates.
- Each time the loop iterates, the `ResultSet` object's internal cursor will be positioned at a specific row. The statement in line 38 gets the value of the `Description` column and displays it.

Line 42 closes the connection to the database.

## More about the `SELECT` Statement

You can specify more than one column in a `SELECT` statement by separating the column names with commas. Here is an example:

```
SELECT Description, Price FROM Coffee
```

This statement will retrieve the `Description` column and the `Price` column for every row in the `Coffee` table. The program shown in [Code Listing 22-3](#) demonstrates.

## Code Listing 22-3 (`ShowDescriptionsAndPrices.java`)

```
1 import java.sql.*; // Needed for JDBC classes
2
3 /**
4 * This program displays the coffee descriptions
5 * and their prices.
6 */
7
```

```

8 public class ShowDescriptionsAndPrices
9 {
10 public static void main(String[] args)
11 {
12 // Create a named constant for the URL.
13 // NOTE: This value is specific for either Java DB or Apache Derby.
14 final String DB_URL = "jdbc:derby:CoffeeDB";
15
16 try
17 {
18 // Create a connection to the database.
19 Connection conn = DriverManager.getConnection(DB_URL);
20
21 // Create a Statement object.
22 Statement stmt = conn.createStatement();
23
24 // Create a string with a SELECT statement.
25 String sqlStatement =
26 "SELECT Description, Price FROM Coffee";
27
28 // Send the statement to the DBMS.
29 ResultSet result = stmt.executeQuery(sqlStatement);
30
31 // Display the contents of the result set.
32 // The result set will have three columns.
33 while (result.next())
34 {
35 System.out.printf("%25s %.2f\n",
36 result.getString("Description"),
37 result.getDouble("Price"));
38 }
39
40 // Close the connection.
41 conn.close();
42 }
43 catch(Exception ex)
44 {
45 System.out.println("ERROR: " + ex.getMessage());
46 }
47 }
48 }

```

# Program Output

|               |      |
|---------------|------|
| Bolivian Dark | 8.95 |
|---------------|------|

|                 |      |
|-----------------|------|
| Bolivian Medium | 8.95 |
|-----------------|------|

|                |      |
|----------------|------|
| Brazilian Dark | 7.95 |
|----------------|------|

|                         |       |
|-------------------------|-------|
| Brazilian Medium        | 7.95  |
| Brazilian Decaf         | 8.55  |
| Central American Dark   | 9.95  |
| Central American Medium | 9.95  |
| Sumatra Dark            | 7.95  |
| Sumatra Decaf           | 8.95  |
| Sumatra Medium          | 7.95  |
| Sumatra Organic Dark    | 11.95 |
| Kona Medium             | 18.45 |
| Kona Dark               | 18.45 |
| French Roast Dark       | 9.65  |
| Galapagos Medium        | 6.85  |
| Guatemalan Dark         | 9.95  |
| Guatemalan Decaf        | 10.45 |
| Guatemalan Medium       | 9.95  |

The program in [Code Listing 22-3](#) is very similar to that in [Code Listing 22-2](#). The differences between the two programs are summarized here:

- Lines 25 and 26 initialize the `sqlStatement` variable with the string `"SELECT Description, Price FROM Coffee"`. This is a `SELECT` statement that will retrieve the `Description` and



```
37 result.getDouble("Price"));
38 }
39
40 // Close the connection.
41 conn.close();
42 }
43 catch(Exception ex)
44 {
45 System.out.println("ERROR: " + ex.getMessage());
46 }
47 }
48 }
```

# Program Output

|                         |              |
|-------------------------|--------------|
| Bolivian Dark           | 14-001 8.95  |
| Bolivian Medium         | 14-002 8.95  |
| Brazilian Dark          | 15-001 7.95  |
| Brazilian Medium        | 15-002 7.95  |
| Brazilian Decaf         | 15-003 8.55  |
| Central American Dark   | 16-001 9.95  |
| Central American Medium | 16-002 9.95  |
| Sumatra Dark            | 17-001 7.95  |
| Sumatra Decaf           | 17-002 8.95  |
| Sumatra Medium          | 17-003 7.95  |
| Sumatra Organic Dark    | 17-004 11.95 |
| Kona Medium             | 18-001 18.45 |

|                   |              |
|-------------------|--------------|
| Kona Dark         | 18-002 18.45 |
| French Roast Dark | 19-001 9.65  |
| Galapagos Medium  | 20-001 6.85  |
| Guatemalan Dark   | 21-001 9.95  |
| Guatemalan Decaf  | 21-002 10.45 |
| Guatemalan Medium | 21-003 9.95  |

SQL statements are free-form, which means that tabs, newlines, and spaces between the key words are ignored. For example, the statement

```
SELECT * FROM Coffee
```

works the same as:

```
SELECT
 *
FROM
 Coffee
```

In addition, SQL key words and table names are case insensitive. The previous statement could be written as:

```
select * from coffee
```

## Specifying Search Criteria with the WHERE Clause

Occasionally you might want to retrieve every row in a table. For example, if you wanted a list of all the coffees in the `Coffee` table, the previous `SELECT` statement would give it to you. Normally, however, you want to narrow the list down to only a few selected rows in the table. That's where the `WHERE` clause comes in. The `WHERE` clause can be used with the `SELECT` statement to specify search criteria. When you use the `WHERE` clause, only the rows that meet the search criteria will be returned in the result set. The general format of a `SELECT` statement with a `WHERE` clause is:

`SELECT Columns FROM Table WHERE Criteria`

In the general format, *Criteria* is a conditional expression. Here is an example of a `SELECT` statement that uses the `WHERE` clause:

```
SELECT * FROM Coffee WHERE Price > 12.00
```

The first part of the statement, `SELECT * FROM Coffee`, specifies that we want to see every column. The `WHERE` clause specifies that we want only the rows in which the content of the `Price` column is greater than 12.00. [Figure 22-9](#) shows the results of this statement. Notice that only two coffees meet this search criterion.

# Figure 22-9 Rows where Price is greater than 12.00

| Description | ProdNum | Price |
|-------------|---------|-------|
|-------------|---------|-------|

|             |        |       |
|-------------|--------|-------|
| Kona Medium | 18-001 | 18.45 |
|-------------|--------|-------|

|           |        |       |
|-----------|--------|-------|
| Kona Dark | 18-002 | 18.45 |
|-----------|--------|-------|

Standard SQL supports the relational operators listed in [Table 22-4](#) for writing conditional expressions in a `WHERE` clause.

## Table 22-4 SQL relational operators

| Operator | Meaning                  |
|----------|--------------------------|
| >        | Greater-than             |
| <        | Less-than                |
| >=       | Greater-than or equal-to |

<=          Less-than or equal-to

=            Equal-to

<>          Not equal-to

Notice that the equal-to and not equal-to operators in SQL are different from those in Java. The equal-to operator is one equal sign, not two equal signs. The not equal-to operator is <>.

Let's look at a few more examples of the `SELECT` statement. The following statement could be used to retrieve the product numbers and prices of all the coffees that are priced at 7.95:

```
SELECT ProdNum, Price FROM Coffee WHERE Price = 7.95
```

The results of this statement are shown in [Figure 22-10](#).

## Figure 22-10 Results of SQL statement

| ProdNum | Price |
|---------|-------|
|---------|-------|

|        |      |
|--------|------|
| 16-001 | 7.95 |
|--------|------|

|        |      |
|--------|------|
| 16-002 | 7.95 |
|--------|------|

|        |      |
|--------|------|
| 17-001 | 7.95 |
|--------|------|

|        |      |
|--------|------|
| 17-003 | 7.95 |
|--------|------|

The following `SELECT` statement retrieves all of the columns for the row where the description is “French Roast Dark”. The results returned from this statement are shown in [Figure 22-11](#).

```
SELECT * FROM Coffee WHERE Description = 'French Roast Dark'
```

## Figure 22-11 Results of SQL statement



| Description | ProdNum | Price |
|-------------|---------|-------|
|-------------|---------|-------|

|                   |        |      |
|-------------------|--------|------|
| French Roast Dark | 19-001 | 9.65 |
|-------------------|--------|------|

If you look carefully at the previous statement you will notice another difference between SQL syntax and Java syntax. In SQL, string literals are enclosed in single quotes, not double quotes.



If you need to include a single quote as part of a string, simply write two single quotes in its place. For example, suppose you wanted to search the `Coffee` table for a coffee named Joe's Special Blend. You could use the following statement:

```
SELECT * FROM Coffee WHERE Description = 'Joe''s Special Blend'
```

Let's look at an example program that uses a `WHERE` clause in a `SELECT` statement. The program in [Code Listing 22-5](#) lets the user enter a minimum price, and then search the `Coffee` table for rows where the `Price` column is greater than or equal to the specified price.

## Code Listing 22-5 (CoffeeMinPrice.java)

```
1 import java.util.Scanner;
2 import java.sql.*;
3
4 /**
5 * This program lets the user search for coffees
6 * priced at a minimum value.
7 */
8
9 public class CoffeeMinPrice
10 {
11 public static void main(String[] args)
12 {
13 double minPrice; // To hold the minimum price
14 int coffeeCount = 0; // To count coffees that qualify
15
16 // Create a named constant for the URL.
17 // NOTE: This value is specific for either Java DB or Apache Derby.
18 final String DB_URL = "jdbc:derby:CoffeeDB";
19
20 // Create a Scanner object for keyboard input.
21 Scanner keyboard = new Scanner(System.in);
22
23 // Get the minimum price from the user.
24 System.out.print("Enter the minimum price: ");
25 minPrice = keyboard.nextDouble();
```

```

26
27 try
28 {
29 // Create a connection to the database.
30 Connection conn = DriverManager.getConnection(DB_URL);
31
32 // Create a Statement object.
33 Statement stmt = conn.createStatement();
34
35 // Create a string containing a SELECT statement.
36 // Note that we are incorporating the user's input
37 // into the string.
38 String sqlStatement =
39 "SELECT * FROM Coffee WHERE Price >= " +
40 Double.toString(minPrice);
41
42 // Send the statement to the DBMS.
43 ResultSet result = stmt.executeQuery(sqlStatement);
44
45 // Display the contents of the result set.
46 // The result set will have three columns.
47 while (result.next())
48 {
49 // Display a row from the result set.
50 System.out.printf("%25s %10s %5.2f\n",
51 result.getString("Description"),
52 result.getString("ProdNum"),
53 result.getDouble("Price"));
54
55 // Increment the counter.
56 coffeeCount++;
57 }
58
59 // Display the number of qualifying coffees.
60 System.out.println(coffeeCount + " coffees found.");
61
62 // Close the connection.
63 conn.close();
64 }
65 catch(Exception ex)
66 {
67 System.out.println("ERROR: " + ex.getMessage());
68 }
69 }
70 }

```

# Program Output

Enter the minimum price: **12.00** [Enter]

Kona Medium 18-001 18.45

Kona Dark 18-002 18.45

2 coffees found.

# Program Output

Enter the minimum price: 10.00 [Enter]

Sumatra Organic Dark 17-004 11.95

Kona Medium 18-001 18.45

Kona Dark 18-002 18.45

Guatemalan Decaf 21-002 10.45

4 coffees found.

# Program Output

Enter the minimum price: 20.00 [Enter]

0 coffees found.

There are a few things in [Code Listing 22-5](#) that deserve some explanation. In lines 24 and 25 the program prompts the user to enter a minimum price, and the user’s input is assigned to the double variable `minPrice`. Then, notice in lines 38 through 40 that the `minPrice` variable is converted to a string, and concatenated onto the string containing the `SELECT` statement. When the program runs, if the user enters 10.00, the `SELECT` statement that is created in lines 38 through 40 will be:

```
SELECT * FROM Coffee WHERE Price >= 10.00
```

Or, if the user enters 12.00, the `SELECT` statement that is created in lines 38 through 40 will be:

```
SELECT * FROM Coffee WHERE Price >= 12.00
```

Programs commonly need to use techniques such as this to create SQL statements that incorporate user input.

# String Comparisons in a SELECT Statement

String comparisons in SQL are case sensitive. If you ran the following statement against the `CoffeeDB` database, you would not get any results:

```
SELECT * FROM Coffee WHERE Description = 'french roast dark'
```

However, you can use the `UPPER()` function to convert a string to uppercase. Here is an example:

```
SELECT * FROM Coffee WHERE UPPER(Description) = 'FRENCH ROAST DARK'
```

This statement will return the same results as shown in [Figure 22-11](#). SQL also provides a `LOWER()` function, which converts its argument to lowercase.

# Using the `LIKE` Operator

Sometimes searching for an exact string will not yield the results you want. For example, suppose we want a list of all the decaf coffees in the `Coffee` table. The following statement will not work. Can you see why?

```
SELECT * FROM Coffee WHERE Description = 'Decaf'
```

This statement will search for rows where the `Description` field is equal to the string “Decaf”. Unfortunately, it will find none. If you glance back at [Table 22-1](#), you will see that none of the rows in the `Coffee` table have a `Description` column that is equal to “Decaf”. You will also see, however, that the word “Decaf” does appear in the `Description` column of some of the rows. For example, in one row you will find “Brazilian Decaf”. In another row you will find “Sumatra Decaf”. In yet another row you will find “Guatemalan Decaf”. In addition to the word “Decaf”, each of these strings contains other characters.

In order to find all of the decaf coffees, we need to search for rows where “Decaf” appears as a substring in the `Description` column. You can perform such a search using the `LIKE` operator. Here is an example of how to use it.

```
SELECT * FROM Coffee WHERE Description LIKE '%Decaf%'
```

The `LIKE` operator is followed by a string that contains a character pattern. In this example, the character pattern is `'%Decaf%'`. The `%` symbol is used as a wildcard character. It represents any sequence of zero or more characters. The pattern `'%Decaf%'` specifies that the string “Decaf” must appear with any sequence of characters before or after it. The results of this statement are shown in [Figure 22-12](#).

Figure 22-12 Results of SQL statement

| Description | ProdNum | Price |
|-------------|---------|-------|
|-------------|---------|-------|

|                 |        |      |
|-----------------|--------|------|
| Brazilian Decaf | 15-003 | 8.55 |
|-----------------|--------|------|

|               |        |      |
|---------------|--------|------|
| Sumatra Decaf | 17-002 | 8.95 |
|---------------|--------|------|

|                  |        |       |
|------------------|--------|-------|
| Guatemalan Decaf | 21-002 | 10.45 |
|------------------|--------|-------|

Likewise, the following statement will result in all the rows where the `Description` column starts with the word “Sumatra”.

```
SELECT * FROM Coffee WHERE Description LIKE 'Sumatra%'
```

The underscore character ( `_` ) is also used as a wildcard. Unlike the `%` character, the underscore represents a single character. For example, look at the following statement.

```
SELECT * FROM Coffee WHERE ProdNum LIKE '2_-00_'
```

This statement will result in all the rows where the `ProdNum` column begins with “2”, followed by any single character, followed by “-00”, followed by any single character. The results of this statement are shown in [Figure 22-13](#).

## Figure 22-13 Results of SQL statement

| Description       | ProdNum | Price |
|-------------------|---------|-------|
| Galapagos Medium  | 20-001  | 6.85  |
| Guatemalan Dark   | 21-001  | 9.95  |
| Guatemalan Decaf  | 21-002  | 10.45 |
| Guatemalan Medium | 21-003  | 9.95  |

You can use the `NOT` operator to disqualify a character pattern in a search criterion. For example, suppose that you want a list of all the coffees that are not decaf. The following statement will yield just those results.

```
SELECT * FROM Coffee WHERE Description NOT LIKE '%Decaf%'
```

# Using AND and OR

You can use the `AND` and `OR` logical operators to specify multiple search criteria in a `WHERE` clause. For example, look at the following statement:

```
SELECT * FROM Coffee WHERE Price > 10.00 AND Price < 14.00
```

The `AND` operator requires that both of the search criteria be true in order for a row to be qualified as a match. The only rows that will be returned from this statement are those where the `Price` column contains a value that is greater than 10.00 and less than 14.00. [Figure 22-14](#) shows the results of the statement.

## Figure 22-14 Results of SQL statement

| Description          | ProdNum | Price |
|----------------------|---------|-------|
| Sumatra Organic Dark | 17-004  | 11.95 |
| Guatemalan Decaf     | 21-002  | 10.45 |

Here’s an example that uses the `OR` operator:

```
SELECT * FROM Coffee
WHERE Description LIKE '%Dark%' OR ProdNum LIKE '16%'
```

The `OR` operator requires that either of the search criteria be true in order for a row to be qualified as a match. This statement searches for rows where the `Description` column contains the string “Dark” at any position, or where the `ProdNum` column starts with “16”. [Figure 22-15](#) shows the results of the query.

## Figure 22-15 Results of SQL statement

| Description   | ProdNum | Price |
|---------------|---------|-------|
| Bolivian Dark | 14-001  | 8.95  |

|                         |        |       |
|-------------------------|--------|-------|
| Brazilian Dark          | 15-001 | 7.95  |
| Central American Dark   | 16-001 | 9.95  |
| Central American Medium | 16-002 | 9.95  |
| Sumatra Dark            | 17-001 | 7.95  |
| Sumatra Organic Dark    | 17-004 | 11.95 |
| Kona Dark               | 18-002 | 18.45 |
| French Roast Dark       | 19-001 | 9.65  |
| Guatemalan Dark         | 21-001 | 9.95  |

# Sorting the Results of a `SELECT` Query

If you wish to sort the results of a `SELECT` query, you can use the `ORDER BY` clause. Here is an example:

```
SELECT * FROM Coffee ORDER BY Price
```

This statement will produce a list of all the rows in the `Coffee` table, ordered by the `Price` column. The list will be sorted in ascending order on the `Price` column, meaning that the lowest-priced coffees will appear first.

Here's a `SELECT` statement that uses both a `WHERE` clause and an `ORDER BY` clause:

```
SELECT * FROM Coffee
WHERE Price > 9.95
ORDER BY Price
```

This statement will produce a list of all the rows in the `Coffee` table where the `Price` column contains a value greater than 9.95, listed in ascending order by price.

If you want the list sorted in descending order (from highest to lowest), use the `DESC` operator, as shown here:

```
SELECT * FROM Coffee
WHERE Price > 9.95
ORDER BY Price DESC
```

# Mathematical Functions

SQL provides several functions for performing calculations. For example, the `AVG` function calculates the average value in a particular column. Here is an example `SELECT` statement using the `AVG` function:

```
SELECT AVG(Price) FROM Coffee
```

This statement produces a single value: the average of all the values in the `Price` column. Because we did not use a `WHERE` clause, it uses all of the rows in the `Coffee` table in the calculation. Here is an example that calculates the average price of all the coffees having a product number that begins with “20”:

```
SELECT AVG(Price) FROM Coffee WHERE ProdNum LIKE '20%'
```

Another of the mathematical functions is `SUM`, which calculates the sum of a column’s values. The following statement, which is probably not very useful, calculates the sum of the values in the `Price` column:

```
SELECT SUM(Price) FROM Coffee
```

The `MIN` and `MAX` functions determine the minimum and maximum values found in a column. The following statement will tell us the minimum value in the `Price` column:

```
SELECT MIN(Price) FROM Coffee
```

The following statement will tell us the maximum value in the `Price` column:

```
SELECT MAX(Price) FROM Coffee
```

The `COUNT` function can be used to determine the number of rows in a table, as demonstrated by the following statement:

```
SELECT COUNT(*) FROM Coffee
```

The `*` simply indicates that you want to count entire rows. Here is another example, which tells us the number of coffees having a price greater than 9.95:

```
SELECT COUNT(*) FROM Coffee WHERE Price > 9.95
```

Queries that use math functions, such as the examples shown here, return only one value. So, when



you submit such a statement to a DBMS using JDBC, the `ResultSet` object that is returned to the program will contain one row with one column. The program shown in [Code Listing 22-6](#) shows an example of how you can use the `MIN`, `MAX`, and `AVG` functions to find the lowest, highest, and average prices in the `Coffee` table.

## Code Listing 22-6 (CoffeeMath.java)

```
1 import java.sql.*;
2
3 /**
4 * This program demonstrates some of the SQL math functions.
5 */
6
7 public class CoffeeMath
8 {
9 public static void main(String[] args)
10 {
11 // Variables to hold the lowest, highest, and
12 // average prices of coffee.
13 double lowest = 0.0,
14 highest = 0.0,
15 average = 0.0;
16
17 // Create a named constant for the URL.
18 // NOTE: This value is specific for either Java DB or Apache Derby.
19 final String DB_URL = "jdbc:derby:CoffeeDB";
20
21 try
22 {
23 // Create a connection to the database.
24 Connection conn = DriverManager.getConnection(DB_URL);
25
26 // Create a Statement object.
27 Statement stmt = conn.createStatement();
28
29 // Create SELECT statements to get the lowest, highest,
30 // and average prices from the Coffee table.
31 String minStatement = "SELECT MIN(Price) FROM Coffee";
32 String maxStatement = "SELECT MAX(Price) FROM Coffee";
33 String avgStatement = "SELECT AVG(Price) FROM Coffee";
34
35 // Get the lowest price.
36 ResultSet minResult = stmt.executeQuery(minStatement);
37 if (minResult.next())
38 lowest = minResult.getDouble(1);
39
40 // Get the highest price.
41 ResultSet maxResult = stmt.executeQuery(maxStatement);
42 if (maxResult.next())
43 highest = maxResult.getDouble(1);
44
45 // Get the average price.
46 ResultSet avgResult = stmt.executeQuery(avgStatement);
47 if (avgResult.next())
```

```

48 average = avgResult.getDouble(1);
49
50 // Display the results.
51 System.out.printf("Lowest price: $%.2f\n", lowest);
52 System.out.printf("Highest price: $%.2f\n", highest);
53 System.out.printf("Average price: $%.2f\n", average);
54
55 // Close the connection.
56 conn.close();
57 }
58 catch(Exception ex)
59 {
60 System.out.println("ERROR: " + ex.getMessage());
61 }
62 }
63 }

```

# Program Output

```

Lowest price: $6.85
Highest price: $18.45
Average price: $10.16

```

Lines 31 through 33 declare three strings: `minStatement`, `maxStatement`, and `avgStatement`. Each of these strings contains a `SELECT` statement that uses a math function.

The code in lines 36 through 38 gets the lowest price in the table. Here is a summary of how the code works:

- Line 36 executes the contents of `minStatement`, and the `ResultSet` reference that is returned is assigned to the `minResult` variable.
- The `if` statement in line 37 advances the `ResultSet` object's cursor, and line 38 gets the value of column 1 and assigns it to the `lowest` variable.

The code in lines 41 through 43 gets the highest price in the table. Here is a summary of how the code works:

- Line 41 executes the contents of `maxStatement`, and the `ResultSet` reference that is returned is assigned to the `maxResult` variable.
- The `if` statement in line 42 advances the `ResultSet` object's cursor, and line 43 gets the value of column 1 and assigns it to the `highest` variable.

The code in lines 46 through 48 gets the average price in the table. Here is a summary of how the code works:

- Line 46 executes the contents of `avgStatement`, and the `ResultSet` reference that is returned is assigned to the `avgResult` variable.

- The `if` statement in line 47 advances the `ResultSet` object's cursor, and line 48 gets the value of column 1 and assigns it to the `average` variable.



## Checkpoint

MyProgrammingLab™ [www.myprogramminglab.com](http://www.myprogramminglab.com)

1. 22.11 What is a `ResultSet` object?
2. 22.12 Look at the following SQL statement.

```
SELECT Id FROM Account
```

What is the name of the table from which this statement is retrieving data?

What is the name of the column that is being retrieved?

3. 22.13 Assume that a database has a table named `Inventory`, with the following columns:

| Column Name | Type |
|-------------|------|
|-------------|------|

|                        |                       |
|------------------------|-----------------------|
| <code>ProductID</code> | <code>CHAR(10)</code> |
|------------------------|-----------------------|

|                        |                  |
|------------------------|------------------|
| <code>QtyOnHand</code> | <code>INT</code> |
|------------------------|------------------|

|                   |                     |
|-------------------|---------------------|
| <code>Cost</code> | <code>DOUBLE</code> |
|-------------------|---------------------|

1. Write a `SELECT` statement that will return all of the columns from every row in table.
  2. Write a `SELECT` statement that will return the `ProductID` column from every row in table.
  3. Write a `SELECT` statement that will return the `ProductID` column and the `QtyOnHand` column from every row in table.
  4. Write a `SELECT` statement that will return the `ProductID` column only from the rows where `Cost` is less than 17.00.
  5. Write a `SELECT` statement that will return all of the columns from the rows where `ProductID` ends with "ZZ".
4. 22.14 What is the purpose of the `LIKE` operator?

5. 22.15 What is the purpose of the `%` symbol in a character pattern used by the `LIKE` operator? What is the purpose of the underline (`_`) character?

6. 22.16 How can you sort the results of a `SELECT` statement on a specific column?

7. 22.17 Assume that the following declarations exist:

```
final String DB_URL = "jdbc:derby:CoffeeDB";
String sql = "SELECT * FROM Coffee";
```

Write code that uses these `String` objects to get a database connection and execute the SQL statement. Be sure to close the connection when done.

8. 22.18 How do you submit a `SELECT` statement to the DBMS?

9. 22.19 Where does a `ResultSet` object's cursor initially point? How do you move the cursor forward in the result set?

10. 22.20 Assume that a valid `ResultSet` object exists, populated with data. What method do you call to retrieve column 3 as a string? What do you pass as an argument to the method?

# 22.4 Inserting Rows

CONCEPT: You use the `INSERT` statement in SQL to insert a new row into a table.

In SQL, the `INSERT` statement is used to insert a row into a database table.

```
INSERT INTO TableName VALUES (Value1, Value2, etc...)
```

In the general format, *TableName* is the name of the database table. *Value1*, *Value2*, etc . . . is a list of column values. After the statement executes, a row containing the specified column values will be inserted into the table. Here is an example that inserts a row into the `Coffee` table, in our `CoffeeDB` database:

```
INSERT INTO Coffee VALUES ('Honduran Dark', '22-001', 8.65)
```

Notice that the string values are enclosed in single-quote marks. Also, notice the order that the values appear in the list. The first value, ‘Honduran Dark’, is inserted into the first column of the table, which is `Description`. The second value, ‘22-001’, is inserted into the second column of the table, which is `ProdNum`. The third value, 8.65, is inserted into the third column of the table, which is `Price`. After this statement executes, a new row will be inserted into the `Coffee` table containing the following column values:

```
Description: 'Honduran Dark'
ProdNum: 22-001
Price: 8.65
```

If you are not sure of the order in which the columns appear in the table, you can use the following general format of the `INSERT` statement to specify the column names and their corresponding values.

```
INSERT INTO TableName
 (ColumnName1, ColumnName2, etc...)
VALUES
 (Value1, Value2, etc...)
```

In this general format *ColumnName1*, *ColumnName2*, *etc...* is a list of column names and *Value1*, *Value2*, *etc...* is a list of corresponding values. In the new row, *Value1* will appear in the column specified by *ColumnName1*, *Value2* will appear in the column specified by *ColumnName2*, and so forth. Here is an example:

```
INSERT INTO Coffee
 (Description, ProdNum, Price)
VALUES
 ('Honduran Dark', '22-001', 8.65)
```

This statement will produce a new row containing the following column values:

```
Description: 'Honduran Dark';
```

ProdNum: 22-001  
Price: 8.65

If we rewrote the `INSERT` statement in the following manner, it would produce a new row with the same values:

```
INSERT INTO Coffee
 (ProdNum, Price, Description)
VALUES
 ('22-001', 8.65, 'Honduran Dark')
```



If a column is a primary key, it must hold a unique value for each row in the table. No two rows in a table can have the same value in the primary key column. Recall that the `ProdNum` column is the primary key in the `Coffee` table. The DBMS will not allow you to insert a new row with the same product number as an existing row.

## Inserting Rows with JDBC

To issue an `INSERT` statement with JDBC, you must first get a `Statement` object from the `Connection` object, using its `createStatement` method. You then use the `Statement` object's `executeUpdate` method. The method returns an `int` value representing the number of rows that were inserted into the table. Here is an example:

```
String sqlStatement = "INSERT INTO Coffee " +
 "(ProdNum, Price, Description) " +
 "VALUES ('22-001', 8.65, 'Honduran Dark')";
int rows = stmt.executeUpdate(sqlStatement);
```

The first statement creates a string containing an `INSERT` statement. The second statement passes this string as an argument to the `executeUpdate` method. The method should return the `int` value 1, indicating that one row was inserted into the table. The program in [Code Listing 22-7](#) shows an example. It prompts the user for the description, product number, and price of a new coffee and inserts that data into the `Coffee` table.

## Code Listing 22-7 (CoffeeInserter.java)

```
1 import java.util.Scanner;
2 import java.sql.*;
3
4 /**
5 * This program lets the user insert a row into the
6 * CoffeeDB database's Coffee table.
```

```

7 */
8
9 public class CoffeeInserter
10 {
11 public static void main(String[] args)
12 {
13 String description; // To hold the coffee description
14 String prodNum; // To hold the product number
15 double price; // To hold the price
16
17 // Create a named constant for the URL.
18 // NOTE: This value is specific for either Java DB or Apache Derby.
19 final String DB_URL = "jdbc:derby:CoffeeDB";
20
21 // Create a Scanner object for keyboard input.
22 Scanner keyboard = new Scanner(System.in);
23
24 try
25 {
26 // Create a connection to the database.
27 Connection conn = DriverManager.getConnection(DB_URL);
28
29 // Get the data for the new coffee.
30 System.out.print("Enter the coffee description: ");
31 description = keyboard.nextLine();
32 System.out.print("Enter the product number: ");
33 prodNum = keyboard.nextLine();
34 System.out.print("Enter the price: ");
35 price = keyboard.nextDouble();
36
37 // Create a Statement object.
38 Statement stmt = conn.createStatement();
39
40 // Create a string with an INSERT statement.
41 String sqlStatement = "INSERT INTO Coffee " +
42 "(ProdNum, Price, Description) " +
43 "VALUES ('" +
44 prodNum + "', " +
45 price + ", '" +
46 description + "')";
47
48 // Send the statement to the DBMS.
49 int rows = stmt.executeUpdate(sqlStatement);
50
51 // Display the results.
52 System.out.println(rows + " row(s) added to the table.");
53
54 // Close the connection.
55 conn.close();
56 }
57 catch(Exception ex)
58 {
59 System.out.println("ERROR: " + ex.getMessage());
60 }
61 }
62 }

```

# Program Output

```
Enter the coffee description: Honduran Dark [Enter]
Enter the product number: 22-001 [Enter]
Enter the price: 8.65 [Enter]
1 row(s) added to the table.
```



MyProgrammingLab™ [www.myprogramminglab.com](http://www.myprogramminglab.com)

1. 22.21 Write an SQL statement to insert a new row into the `Coffee` table containing the following data:

```
Description: Eastern Blend
ProdNum: 30-001
Price: 18.95
```

2. 22.22 Rewrite the following `INSERT` statement so that it specifies the `Coffee` table's column names.

```
INSERT INTO Coffee
VALUES ('Honduran Dark', '22-001', 8.65)
```



# 22.5 Updating and Deleting Existing Rows

**CONCEPT:** You use the `UPDATE` statement in SQL to change the value of an existing row. You use the `DELETE` statement to delete rows from a table.

In SQL, the `UPDATE` statement is used to change the contents of an existing row in a table. For example, if the price of Brazilian Decaf coffee changes, we could use an `UPDATE` statement to change the `Price` column for that row. Here is the general format of the `UPDATE` statement:

```
UPDATE Table
 SET Column = Value
 WHERE Criteria
```

In the general format, *Table* is a table name, *Column* is a column name, *Value* is a value to store in the column, and *Criteria* is a conditional expression. Here is an `UPDATE` statement that will change the price of Brazilian Decaf coffee to 9.95:

```
UPDATE Coffee
 SET Price = 9.95
 WHERE Description = 'Brazilian Decaf'
```

Here is another example:

```
UPDATE Coffee
 SET Description = 'Galapagos Organic Medium'
 WHERE ProdNum = '20-001'
```

This statement locates the row where `ProdNum` is '20-001' and sets the `Description` field to 'Galapagos Organic Medium'.

It is possible to update more than one row. For example, suppose we wish to change the price of every Guatemalan coffee to 12.95. If you look back at [Table 22-1](#) you will see that the product number for each of the Guatemalan coffees begins with '21'. All we need is an `UPDATE` statement that locates all the rows where the `ProdNum` column begins with '21', and changes the `Price` column of those rows to 12.95. Here is such a statement:

```
UPDATE Coffee
 SET Price = 12.95
 WHERE ProdNum LIKE '21%'
```



Be careful that you do not leave out the `WHERE` clause and the conditional expression when using an `UPDATE` statement. You could change the contents of every row in the table! For example, look at the

following statement:

```
UPDATE Coffee
 SET Price = 4.95
```

Because this statement does not have a `WHERE` clause, it will change the `Price` column for every row in the `Coffee` table to 4.95!

## Updating Rows with JDBC

The process of issuing an `UPDATE` statement in JDBC is similar to that of issuing an `INSERT` statement. First, you get a `Statement` object from the `Connection` object, using its `createStatement` method. You then use the `Statement` object's `executeUpdate` method to issue the `UPDATE` statement. The method returns an `int` value representing the number of rows that were affected by the `UPDATE` statement. Here is an example:

```
String sqlStatement = "UPDATE Coffee " +
 "SET Price = 9.95 " +
 "WHERE Description = 'Brazilian Decaf'";
int rows = stmt.executeUpdate(sqlStatement);
```

The first statement creates a string containing an `UPDATE` statement. The second statement passes this string as an argument to the `executeUpdate` method. The method returns an `int` value indicating the number of rows that were changed.

[Code Listing 22-8](#) demonstrates how to update a row in the `Coffee` table. The user enters an existing product number, and the program displays that product's data. The user then enters a new price for the specified product, and the program updates the row with the new price.

## Code Listing 22-8 (CoffeePriceUpdater.java)

```
1 import java.util.Scanner;
2 import java.sql.*;
3
4 /**
5 * This program lets the user change the price of a
6 * coffee in the CoffeeDB database's Coffee table.
7 */
8
9 public class CoffeePriceUpdater
10 {
11 public static void main(String[] args)
12 {
13 String prodNum; // To hold the product number
14 double price; // To hold the price
15
16 // Create a named constant for the URL.
17 // NOTE: This value is specific for either Java DB or Apache Derby.
```

```

18 final String DB_URL = "jdbc:derby:CoffeeDB";
19
20 // Create a Scanner object for keyboard input.
21 Scanner keyboard = new Scanner(System.in);
22
23 try
24 {
25 // Create a connection to the database.
26 Connection conn = DriverManager.getConnection(DB_URL);
27
28 // Create a Statement object.
29 Statement stmt = conn.createStatement();
30
31 // Get the product number for the desired coffee.
32 System.out.print("Enter the product number: ");
33 prodNum = keyboard.nextLine();
34
35 // Display the coffee's current data.
36 if (findAndDisplayProduct(stmt, prodNum))
37 {
38 // Get the new price.
39 System.out.print("Enter the new price: ");
40 price = keyboard.nextDouble();
41
42 // Update the coffee with the new price.
43 updatePrice(stmt, prodNum, price);
44 }
45 else
46 {
47 // The specified product number was not found.
48 System.out.println("That product was not found.");
49 }
50
51 // Close the connection.
52 conn.close();
53 }
54 catch(Exception ex)
55 {
56 System.out.println("ERROR: " + ex.getMessage());
57 }
58 }
59
60 /**
61 The findAndDisplayProduct method finds a specified coffee's
62 data and displays it.
63 @param stmt A Statement object for the database.
64 @param prodNum The product number for the desired coffee.
65 @return true/false indicating whether the product was found.
66 */
67
68 public static boolean findAndDisplayProduct(Statement stmt,
69 String prodNum)
70 throws SQLException
71 {
72 boolean productFound; // Flag
73
74 // Create a SELECT statement to get the specified

```

```

75 // row from the Coffee table.
76 String sqlStatement =
77 "SELECT * FROM Coffee WHERE ProdNum = '" +
78 prodNum + "'";
79
80 // Send the SELECT statement to the DBMS.
81 ResultSet result = stmt.executeQuery(sqlStatement);
82
83 // Display the contents of the result set.
84 if (result.next())
85 {
86 // Display the product.
87 System.out.println("Description: " +
88 result.getString("Description"));
89 System.out.println("Product Number: " +
90 result.getString("ProdNum"));
91 System.out.println("Price: $" +
92 result.getDouble("Price"));
93
94 // Set the flag to indicate the product was found.
95 productFound = true;
96 }
97 else
98 {
99 // Indicate the product was not found.
100 productFound = false;
101 }
102
103 return productFound;
104 }
105
106 /**
107 The updatePrice method updates a specified coffee's price.
108 @param stmt A Statement object for the database.
109 @param prodNum The product number for the desired coffee.
110 @param price The product's new price.
111 */
112
113 public static void updatePrice(Statement stmt, String prodNum,
114 double price) throws SQLException
115 {
116 // Create an UPDATE statement to update the price
117 // for the specified product number.
118 String sqlStatement = "UPDATE Coffee " +
119 "SET Price = " + Double.toString(price) +
120 "WHERE ProdNum = '" + prodNum + "'";
121
122 // Send the UPDATE statement to the DBMS.
123 int rows = stmt.executeUpdate(sqlStatement);
124
125 // Display the results.
126 System.out.println(rows + " row(s) updated.");
127 }
128 }

```

# Program Output

```
Enter the product number: 17-001 [Enter]
Description: Sumatra Dark
Product Number: 17-001
Price: $7.95
Enter the new price: 9.95 [Enter]
1 row(s) updated.
```

In the `main` method, line 26 gets a connection to the database and line 29 creates a `Statement` object. Lines 32 and 33 prompt the user for a product number, which is assigned to the `prodNum` variable.

Before we let the user change the specified product's price, we want to display the product's current information. So, line 36 calls a method named `findAndDisplayProduct`, passing the `Statement` object and the `prodNum` variable as arguments. The `findAndDisplayProduct` method (which is shown in lines 68 through 104) queries the database table for the row with the specified product number. If the row is found, the method displays the row's contents and then returns `true`. If the row is not found, the method simply returns `false`.

If the specified product is found, lines 39 and 40 prompt the user for the product's new price, and the user's input is assigned to the `price` variable. Then, line 43 calls a method named `updatePrice`, passing the `Statement` object, the `prodNum` variable, and the `price` variable as arguments. The `updatePrice` method (which is shown in lines 113 through 127) updates the row containing the specified product number with the new price.

Notice that neither the `findAndDisplayProduct` method nor the `updatePrice` method handles any `SQLExceptions` that might occur. If an `SQLException` happens in either of those methods, it gets passed up to the `main` method, where it is handled by the `try-catch` statement.

## Deleting Rows with the DELETE Statement

In SQL you use the `DELETE` statement to delete one or more rows from a table. The general format of the `DELETE` statement is:

```
DELETE FROM Table WHERE Criteria
```

In the general format, *Table* is a table name and *Criteria* is a conditional expression. Here is a `DELETE` statement that will delete the row where `ProdNum` is 20-001:

```
DELETE FROM Coffee WHERE ProdNum = '20-001'
```

This statement locates the row in the `Coffee` table where the `ProdNum` column is set to the value '20-001', and deletes that row.

It is possible to delete multiple rows with the `DELETE` statement. For example, look at the following

statement:

```
DELETE FROM Coffee WHERE Description LIKE 'Sumatra%'
```

This statement will delete all rows in the `Coffee` table where the `Description` column begins with ‘Sumatra’. If you glance back at [Table 22-1](#) you will see that four rows will be deleted.



Be careful that you do not leave out the `WHERE` clause and the conditional expression when using a `DELETE` statement. You could delete every row in the table! For example, look at the following statement:

```
DELETE FROM Coffee
```

Because this statement does not have a `WHERE` clause, it will delete every row in the `Coffee` table!

## Deleting Rows with JDBC

The process of issuing a `DELETE` statement in JDBC is similar to that of issuing an `INSERT` statement or an `UPDATE` statement. First, you get a `Statement` object from the `Connection` object, using its `createStatement` method. You then use the `Statement` object’s `executeUpdate` method to issue the `DELETE` statement. The method returns an `int` value representing the number of rows that were deleted. Here is an example:

```
String sqlStatement = "DELETE FROM Coffee " +
 "WHERE ProdNum = '20-001'";
int rows = stmt.executeUpdate(sqlStatement);
```

The first statement creates a string containing a `DELETE` statement. The second statement passes this string as an argument to the `executeUpdate` method. The method returns an `int` value indicating the number of rows that were deleted.

The program shown in [Code Listing 22-9](#) demonstrates how a row can be deleted from the `Coffee` table.

## Code Listing 22-9 (CoffeeDeleter.java)

```
1 import java.util.Scanner;
2 import java.sql.*;
3
4 /**
5 This program lets the user delete a coffee
6 from the CoffeeDB database's Coffee table.
```

```

7 */
8
9 public class CoffeeDeleter
10 {
11 public static void main(String[] args)
12 {
13 String prodNum; // To hold the product number
14 String sure; // To make sure the user wants to delete
15
16 // Create a named constant for the URL.
17 // NOTE: This value is specific for either Java DB or Apache Derby.
18 final String DB_URL = "jdbc:derby:CoffeeDB";
19
20 // Create a Scanner object for keyboard input.
21 Scanner keyboard = new Scanner(System.in);
22
23 try
24 {
25 // Create a connection to the database.
26 Connection conn = DriverManager.getConnection(DB_URL);
27
28 // Create a Statement object.
29 Statement stmt = conn.createStatement();
30
31 // Get the product number for the desired coffee.
32 System.out.print("Enter the product number: ");
33 prodNum = keyboard.nextLine();
34
35 // Display the coffee's current data.
36 if (findAndDisplayProduct(stmt, prodNum))
37 {
38 // Make sure the user wants to delete this product.
39 System.out.print("Are you sure you want to delete " +
40 "this item? (y/n): ");
41 sure = keyboard.nextLine();
42
43 if (Character.toUpperCase(sure.charAt(0)) == 'Y')
44 {
45 // Delete the specified coffee.
46 deleteCoffee(stmt, prodNum);
47 }
48 else
49 {
50 System.out.println("The item was not deleted.");
51 }
52 }
53 else
54 {
55 // The specified product number was not found.
56 System.out.println("That product was not found.");
57 }
58
59 // Close the connection.
60 conn.close();
61 }
62 catch(Exception ex)
63 {
64 System.out.println("ERROR: " + ex.getMessage());

```

```

65 }
66 }
67
68 /**
69 The findAndDisplayProduct method finds a specified coffee's
70 data and displays it.
71 @param stmt A Statement object for the database.
72 @param prodNum The product number for the desired coffee.
73 @return true/false to indicate whether the product was found.
74 */
75
76 public static boolean findAndDisplayProduct(Statement stmt,
77 String prodNum)
78 throws SQLException
79 {
80 boolean productFound; // Flag
81
82 // Create a SELECT statement to get the specified
83 // row from the Coffee table.
84 String sqlStatement =
85 "SELECT * FROM Coffee WHERE ProdNum = '" +
86 prodNum + "'";
87
88 // Send the SELECT statement to the DBMS.
89 ResultSet result = stmt.executeQuery(sqlStatement);
90
91 // Display the contents of the result set.
92 if (result.next())
93 {
94 // Display the product.
95 System.out.println("Description: " +
96 result.getString("Description"));
97 System.out.println("Product Number: " +
98 result.getString("ProdNum"));
99 System.out.println("Price: $" +
100 result.getDouble("Price"));
101
102 // Set the flag to indicate the product was found.
103 productFound = true;
104 }
105 else
106 {
107 // Indicate the product was not found.
108 productFound = false;
109 }
110
111 return productFound;
112 }
113
114 /**
115 The deleteCoffee method deletes a specified coffee.
116 @param stmt A Statement object for the database.
117 @param prodNum The product number for the desired coffee.
118 */
119
120 public static void deleteCoffee(Statement stmt, String prodNum)
121 throws SQLException

```



```

122 {
123 // Create a DELETE statement to delete the
124 // specified product number.
125 String sqlStatement = "DELETE FROM Coffee " +
126 "WHERE ProdNum = '" + prodNum + "'";
127
128 // Send the DELETE statement to the DBMS.
129 int rows = stmt.executeUpdate(sqlStatement);
130
131 // Display the results.
132 System.out.println(rows + " row(s) deleted.");
133 }
134 }

```

## Program Output

```

Enter the product number: 20-001 [Enter]
Description: Galapagos Medium
Product Number: 20-001
Price: $6.85
Are you sure you want to delete this item? (y/n): y [Enter]
1 row(s) deleted.

```



MyProgrammingLab™ [www.myprogramminglab.com](http://www.myprogramminglab.com)

1. 22.23 The Midnight Coffee Roastery is running a special on decaf coffee. Write an SQL statement that changes the price of all decaf coffees to 4.95.
2. 22.24 The sale on decaf coffee didn't do too well, so the Midnight Coffee Roastery has decided to stop selling decaf. Write an SQL statement that will delete all decaf coffees from the `Coffee` table.

# 22.6 Creating and Deleting Tables

**CONCEPT:** In SQL, the `CREATE TABLE` statement can be used to create a database table. The `DROP TABLE` statement can be used to delete a table.

The `CoffeeDB` database that we have been using as our example is very simple. It has only one table, `Coffee`, which holds product information. The usefulness of this database is limited to looking up coffee descriptions, product numbers, and prices.

Suppose we want to store other data in the database, such as a list of customers. To do so, we would have to add another table to the database. In SQL you use the `CREATE TABLE` statement to create a table. Here is the general format of the `CREATE TABLE` statement:

```
CREATE TABLE TableName
 (ColumnName1 DataType1,
 ColumnName2 DataType2,
 etc...)
```

In the general format, *TableName* is the name of the table. *ColumnName1* is the name of the first column, and *DataType1* is the SQL data type for the first column. *ColumnName2* is the name of the second column, and *DataType2* is the SQL data type for the second column. This sequence repeats for all of the columns in the table. Here is an example:

```
CREATE TABLE Customer
(Name CHAR(25),
 Address CHAR(25),
 City CHAR(12),
 State CHAR(2),
 Zip CHAR(5))
```

This statement creates a new table named `Customer`. The columns in the `Customer` table are `Name`, `Address`, `City`, `State`, and `Zip`.

You may also specify that a column is a primary key by listing the `PRIMARY KEY` qualifier after the column's data type. Recall from our earlier discussion on database organization that a primary key is a column that holds a unique value for each row, and can be used to identify specific rows. When you use the `PRIMARY KEY` qualifier with a column, you should also use the `NOT NULL` qualifier. The `NOT NULL` qualifier specifies that the column must contain a value for every row. Here is an example of how we can create a `Customer` table, using the `CustomerNumber` column as the primary key:

```
CREATE TABLE Customer
(CustomerNumber CHAR(10) NOT NULL PRIMARY KEY,
 Name CHAR(25),
 Address CHAR(25),
 City CHAR(12),
 State CHAR(2),
 Zip CHAR(5))
```

This statement creates a new table named `Customer`. It has the same structure as the table created by the previous example, with one additional column, `CustomerNumber`, which is the primary key. Because `CustomerNumber` is the primary key, this column must hold a unique value for each row in the table.



Remember, a primary key is used to identify a specific row in a table. When selecting a column as a primary key, make sure it holds a unique value that cannot be duplicated for two rows in the table.

Take a look at the program in [Code Listing 22-10](#). When you run this program, it creates the `Customer` table in the `CoffeeDB` database, and then inserts the three rows shown in [Table 22-5](#).

## Table 22-5 Rows inserted into the Customer table

| CustomerNumber | Name                | Address            | City        | State | Zip   |
|----------------|---------------------|--------------------|-------------|-------|-------|
| 101            | Downtown Cafe       | 17 N. Main Street  | Asheville   | NC    | 55515 |
| 102            | Main Street Grocery | 110 E. Main Street | Canton      | NC    | 55555 |
| 103            | The Coffee Place    | 101 Center Plaza   | Waynesville | NC    | 55516 |

## Code Listing 22-10 (CreateCustomerTable.java)

```
1 import java.sql.*; // Needed for JDBC classes
2
3 /**
4 * This program creates a Customer
5 * table in the CoffeeDB database.
6 */
7
8 public class CreateCustomerTable
9 {
```

```

10 public static void main(String[] args)
11 {
12 // Create a named constant for the URL.
13 // NOTE: This value is specific for either Java DB or Apache Derby.
14 final String DB_URL = "jdbc:derby:CoffeeDB";
15
16 try
17 {
18 // Create a connection to the database.
19 Connection conn = DriverManager.getConnection(DB_URL);
20
21 // Get a Statement object.
22 Statement stmt = conn.createStatement();
23
24 // Make an SQL statement to create the table.
25 String sql = "CREATE TABLE Customer" +
26 "(CustomerNumber CHAR(10) NOT NULL PRIMARY KEY, " +
27 " Name CHAR(25), " +
28 " Address CHAR(25), " +
29 " City CHAR(12), " +
30 " State CHAR(2), " +
31 " Zip CHAR(5))";
32
33 // Execute the statement.
34 stmt.execute(sql);
35
36 // Add some rows to the new table.
37 sql = "INSERT INTO Customer VALUES" +
38 "('101', 'Downtown Cafe', '17 N. Main Street', " +
39 " 'Asheville', 'NC', '55515')";
40 stmt.executeUpdate(sql);
41
42 sql = "INSERT INTO Customer VALUES" +
43 "('102', 'Main Street Grocery', " +
44 " '110 E. Main Street', " +
45 " 'Canton', 'NC', '55555')";
46 stmt.executeUpdate(sql);
47
48 sql = "INSERT INTO Customer VALUES" +
49 "('103', 'The Coffee Place', '101 Center Plaza', " +
50 " 'Waynesville', 'NC', '55516')";
51 stmt.executeUpdate(sql);
52
53 // Close the connection.
54 conn.close();
55 }
56 catch (Exception ex)
57 {
58 System.out.println("ERROR: " + ex.getMessage());
59 }
60 }
61 }

```

# Removing a Table with the DROP TABLE

# Statement

Should the need arise to delete a table from a database, you can use the `DROP TABLE` statement. Here is the statement's general format:

```
DROP TABLE TableName
```

In the general format, *TableName* is the name of the table you wish to delete. For example, suppose that after we created the `Customer` table, we discovered that we selected the wrong data type for many of the columns. We could delete the table, and then recreate it with the proper data types. The SQL statement to delete the table would be:

```
DROP TABLE Customer
```



MyProgrammingLab™ [www.myprogramminglab.com](http://www.myprogramminglab.com)

1. 22.25 Write the SQL statement to create a table named `Book`. The `Book` table should have the columns to hold the name of the publisher, the name of the author, the number of pages, and a 10-character ISBN number.
2. 22.26 Write a statement to delete the `Book` table you created in Checkpoint 22.25.

# 22.7 Creating a New Database with JDBC

CONCEPT: Creating a new database with JDBC is as simple as adding an attribute to the database URL and then using SQL to create a table in the database.

In the previous section you learned about the `CREATE TABLE` statement, which is used to create a new table in an existing database. But, suppose you wish to create a completely new database. With JDBC, all you must do is append the attribute `;create=true` to the database URL. For example, suppose you wish to create a new database named `EntertainmentDB`, to hold data on your collection of DVDs. In either Java DB or Apache Derby, the URL you would use would be:

```
"jdbc:derby:EntertainmentDB;create=true"
```

Because we have appended the attribute `;create=true` to the database URL, the program will create the database when it runs. Then, we can use a `CREATE TABLE` statement to create a table in the database. The program in [Code Listing 22-11](#) demonstrates.

## Code Listing 22-11 (BuildEntertainmentDB.java)

```
1 import java.sql.*;
2
3 /**
4 * This program shows how to create a new database
5 * using either Java DB or Apache Derby.
6 */
7
8 public class BuildEntertainmentDB
9 {
10 public static void main(String[] args)
11 throws Exception
12 {
13 final String DB_URL =
14 "jdbc:derby:EntertainmentDB;create=true";
15
16 try
17 {
18 // Create a connection to the database.
19 Connection conn =
20 DriverManager.getConnection(DB_URL);
21
22 // Create a Statement object.
23 Statement stmt = conn.createStatement();
24
25 // Create the Dvd table.
26 System.out.println("Creating the Dvd table...");
27 stmt.execute("CREATE TABLE Dvd (" +
```

```

28 "Title CHAR(25), " +
29 "Minutes INTEGER, " +
30 "Price DOUBLE)");
31
32 // Close the resources.
33 stmt.close();
34 conn.close();
35 System.out.println("Done");
36 }
37 catch(Exception ex)
38 {
39 System.out.println("ERROR: " + ex.getMessage());
40 }
41 }
42 }

```

When this program runs, the `EntertainmentDB` database will be created. This is because the database URL, in line 14, has the `;create=true` attribute. Lines 27 through 30 then create a table named `Dvd`.

## **NOTE:**

When you create a new database using either Java DB or Apache Derby, you will see a folder appear on your system with the same name as the database. This folder holds the database. To delete the entire database, simply delete the folder.

# 22.8 Scrollable Result Sets

CONCEPT: A scrollable result set allows random cursor movement. By default, a result set is not scrollable.

By default, `ResultSet` objects allow you to move the cursor forward only. Once the cursor has moved past a row, you cannot move the cursor backward to read that row again. If you need to move the cursor backward through the result set, you can create a *scrollable result set*. You do this when you create a `Statement` object by using an overloaded version of a `Connection` object's `createStatement` method. The method accepts two arguments. The first specifies the result set's scrolling type. You can use any of the following constants for this argument:

```
ResultSet.TYPE_FORWARD_ONLY
```

This is the default scrolling type. It specifies that the result set's cursor should move forward only.

```
ResultSet.TYPE_SCROLL_INSENSITIVE
```

This specifies that the result set should be scrollable, allowing the cursor to move forward and backward through the result set. In addition, this result set is insensitive to changes made to the database. This means that if another program or process makes changes to the database, those changes will not appear in this result set.

```
ResultSet.TYPE_SCROLL_SENSITIVE
```

This specifies that the result set should be scrollable, allowing the cursor to move forward and backward through the result set. In addition, this result set is sensitive to changes made to the database. This means that if another program or process makes changes to the database, those changes will appear in this result set as soon as they are made.

The second argument specifies the result set's concurrency level. You can use any of the following constants for this argument:

```
ResultSet.CONCUR_READ_ONLY
```

This is the default concurrency level. It specifies that the result set contains a read-only version of data from the database. You cannot change the contents of the database by altering the contents of the result set.

```
ResultSet.CONCUR_UPDATABLE
```

This specifies that the result set should be updateable. Changes can be made to the result set, and then those changes can be saved to the database. The `ResultSet` interface specifies several methods that may be used to update the result set and then save those updates to the database. These methods allow you to make changes to the database without issuing SQL statements. For more information on these



methods, see the Java API documentation.

Assuming that `conn` references a `Connection` object, here is an example of the method call:

```
Statement stmt =
 conn.createStatement(ResultSet.TYPE_SCROLL_INSENSITIVE,
 ResultSet.CONCUR_READ_ONLY);
```

The `Statement` object created by this code will be scrollable, insensitive to changes made to the database by other processes, and will not be updateable.

## ResultSet Navigation Methods

Once you have created a scrollable result set, you can use the following `ResultSet` methods to move the cursor:

`first()` Moves the cursor to the first row in the result set.

`last()` Moves the cursor to the last row in the result set.

`next()` Moves the cursor to the next row in the result set.

`previous()` Moves the cursor to the previous row in the result set.

`relative(rows)` Moves the cursor the number of rows specified by the argument *rows*, relative to the current row. For example, the call `relative(2)` will move the cursor 2 rows forward from the current row, and `relative(-1)` will move the cursor 1 row backward from the current row.

`absolute(row)` Moves the cursor to the row specified by the integer *row*. Remember, row numbering begins at 1, so the call `absolute(1)` will move the cursor to the first row in the result set.

### NOTE:

Scrollable result sets are not supported by all JDBC drivers. If your driver does not support

scrollable result sets, it will throw an exception when you try to use an unsupported navigation method.

The following code shows a simple, yet practical use of some of these methods:

```
resultSet.last(); // Move to the last row
int numRows = resultSet.getRow(); // Get the current row number
resultSet.first(); // Move back to the first row
```

This code would be useful when you need to determine the number of rows in the result set before processing any of its data. The first statement moves the cursor to the last row. The second statement calls the `ResultSet` method `getRow`, which returns the row number of the current row. The third statement then moves the cursor to the first row for subsequent processing.

# 22.9 Result Set Metadata

CONCEPT: Result set metadata describes the contents of a result set. The metadata can be used to determine which columns were returned when a query that is not known in advance is executed.

The term *metadata* refers to data that describes other data. A `ResultSet` object has metadata, which describes the data stored in the `ResultSet`. You can use result set metadata to determine several things about a result set, including the number of columns it contains, the names of the columns, the types of each column, and much more. Result set metadata can be very useful if you are writing an application that will submit an SQL query, and you don't know in advance what columns will be returned.

Once you have a `ResultSet` object, you can call its `getMetaData` method to get a reference to a `ResultSetMetaData` object. Assuming that `resultSet` references a `ResultSet` object, here is an example:

```
ResultSetMetaData meta = resultSet.getMetaData();
```

`ResultSetMetaData` is an interface in the `java.sql` package. It specifies numerous methods, a few of which are described in [Table 22-6](#).

## Table 22-6 A few `ResultSetMetaData` methods

| Method                                             | Description                                                                                                                                                                                   |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>int getColumnCount()</code>                  | Returns the number of columns in the result set.                                                                                                                                              |
| <code>String getColumnName(int col)</code>         | Returns the name of the column specified by the integer <code>col</code> . The first column is column 1.                                                                                      |
| <code>String<br/>getColumnTypeName(int col)</code> | Returns the name of the data type of the column specified by the integer <code>col</code> . The first column is column 1. The data type name returned is the database-specific SQL data type. |
| <code>int</code>                                   | Returns the display width, in characters, of the column                                                                                                                                       |

`getColumnDisplaySize(int col)` specified by the integer `col`. The first column is column 1.

`String getTableName(int col)` Returns the name of the table associated with the column - specified by the integer `col`. The first column is column 1.

The program in [Code Listing 22-12](#) demonstrates how metadata can be used. It asks the user to enter a `SELECT` statement for the `CoffeeDB` database, then displays information about the result set as well as the result set's contents.

## Code Listing 22-12 (MetaDataDemo.java)

```
1 import java.sql.*;
2 import java.util.Scanner;
3
4 /**
5 * This program demonstrates result set metadata.
6 */
7
8 public class MetaDataDemo
9 {
10 public static void main(String[] args) throws Exception
11 {
12 // Create a named constant for the URL.
13 // NOTE: This value is specific for either Java DB or Apache Derby.
14 final String DB_URL = "jdbc:derby:CoffeeDB";
15
16 try
17 {
18 // Create a Scanner object for keyboard input.
19 Scanner keyboard = new Scanner(System.in);
20
21 // Get a SELECT statement from the user.
22 System.out.println("Enter a SELECT statement for " +
23 "the CoffeeDB database:");
24 String sql = keyboard.nextLine();
25
26 // Create a connection to the database.
27 Connection conn =
28 DriverManager.getConnection(DB_URL);
29
30 // Create a Statement object.
31 Statement stmt = conn.createStatement();
32
33 // Execute the query.
34 ResultSet resultSet = stmt.executeQuery(sql);
35
36 // Get the result set metadata.
37 ResultSetMetaData meta = resultSet.getMetaData();
38
39 // Display the number of columns returned.
```

```

40 System.out.println("The result set has " +
41 meta.getColumnCount() +
42 " column(s).");
43
44 // Display the column names and types.
45 for (int i = 1; i <= meta.getColumnCount(); i++)
46 {
47 System.out.println(meta.getColumnName(i) + ", " +
48 meta.getColumnTypeName(i));
49 }
50
51 // Display the contents of the rows.
52 System.out.println("\nHere are the result set rows:");
53 while (resultSet.next())
54 {
55 // Display a row.
56 for (int i = 1; i <= meta.getColumnCount(); i++)
57 {
58 System.out.print(resultSet.getString(i));
59 }
60 System.out.println();
61 }
62
63 // Close the statement and the connection.
64 stmt.close();
65 conn.close();
66 }
67 catch(Exception ex)
68 {
69 System.out.println("ERROR: " + ex.getMessage());
70 }
71 }
72 }

```

# Program Output with Example Input Shown in Bold

Enter a SELECT statement for the CoffeeDB database:

**SELECT \* FROM Coffee WHERE Price > 10.00 [Enter]**

The result set has 3 column(s).

DESCRIPTION, CHAR

PRODNUM, CHAR

PRICE, DOUBLE

Here are the result set rows:

Sumatra Organic Dark17-00411.95

Kona Medium                      18-00118.45

Kona Dark 18-002 18.45

Guatemalan Decaf 21-002 10.45

# Program Output with Example Input Shown in Bold

```
Enter a SELECT statement for the CoffeeDB database:
SELECT ProdNum FROM Coffee WHERE Price > 10.00 [Enter]
The result set has 1 column(s).
PRODNUM, CHAR
```

```
Here are the result set rows:
17-004
18-001
18-002
21-002
```

Line 34 submits the query to the DBMS and gets a `ResultSet` object. Line 37 gets a `ResultSetMetaData` object. The statement in lines 40 through 42 displays the number of columns contained in the result set. It uses the `ResultSetMetaData` object's `getColumnCount` method to get this value. The loop in lines 45 through 49 iterates once for each column in the result set. Each iteration displays the column name and column data type. The `ResultSetMetaData` object's `getColumnName` and `getColumnTypeName` methods are used to retrieve this information. The `while` loop in lines 53 through 61 displays the contents of the result set. It has a nested `for` loop, in lines 56 through 59, which iterates once for each column in the result set. Each iteration gets the column value as a string and displays it.

Now let's look at a JavaFX GUI application that we can use as an interactive utility to view the tables in the CoffeeDB database. The DBViewer application, which is shown in [Code Listing 22-13](#), displays the GUI shown in [Figure 22-16](#). The user enters a `SELECT` statement in the topmost `TextArea` control, and then clicks the Submit button. The application executes the query and displays the results in the bottom `TextArea`.




**Figure 22-16 The DBViewer application**

(Oracle Corporate Counsel)

## Code Listing 22-13 (DBViewer.java)

```
1 import javafx.application.Application;
2 import javafx.stage.Stage;
3 import javafx.scene.Scene;
4 import javafx.scene.control.TextArea;
5 import javafx.scene.control.Button;
6 import javafx.scene.layout.VBox;
7 import javafx.geometry.Insets;
8 import javafx.geometry.Pos;
9 import javafx.event.EventHandler;
10 import javafx.event.ActionEvent;
11 import java.sql.*;
12
13 public class DBViewer extends Application
14 {
15 // Database URL.
16 final String DB_URL = "jdbc:derby:CoffeeDB";
17
18 // Controls for input and output
19 TextArea queryTextArea;
20 TextArea resultsTextArea;
```

```

21
22 public static void main(String[] args)
23 {
24 // Launch the application
25 launch(args);
26 }
27
28 @Override
29 public void start(Stage primaryStage)
30 {
31 final int COL_SIZE = 50;
32 final int ROW_SIZE = 10;
33 final double SPACING = 10.0;
34
35 // Build the query entry area.
36 queryTextArea = new TextArea();
37 queryTextArea.setPrefColumnCount(COL_SIZE);
38 queryTextArea.setPrefRowCount(ROW_SIZE);
39
40 // Build the query results area.
41 resultsTextArea = new TextArea();
42 resultsTextArea.setPrefColumnCount(COL_SIZE);
43 resultsTextArea.setPrefRowCount(ROW_SIZE);
44
45 // Create the Submit button.
46 Button submitButton = new Button("Submit");
47 submitButton.setOnAction(new SubmitButtonHandler());
48
49 // Put the controls in a VBox.
50 VBox vbox = new VBox(SPACING, queryTextArea,
51 submitButton, resultsTextArea);
52 vbox.setAlignment(Pos.CENTER);
53 vbox.setPadding(new Insets(SPACING));
54
55 // Set the title.
56 primaryStage.setTitle("Enter a SELECT Query");
57
58 // Create a scene and set it to the stage.
59 Scene scene = new Scene(vbox);
60 primaryStage.setScene(scene);
61
62 // Show the stage.
63 primaryStage.show();
64 }
65
66 /**
67 * Event handler class for submitButton
68 */
69
70 class SubmitButtonHandler implements EventHandler<ActionEvent>
71 {
72 @Override
73 public void handle(ActionEvent event)
74 {
75 try
76 {
77 // Clear the previous results.

```



```

78 resultsTextArea.setText("");
79
80 // Create a connection to the database.
81 Connection conn = DriverManager.getConnection(DB_URL);
82
83 // Create a Statement object.
84 Statement stmt = conn.createStatement();
85
86 // Execute the query.
87 ResultSet resultSet =
88 stmt.executeQuery(queryTextArea.getText());
89
90 // Get the result set meta data.
91 ResultSetMetaData meta = resultSet.getMetaData();
92
93 // Create a string to hold the results
94 String output = "";
95
96 // Get the results.
97 while (resultSet.next())
98 {
99 // Get all of the columns in a row.
100 for (int i = 1; i <= meta.getColumnCount(); i++)
101 {
102 output += resultSet.getString(i);
103 output += '\t';
104 }
105 output += '\n';
106 }
107
108 // Display the results.
109 resultsTextArea.setText(output);
110
111 // Close the statement and the connection.
112 stmt.close();
113 conn.close();
114
115 }
116 catch (SQLException e)
117 {
118 e.printStackTrace();
119 System.exit(0);
120 }
121 }
122 }
123 }

```

# 22.10 Relational Data

CONCEPT: In a relational database, a column from one table can be associated with a column from other tables. This association creates a relationship between the tables.

In the [section 22.6](#) we added a `Customer` table to the `CoffeeDB` database (see [Code Listing 22-10](#)). This made the database more useful by giving us the ability to look up customer information, as well as the product information held in the `Coffee` table.

Suppose we want to make the database even more useful by storing information about unpaid customer orders. That way, we can get a list of all the customers with outstanding balances. To do this, we will need to add an additional table and more data to the database. Here is a summary of the `UnpaidOrder` table, which we will create to hold order data. (We will explain what a foreign key is momentarily.)

`UnpaidOrder`

`CustomerNumber` `CHAR(10)` *Foreign Key*

`ProdNum` `CHAR(10)` *Foreign Key*

`OrderDate` `CHAR(10)`

`Quantity` `DOUBLE`

`Cost` `DOUBLE`

The first column, `CustomerNumber`, identifies the customer that placed the order. Notice, however, that the `UnpaidOrder` table does not hold any other customer data. It holds only the customer number. When designing a database, it is important that you avoid unnecessary duplication of data. Because the customer data is already stored in the `Customer` table, we need only to store the customer number in the `Order` table. We can use that number to look up the rest of the customer's data in the `Customer` table.

In the `UnpaidOrder` table, the `CustomerNumber` column is considered a foreign key. A *foreign key* is a column in one table that references a primary key in another table. Recall that `CustomerNumber` is the primary key in the `Customer` table. When we add a row to the `UnpaidOrder` table, the value that we store in the `CustomerNumber` column must match a value in the `CustomerNumber` column of the `Customer` table. This creates a relationship between the rows in the `UnpaidOrder` table and the rows

in the `Customer` table.

The next column, `ProdNum`, is also a foreign key because it identifies a product in the `Coffee` table. This is the item that the customer ordered. Once again, it is not necessary to store all of the product data in the `UnpaidOrder` table. We need only to store the product number, and then we can use that number to look up the product data in the `Coffee` table.

The next column in the `UnpaidOrder` table is `OrderDate`. This will hold the date that the order was placed.<sup>1</sup> The `Quantity` column holds the number of pounds of coffee that the customer ordered. The `Cost` column holds the total cost of the item. We will use the following SQL statement to create the `Order` table:

<sup>1</sup>In SQL there is a `DATE` data type, which is used to hold dates. It corresponds to the `java.sql.Date` class. To keep the example simple, however, we will merely store the invoice date as a string.

```
CREATE TABLE UnpaidOrder
(CustomerNumber CHAR(10) NOT NULL REFERENCES Customer(CustomerNumber),
 ProdNum CHAR(10) NOT NULL REFERENCES Coffee(ProdNum),
 OrderDate CHAR(10),
 Quantity DOUBLE,
 Cost DOUBLE)
```

Notice that this statement introduces a new qualifier, `REFERENCES`, which is used with both the `CustomerNumber` and `ProdNum` columns. Here is the way it is used with the `CustomerNumber` column:

```
REFERENCES Customer(CustomerNumber)
```


This indicates that the column references the `CustomerNumber` column in the `Customer` table. Because of this qualifier, the DBMS performs a check when you insert a row into the `UnpaidOrder` table. It will allow you to insert a row only if the `CustomerNumber` column contains a valid value from the `CustomerNumber` column of the `Customer` table. This ensures *referential integrity* between the two tables.

The `REFERENCES` qualifier is also used with the `ProdNum` column:

```
REFERENCES Coffee(ProdNum)
```

This indicates that the column references the `ProdNum` column in the `Coffee` table. When you insert a row into the `Order` table, its `ProdNum` column must contain a valid value from the `ProdNum` column of the `Coffee` table.

System designers commonly use *entity relationship diagrams* to show the relationships between database tables. [Figure 22-17](#) shows an entity relationship diagram for the `CoffeeDB` database. In the diagram, the primary keys are denoted with (PK). The lines that are drawn between the tables show how the tables are related. In this diagram, there are two types of relationships:



## Figure 22-17 Entity relationship diagram

[Figure 22-17 Full Alternative Text](#)

- A *one to many relationship* means that for each row in table *A* there can be many rows in table *B* that reference it.
- A *many to one relationship* means that many rows in table *A* can reference a single row in table *B*.

Notice that the ends of each line show either a 1 or an infinity symbol [ $\infty$ ] You can interpret the infinity symbol as meaning *many*, and the number 1 as meaning *one*. Look at the line that connects the `Customer` table to the `UnpaidOrder` table. The 1 is at the end of the line near the `Customer` table, and the infinity symbol is at the end near the `UnpaidOrder` table. This means that one row in the `Customer` table may be referenced by many rows in the `UnpaidOrder` table. This makes sense because a customer can place many orders. (In fact, this is what management hopes for!)

If we look at the relationship in the other direction, we see that many of the rows in the `UnpaidOrder` table can reference one row in the `Customer` table. Here is a summary of all the relationships shown in the diagram:

- There is a one to many relationship between the `Customer` table and the `UnpaidOrder` table. One row in the `Customer` table may be referenced by many rows in the `UnpaidOrder` table.
- There is a many to one relationship between the `UnpaidOrder` table and the `Customer` table. Many rows in the `UnpaidOrder` table may reference a single row in the `Customer` table.
- There is a one to many relationship between the `Coffee` table and the `UnpaidOrder` table. One row in the `Coffee` table may be referenced by many rows in the `UnpaidOrder` table.
- There is a many to one relationship between the `UnpaidOrder` table and the `Coffee` table. Many rows in the `UnpaidOrder` table may reference a single row in the `Coffee` table.

## Joining Data from Multiple Tables

When related data is stored in multiple tables, as in the `CoffeeDB` database, it is often necessary to pull data from different tables in a `SELECT` statement. For example, suppose we want to see information about all the unpaid orders. Specifically, for each unpaid order, we want to see the customer number, customer name, order date, coffee description, and cost. This involves columns from the `Customer`, `UnpaidOrder`, and `Coffee` tables. Because some of these tables have columns with the same name, we have to use qualified column names in our `SELECT` statement. A *qualified column name* takes the following form:

*TableName.ColumnName*

For example, `Customer.CustomerNumber` specifies the `CustomerNumber` column in the `Customer` table, and `UnpaidOrder.CustomerNumber` specifies the `CustomerNumber` column in the `UnpaidOrder` table. Take a look at the following query:

```
SELECT
 Customer.CustomerNumber,
 Customer.Name,
 UnpaidOrder.OrderDate,
 Coffee.Description,
 UnpaidOrder.Cost
FROM
 Customer, UnpaidOrder, Coffee
WHERE
 UnpaidOrder.CustomerNumber = Customer.CustomerNumber AND
 UnpaidOrder.ProdNum = Coffee.ProdNum
```

The first part of the query specifies the columns we want:

```
SELECT
 Customer.CustomerNumber,
 Customer.Name,
 UnpaidOrder.OrderDate,
 Coffee.Description,
 UnpaidOrder.Cost
```

The second part of the query, which uses the `FROM` clause, specifies the tables we want to pull the data from:

```
FROM
 Customer, UnpaidOrder, Coffee
```

Notice that the table names are separated by commas. The third part of the query, which uses the `WHERE` clause, specifies search criteria:

```
WHERE
 UnpaidOrder.CustomerNumber = Customer.CustomerNumber AND
 UnpaidOrder.ProdNum = Coffee.ProdNum
```

The search criteria tell the DBMS how to link the rows in the tables. Recall from our earlier discussion that the `UnpaidOrder.CustomerNumber` column references the `Customer.CustomerNumber` column, and the `UnpaidOrder.ProdNum` column references the

Coffee.ProdNum column.



# WARNING!

When joining data from multiple tables, be sure to use a `WHERE` clause to specify search criteria that link the appropriate columns. Failure to do so will result in a large set of unrelated data.

## An Order Entry System

Now, let’s look at an example application that uses a relational database. In order to use this application, you will need the `Coffee` table, the `Customer` table, and the `UnpaidOrder` table in the `CoffeeDB` database. Back in [Section 22.6](#) you saw [Code Listing 22-10](#), which is a program named `CreateCustomerTable.java`. This program created the `Customer` table in the `CoffeeDB` database, and added three sample rows. If you haven’t already run that program, do so now. After running the program, you can run the `DBViewer` application presented earlier in this chapter, and enter the statement `SELECT * FROM Customer`. You should see the data shown in [Figure 22-18](#).

Enter a SELECT Query

SELECT \* FROM Customer

Submit

|     |                     |                    |             |    |       |
|-----|---------------------|--------------------|-------------|----|-------|
| 101 | Downtown Cafe       | 17 N. Main Street  | Asheville   | NC | 55515 |
| 102 | Main Street Grocery | 110 E. Main Street | Canton      | NC | 55555 |
| 103 | The Coffee Place    | 101 Center Plaza   | Waynesville | NC | 55516 |

# Figure 22-18 Customer table displayed in the DBViewer application

(Oracle Corporate Counsel)

Next, you should run the program `CreateUnpaidOrderTable.java`. This program is in the source code folder for this chapter, and it will create the `UnpaidOrder` table we discussed earlier in this chapter. The table will have no data stored in it, however.

Now that we have the necessary tables set up in our database, we will examine an order entry application that allows the user to place an order for coffee. The application is built from two classes: one to manage the database, and another to display the GUI. The first class we will look at is the `CoffeeDBManager` class, shown in [Code Listing 22-14](#). This class performs a variety of operations on the `CoffeeDB` database, which we will need in our order entry system.

## Code Listing 22-14 (`CoffeeDBManager.java`)

```
1 import java.util.ArrayList;
2 import java.sql.*;
3
4 /**
5 * The CoffeeDBManager class performs operations
6 * the CoffeeDB database.
7 */
8
9 public class CoffeeDBManager
10 {
11 // Constant for database URL.
12 public final static String DB_URL = "jdbc:derby:CoffeeDB";
13
14 /**
15 * The getCustomerNames method returns an ArrayList
16 * of Strings containing all the customer names.
17 */
18 public static ArrayList<String> getCustomerNames() throws SQLException
19 {
20 // Create a connection to the database.
21 Connection conn = DriverManager.getConnection(DB_URL);
22
23 // Create a connection to the database.
24 conn = DriverManager.getConnection(DB_URL);
25
26 // Create a Statement object for the query.
27 Statement stmt =
28 conn.createStatement(
29 ResultSet.TYPE_SCROLL_SENSITIVE,
30 ResultSet.CONCUR_READ_ONLY);
31 }
```

```

32 // Execute the query.
33 ResultSet resultSet = stmt.executeQuery(
34 "SELECT Name FROM Customer");
35
36 // Get the number of rows
37 resultSet.last(); // Move last row
38 int numRows = resultSet.getRow(); // Get row number
39 resultSet.first(); // Move to first row
40
41 // Create an ArrayList for the customer names.
42 ArrayList<String> listData = new ArrayList<>();
43
44 // Populate the ArrayList with customer names.
45 for (int index = 0; index < numRows; index++)
46 {
47 // Store the customer name in the array.
48 listData.add(resultSet.getString(1));
49
50 // Go to the next row in the result set.
51 resultSet.next();
52 }
53
54 // Close the connection and statement objects.
55 conn.close();
56 stmt.close();
57
58 // Return the listData array.
59 return listData;
60 }
61
62 /**
63 The getCoffeeNames method returns an array
64 of Strings containing all the coffee names.
65 */
66 public static ArrayList<String> getCoffeeNames() throws SQLException
67 {
68 // Create a connection to the database.
69 Connection conn = DriverManager.getConnection(DB_URL);
70
71 // Create a Statement object for the query.
72 Statement stmt =
73 conn.createStatement(
74 ResultSet.TYPE_SCROLL_SENSITIVE,
75 ResultSet.CONCUR_READ_ONLY);
76
77 // Execute the query.
78 ResultSet resultSet = stmt.executeQuery(
79 "SELECT Description FROM Coffee");
80
81 // Get the number of rows
82 resultSet.last(); // Move to last row
83 int numRows = resultSet.getRow(); // Get row number
84 resultSet.first(); // Move to first row
85
86 // Create an array for the coffee names.
87 ArrayList<String> listData = new ArrayList<>();
88

```



```

89 // Populate the array with coffee names.
90 for (int index = 0; index < numRows; index++)
91 {
92 // Store the coffee name in the array.
93 listData.add(resultSet.getString(1));
94
95 // Go to the next row in the result set.
96 resultSet.next();
97 }
98
99 // Close the connection and statement objects.
100 conn.close();
101 stmt.close();
102
103 // Return the listData array.
104 return listData;
105 }
106
107 /**
108 * The getProdNum method returns a specific
109 * coffee's product number.
110 */
111 public static String getProdNum(String coffeeName) throws SQLException
112 {
113 String prodNum = ""; // Product number
114
115 // Create a connection to the database.
116 Connection conn = DriverManager.getConnection(DB_URL);
117
118 // Create a Statement object for the query.
119 Statement stmt = conn.createStatement();
120
121 // Execute the query.
122 ResultSet resultSet = stmt.executeQuery(
123 "SELECT ProdNum " +
124 "FROM Coffee " +
125 "WHERE Description = '" +
126 coffeeName + "'");
127
128 // If the result set has a row, go to it
129 // and retrieve the product number.
130 if (resultSet.next())
131 prodNum = resultSet.getString(1);
132
133 // Close the Connection and Statement objects.
134 conn.close();
135 stmt.close();
136
137 // Return the product number.
138 return prodNum;
139 }
140
141 /**
142 * The getCoffeePrice method returns the price
143 * of a coffee.
144 */
145 public static double getCoffeePrice(String prodNum) throws SQLException

```

```

146 {
147 double price = 0.0; // Coffee price
148
149 // Create a connection to the database.
150 Connection conn = DriverManager.getConnection(DB_URL);
151
152 // Create a Statement object for the query.
153 Statement stmt = conn.createStatement();
154
155 // Execute the query.
156 ResultSet resultSet = stmt.executeQuery(
157 "SELECT Price " +
158 "FROM Coffee " +
159 "WHERE ProdNum = " +
160 prodNum + "'");
161
162 // If the result set has a row, go to it
163 // and retrieve the price.
164 if (resultSet.next())
165 price = resultSet.getDouble(1);
166
167 // Close the connection and statement objects.
168 conn.close();
169 stmt.close();
170
171 // Return the price.
172 return price;
173 }
174
175 /**
176 * The getCustomerNum method returns a specific
177 * customer's number.
178 */
179 public static String getCustomerNum(String name) throws SQLException
180 {
181 String customerNumber = "";
182
183 // Create a connection to the database.
184 Connection conn = DriverManager.getConnection(DB_URL);
185
186 // Create a Statement object for the query.
187 Statement stmt = conn.createStatement();
188
189 // Execute the query.
190 ResultSet resultSet = stmt.executeQuery(
191 "SELECT CustomerNumber " +
192 "FROM Customer " +
193 "WHERE Name = '" + name + "'");
194
195 if (resultSet.next())
196 customerNumber = resultSet.getString(1);
197
198 // Close the connection and statement objects.
199 conn.close();
200 stmt.close();
201
202 // Return the customer number.

```

```

203 return customerNumber;
204 }
205
206 /**
207 The submitOrder method submits an order to
208 the UnpaidOrder table in the CoffeeDB database.
209 */
210 public static void submitOrder(String custNum, String prodNum,
211 int quantity, double price,
212 String orderDate) throws SQLException
213 {
214 // Calculate the cost of the order.
215 double cost = quantity * price;
216
217 // Create a connection to the database.
218 Connection conn = DriverManager.getConnection(DB_URL);
219
220 // Create a Statement object for the query.
221 Statement stmt = conn.createStatement();
222
223 // Execute the query.
224 stmt.executeUpdate("INSERT INTO UnpaidOrder VALUES('" +
225 custNum + "', '" +
226 prodNum + "', '" + orderDate + "', " +
227 quantity + ", " + cost + ")");
228
229 // Close the connection and statement objects.
230 conn.close();
231 stmt.close();
232 }
233 }

```

Here is a summary of the methods in the `CoffeeDBManager` class:

- The `getCustomerNames` method, in lines 18 through 60, returns an `ArrayList` of strings containing the names of all the customers in the `Customer` table.
- The `getCoffeeNames` method, in lines 66 through 105, returns an `ArrayList` of strings containing the names of all the coffees in the `Coffee` table.
- The `getProdNum` method, in lines 111 through 139, accepts a `String` argument containing the name of a coffee. The method returns the coffee's product number.
- The `getCoffeePrice` method, in lines 145 through 173, accepts a `String` argument containing a coffee's product number. The method returns the price of the specified coffee.
- The `getCustomerNum` method, in lines 179 through 204, accepts a `String` argument containing a customer's name. The method returns that customer's customer number.
- The `submitOrder` method, in lines 210 through 232, creates a row in the `UnpaidOrder` table. It accepts arguments for the customer number, the product number of the coffee being ordered, the quantity being ordered, the coffee's price per pound, and the order date. Line 215 calculates the cost of the order by multiplying the quantity by the price per pound. Line 218 opens a connection

to the database, and line 221 creates a `Statement` object. Lines 224 through 227 execute an `INSERT` statement on the `UnpaidOrders` table.

Next, we will look at the `OrderEntrySystem` class, shown in [Code Listing 22-15](#). This class builds the GUI and interacts with the user. When you run the application, it displays the GUI shown in [Figure 22-17](#). To place an order, you select a customer and a coffee from the `ListView` controls, enter the quantity and the date in the `TextField` controls, and click the Submit button. The system then inserts a new row for the order into the `UnpaidOrder` table.

## Code Listing 22-15 (OrderEntrySystem.java)

```
1 import javafx.application.Application;
2 import javafx.stage.Stage;
3 import javafx.scene.Scene;
4 import javafx.scene.control.Label;
5 import javafx.scene.control.ListView;
6 import javafx.scene.control.TextField;
7 import javafx.scene.control.Button;
8 import javafx.scene.layout.VBox;
9 import javafx.scene.layout.HBox;
10 import javafx.scene.layout.BorderPane;
11 import javafx.geometry.Insets;
12 import javafx.geometry.Pos;
13 import java.sql.SQLException;
14 import javafx.event.EventHandler;
15 import javafx.event.ActionEvent;
16
17 public class OrderEntrySystem extends Application
18 {
19 // Fields for input
20 private ListView<String> customerListView;
21 private ListView<String> coffeeListView;
22 private TextField qtyTextField;
23 private TextField orderDateTextField;
24
25 // Constants
26 final int LV_WIDTH = 150;
27 final int LV_HEIGHT = 75;
28 final double SPACING = 10.0;
29
30 private void buildGUI(Stage stage)
31 {
32 // Build the customer selection area.
33 Label customerPrompt = new Label("Select a Customer");
34 customerListView = new ListView<String>();
35 customerListView.setPrefSize(LV_WIDTH, LV_HEIGHT);
36 VBox customerVBox = new VBox(SPACING, customerPrompt,
37 customerListView);
38 customerVBox.setAlignment(Pos.CENTER);
39 customerVBox.setPadding(new Insets(SPACING));
40
41 // Build the coffee selection area
42 Label coffeePrompt = new Label("Select a Coffee");
```

```

43 coffeeListView = new ListView<String>();
44 coffeeListView.setPrefSize(LV_WIDTH, LV_HEIGHT);
45 VBox coffeeVBox = new VBox(10, coffeePrompt,
46 coffeeListView);
47 coffeeVBox.setAlignment(Pos.CENTER);
48 coffeeVBox.setPadding(new Insets(SPACING));
49
50 // Build the quantity and date entry areas.
51 Label qtyPrompt = new Label("Quantity");
52 qtyTextField = new TextField();
53
54 Label datePrompt = new Label("Order Date");
55 orderDateTextField = new TextField();
56 VBox orderVBox = new VBox(SPACING, qtyPrompt, qtyTextField,
57 datePrompt, orderDateTextField);
58 orderVBox.setAlignment(Pos.CENTER);
59 orderVBox.setPadding(new Insets(SPACING));
60
61 // Build the button area.
62 Button submitButton = new Button("Submit");
63 Button exitButton = new Button("Exit");
64 HBox buttonHBox = new HBox(SPACING, submitButton, exitButton);
65 buttonHBox.setAlignment(Pos.CENTER);
66 buttonHBox.setPadding(new Insets(SPACING));
67
68 // Register event handlers for the Button controls.
69 submitButton.setOnAction(new SubmitButtonHandler());
70 exitButton.setOnAction(e ->
71 {
72 // Close the stage, end the application.
73 stage.close();
74 });
75
76 // Put everything inside a BorderPane.
77 BorderPane borderPane = new BorderPane();
78 borderPane.setLeft(customerVBox);
79 borderPane.setCenter(coffeeVBox);
80 borderPane.setRight(orderVBox);
81 borderPane.setBottom(buttonHBox);
82
83 // Set the title bar text.
84 stage.setTitle("Order Entry System");
85
86 // Create a scene and add it to the stage
87 Scene scene = new Scene(borderPane);
88 stage.setScene(scene);
89 }
90
91 /**
92 The loadData method loads customer names and
93 coffee names into the ListView controls.
94 */
95
96 private void loadData()
97 {
98 try
99 {
100 // Load customer names into the ListView.

```

```

101 customerListView.getItems().setAll(
102 CoffeeDBManager.getCustomerNames());
103
104 // Load coffee names into the ListView.
105 coffeeListView.getItems().setAll(
106 CoffeeDBManager.getCoffeeNames());
107 }
108 catch (SQLException e)
109 {
110 e.printStackTrace();
111 System.exit(0);
112 }
113 }
114
115 public static void main(String[] args)
116 {
117 // Launch the application
118 launch(args);
119 }
120
121 @Override
122 public void start(Stage primaryStage)
123 {
124 // Build the GUI.
125 buildGUI(primaryStage);
126
127 // Load data from the database.
128 loadData();
129
130 // Show the stage.
131 primaryStage.show();
132 }
133
134 /**
135 * Event handler class for submitButton
136 */
137
138 class SubmitButtonHandler implements EventHandler<ActionEvent>
139 {
140 @Override
141 public void handle(ActionEvent event)
142 {
143 try
144 {
145 // Get the customer name from the ListView
146 String customerName =
147 customerListView.getSelectionModel().getSelectedItem();
148
149 // Get the coffee name from the ListView
150 String coffeeName =
151 coffeeListView.getSelectionModel().getSelectedItem();
152
153 // Get the quantity
154 int qty = Integer.parseInt(qtyTextField.getText());
155
156 // Get the order date
157 String orderDate = orderDateTextField.getText();

```

```

158
159 // Get the customer number from the DB
160 String customerNum =
161 CoffeeDBManager.getCustomerNum(customerName);
162
163 // Get the coffee product number from the DB
164 String prodNum = CoffeeDBManager.getProdNum(coffeeName);
165
166 // Get the price per pound from the DB
167 double price = CoffeeDBManager.getCoffeePrice(prodNum);
168
169 // Submit the order to the DB
170 CoffeeDBManager.submitOrder(customerNum, prodNum, qty,
171 price, orderDate);
172
173 // Clear the text fields
174 qtyTextField.clear();
175 orderDateTextField.clear();
176 }
177 catch (SQLException e)
178 {
179 e.printStackTrace();
180 System.exit(0);
181 }
182 }
183 }
184 }


```

Let's take a closer look at the code.

- Lines 20 through 23 declare the `ListView` and `TextField` fields that will be used for input.
- Lines 26 through 28 declare constants for the `ListView` width and height, and for the spacing that we will place around controls in our layout containers.
- The `buildGUI` method in lines 30 through 89 creates all of the controls, places them in layout containers, and registers event handlers for the buttons. A `Scene` is created, and it is set to the stage.
- The `loadData` method, in lines 96 through 113, uses the `CoffeeDBManager` class to get the customer names and coffee names from the database and load them into the `ListView` controls.
- The `start` method, in lines 122 through 132, calls the `buildGUI` method (line 125), calls the `loadData` method (line 128), and shows the stage (line 131).
- The `SubmitButtonHandler` class, in lines 138 through 183, is the event handler for the `submitButton` control. Its `handle` method performs the following steps:
  - It gets the selected customer name from the `customerListView` control in lines 146 and 147.
  - It gets the selected coffee name from the `coffeeListView` control in lines 150 and 151.

- It gets the quantity from the `qtyTextField` control in line 154.
- It gets the order date from the `orderDateTextField` in line 157.
- At this point, we have the name of the customer placing the order and the name of the coffee being ordered, but to submit an order, we need the customer number and the product number. Lines 160 and 161 call the `CoffeeDBManager` class's `getCustomerNum` method to retrieve the customer number from the database. Line 164 calls the `CoffeeDBManager` class's `getProdNum` method to retrieve the product number. We also need the price of the coffee. Line 167 calls the `CoffeeDBManager` class's `getCoffeePrice` method to retrieve this information.
- Lines 170 and 171 call the `CoffeeDBManager` class's `submitOrder` method to submit the order. After the order is submitted, lines 174 and 175 clear the text fields holding the quantity and order date, making it easier to enter the next order.

[Figure 22-19](#) shows the order entry GUI with a customer selected, a coffee selected, a quantity entered, and an order date entered.



## Figure 22-19 Order data entered

(Oracle Corporate Counsel)

[Figure 22-19 Full Alternative Text](#)

After we submit the order shown in [Figure 22-19](#), we can run the `DBViewer` application and enter the following `SELECT` statement to pull data from various tables relating to the order. [Figure 22-20](#) shows the `DBViewer` application's screen with the `SELECT` statement already filled in and the results of the statement.






**Figure 22-20 Order information viewed in DBViewer**

(Oracle Corporate Counsel)

```
SELECT
 Customer.CustomerNumber,
 Customer.Name,
 UnpaidOrder.OrderDate,
 Coffee.Description,
 UnpaidOrder.Cost
FROM
 Customer, UnpaidOrder, Coffee
WHERE
 UnpaidOrder.CustomerNumber = Customer.CustomerNumber AND
 UnpaidOrder.ProdNum = Coffee.ProdNum
```

# 22.11 Advanced Topics

## Transactions

Sometimes an application must perform several database updates to carry out a single task. For example, suppose you have a checking account and a car loan at your bank. Each month, your car payments are automatically taken from your checking account. For this operation to take place, the balance of your checking account must be reduced by the amount of the car payment, and the balance of the car loan must also be reduced.

An operation that requires multiple database updates is known as a *transaction*. In order for a transaction to be complete, all of the steps involved in the transaction must be performed. If any single step within a transaction fails, then none of the steps in the transaction should be performed. For example, imagine that the bank system has begun the process of making your car payment. The amount of the payment is subtracted from your checking account balance, but then some sort of system failure occurs before the balance of the car loan is reduced. You would be quite upset to learn that the amount for your car payment was withdrawn from your checking account, but never applied to your loan!

Most database systems provide a means for undoing the partially completed steps in a transaction when a failure occurs. When you write transaction-processing code, there are two concepts you must understand: commit and rollback. The term *commit* refers to making a permanent change to a database, and the term *rollback* refers to undoing changes to a database.

By default, the `JDBC Connection` class operates in auto commit mode. In *auto commit* mode, all updates that are made to the database are made permanent as soon as they are executed. When auto commit mode is turned off, however, changes do not become permanent until a commit command is executed. This makes it possible to use a rollback command to undo changes. A rollback command will undo all database changes since the last commit command.

In JDBC, you turn auto commit mode off with the `Connection` class's `setAutoCommit` method, passing the argument `false`. Here is an example:

```
conn.setAutoCommit(false);
```

You execute a commit command by calling the `Connection` class's `commit` method, as shown here:

```
conn.commit();
```

A rollback command can be executed with the `Connection` class's `rollback` method, as shown here:

```
conn.rollback();
```

Let's look at an example. Suppose we add a new table named `Inventory` to the `CoffeeDB` database, for the purpose of storing the quantity of each type of coffee in inventory. The table has two rows: `ProdNum`, which is a coffee product number, and `Qty`, which is the quantity of each type of coffee. When an order is placed, we want to update both the `Inventory` table and the `UnpaidOrder` table. In the `Inventory` table we want to subtract the quantity being ordered from the quantity in inventory. In the `UnpaidOrder` table we want to insert a new row representing the order. Here is some example code that might be used to process this as a transaction.

```
Connection conn = DriverManager.getConnection(DB_URL);
conn.setAutoCommit(false);
Statement stmt = conn.createStatement();
// Attempt the transaction.
try
{
 // Update the inventory records.
 stmt.executeUpdate("UPDATE Inventory SET Qty = Qty - " +
 qtyOrdered + " WHERE ProdNum = " + prodNum);
 // Add the order to the UnpaidOrder table.
 stmt.executeUpdate("INSERT INTO UnpaidOrder VALUES('" +
 custNum + "', '" +
 prodNum + "', '" + orderDate + "', " +
 qtyOrdered + ", " + cost + ")");
 // Commit all these updates.
 conn.commit();
}
catch (SQLException ex)
{
 // Roll back the changes.
 conn.rollback();
}
```

Notice that inside the `try` block, after the statements to update the database have been executed, the `Connection` class's `commit` method is executed. In the `catch` block, the `rollback` method is executed in the event of an error.

## Stored Procedures

Many commercial database systems allow you to create SQL statements and store them in the DBMS itself. These SQL statements are called *stored procedures*, and they can be executed by other applications using the DBMS. If you have written an SQL statement that is used often in a variety of applications, it might be helpful to store it as a stored procedure in the DBMS. Then, you can call the stored procedure from any of your applications when you need to execute the SQL statement. Because stored procedures are already in the DBMS, they usually execute faster than SQL statements that are submitted from applications outside the DBMS.

We won't go into the details of stored procedures in this book, but we will point you in the right direction if you want to learn more. Each DBMS has its own syntax for creating a stored procedure in SQL, so you will have to consult your DBMS documentation to determine the format. Once you have properly written a stored procedure in SQL, you simply submit it to the DBMS using the `Statement`

class's `execute` method. To execute a stored procedure, you must create a `CallableStatement` object. `CallableStatement` is an interface in the `java.sql` package. To create a `CallableStatement` object, you call the `Connection` class's `prepareCall` statement.

## 22.12 Common Errors to Avoid

- Using the `==` operator instead of the `=` operator in an SQL statement. The equal-to operator in SQL is one `=` sign, instead of two.
- Using double quotes around strings instead of single quotes. String literals in SQL are enclosed in single quotes instead of double quotes.
- Using `&&` and `||` in an SQL statement. The logical AND and logical OR operators in SQL are the words `AND` and `OR`, not the `&&` and `||` symbols.
- Not using the correct `WHERE` clause in an `UPDATE` statement. Be careful that you do not leave out the `WHERE` clause and the conditional expression when using an `UPDATE` statement. You could change the contents of every row in the table!
- Not using the correct `WHERE` clause in a `DELETE` statement. Be careful that you do not leave out the `WHERE` clause and the conditional expression when using a `DELETE` statement. You could delete every row in the table!
- Not using the correct `WHERE` clause when joining data. When joining data from multiple tables, be sure to use a `WHERE` clause to specify search criteria that link the appropriate columns. Failure to do so will result in a large set of unrelated data.

# Review Questions and Exercises

## Multiple Choice and True/False

1. This is the technology that makes it possible for a Java application to communicate with a DBMS.
  1. DBMSC
  2. JDBC
  3. JDBMS
  4. JDSQL
2. This is a standard language for working with database management systems.
  1. Java
  2. COBOL
  3. SQL
  4. BASIC
3. The data that is stored in a table is organized in .
  1. rows
  2. files
  3. folders
  4. pages
4. The data that is stored in a row is divided into .
  1. sections
  2. bytes
  3. columns
  4. tables

5. This is a column that holds a unique value for each row, and can be used to identify specific rows.

1. ID column
2. public key
3. designator column
4. primary key

6. This type of SQL statement is used to retrieve rows from a table.

1. RETRIEVE
2. GET
3. SELECT
4. READ

7. This contains the results of an SQL `SELECT` statement.

1. select set
2. result set
3. SQL set
4. collection set

8. This clause allows you to specify search criteria with the `SELECT` statement.

1. SEARCH
2. WHERE
3. AS
4. CRITERIA

9. This is a Java class that is designed to communicate with a specific DBMS.

1. JDBC driver
2. DBMS Superclass
3. DBMS Subclass

4. Stream converter
10. This is a string listing the protocol that should be used to access a database, the name of the database, and potentially other items.
  1. JDBC driver
  2. JDBC locator
  3. Database URL
  4. Database specifier
11. This method is specified in the `Statement` interface, and should be used to execute a `SELECT` statement.
  1. `execute`
  2. `executeUpdate`
  3. `executeQuery`
  4. `executeSelect`
12. This method is specified in the `Statement` interface, and should be used to execute an `UPDATE` statement.
  1. `execute`
  2. `executeUpdate`
  3. `executeQuery`
  4. `executeSelect`
13. This method is specified in the `Statement` interface, and should be used to execute an `INSERT` statement.
  1. `execute`
  2. `executeUpdate`
  3. `executeQuery`
  4. `executeSelect`
14. This SQL statement is used to insert rows into a table.



1. INSERT
2. ADD
3. CREATE
4. UPDATE

15. This SQL statement is used to remove rows from a table.

1. REMOVE
2. ERASE
3. PURGE
4. DELETE

16. This SQL statement is used to delete an entire table.

1. REMOVE
2. DROP
3. PURGE
4. DELETE

17. This is a column in one table that references a primary key in another table.

1. secondary key
2. fake key
3. foreign key
4. duplicate key

18. **True/False:** Java comes with its own built-in DBMS.

19. **True/False:** A Java programmer that uses a DBMS to store data does not need to know about the physical structure of the data.

20. **True/False:** You use SQL instead of Java to write entire applications, including the user interface.

21. **True/False:** In SQL, the not-equal-to operator is `!=`, which is the same as in Java.

22. **True/False:** When a `ResultSet` object is initially created, its cursor is pointing at the first row in the result set.
23. **True/False:** In a transaction, it is permissible for only some of the database updates to be made.
24. **True/False:** The term *rollback* refers to undoing changes to a database.

## Find the Error

1. Find the error in the following SQL statement.

```
SELECT * FROM Coffee WHERE Description = "French Roast Dark"
```

2. Find the error in the following SQL statement.

```
SELECT * FROM Coffee WHERE ProdNum != '14-001'
```

3. Find the error in the following Java code. Assume that `conn` references a valid `Connection` object.

```
// This code has an error!!!
String sql = "SELECT * FROM Coffee";
Statement stmt = conn.createStatement();
ResultSet result = stmt.execute(sql);
```

## Algorithm Workbench

1. What SQL data types correspond with the following Java types?

- `int`
- `float`
- `String`
- `double`

2. Look at the following SQL statement.

```
SELECT Name FROM Employee
```

- What is the name of the table from which this statement is retrieving data?
- What is the name of the column that is being retrieved?

For questions 3 through 12, assume that a database has a table named `Stock`, with the following

columns:

| Column Name | Type |
|-------------|------|
|-------------|------|

|               |          |
|---------------|----------|
| TradingSymbol | CHAR(10) |
|---------------|----------|

|             |          |
|-------------|----------|
| CompanyName | CHAR(25) |
|-------------|----------|

|           |     |
|-----------|-----|
| NumShares | INT |
|-----------|-----|

|               |        |
|---------------|--------|
| PurchasePrice | DOUBLE |
|---------------|--------|

|              |        |
|--------------|--------|
| SellingPrice | DOUBLE |
|--------------|--------|

- Write a `SELECT` statement that will return all of the columns from every row in table.
- Write a `SELECT` statement that will return the `TradingSymbol` column from every row in table.
- Write a `SELECT` statement that will return the `TradingSymbol` column and the `NumShares` column from every row in table.
- Write a `SELECT` statement that will return the `TradingSymbol` column only from the rows where `PurchasePrice` is greater than 25.00.
- Write a `SELECT` statement that will return all of the columns from the rows where `TradingSymbol` starts with "SU".
- Write a `SELECT` statement that will return the `TradingSymbol` column only from the rows where `SellingPrice` is greater than `PurchasePrice`, and `NumShares` is greater than 100.
- Write a `SELECT` statement that will return the `TradingSymbol` column and the `NumShares` column only from the rows where `SellingPrice` is greater than `PurchasePrice`, and `NumShares` is greater than 100. The results should be sorted by the `NumShares` column, in ascending order.
- Write an SQL statement that will insert a new row into the `stock` table. The row should have the following column values:
  - `TradingSymbol`: XYZ
  - `CompanyName`: "XYZ Company"

- NumShares: 150
- PurchasePrice: 12.55
- SellingPrice: 22.47

11. Write an SQL statement that does the following: For each row in the `Stock` table, if the `TradingSymbol` column is “XYZ”, change it to “ABC”.
12. Write an SQL statement that will delete rows in the `Stock` table where the number of shares is less than 10.
13. Assume that the following declaration exists.

```
final String DB_URL = "jdbc:derby:CoffeeDB";
```

The string referenced by `DB_URL` is a database URL. Write a statement that uses this string to get a connection to the database.

14. Assuming that `conn` references a valid `Connection` object, write code to create a `Statement` object. (Do not be concerned about result set scrolling or concurrency.)
15. Look at the following declaration.

```
String sql = "SELECT * FROM Coffee WHERE Price < 10.00";
```

Assume also that `stmt` references a valid `Statement` object. Write code that executes the SQL statement referenced by the `sql` variable.

16. Assume that the following code is used to retrieve data from the `CoffeeDB` database’s `Coffee` table. Write the code that should appear inside the loop to display the contents of the result set.

```
String sql = "SELECT * FROM Coffee";
Connection conn = DriverManager.getConnection(DB_URL);
Statement stmt = conn.createStatement();
ResultSet result = stmt.executeQuery(sql);
while (result.next())
{
 // Finish this code!!
}
stmt.close();
conn.close();
```

17. Write an SQL statement to create a table named `Car`. The `Car` table should have the columns to hold a car’s manufacturer, year model, and a 20-character vehicle ID number.
18. Write an SQL statement to delete the `Car` table you created in Algorithm Workbench 17.

# Short Answer

1. If you are writing an application to store the customer and inventory records for a large business, why would you not want to use traditional text or binary files?
2. You hear a fellow classmate say the following: “JDBC is a standard language for working with database management systems. It was invented at IBM.” Are these statements correct, or is he confusing JDBC with something else?
3. When we speak of database organization, we speak of such things as rows, tables, and columns. Describe how the data in a database is organized into these conceptual units.
4. What is a primary key?
5. What is a result set?
6. What are the relational operators in SQL for the following comparisons?
  - Greater-than
  - Less-than
  - Greater-than or equal-to
  - Less-than or equal-to
  - Equal-to
  - Not equal-to
7. What is the number of the first row in a table? What is the number of the first column in a table?
8. What is metadata? What is result set metadata? When is result set metadata useful?
9. What is a foreign key?

# Programming Challenges

MyProgrammingLab™ Visit [www.myprogramminglab.com](http://www.myprogramminglab.com) to complete many of these Programming Challenges online and get instant feedback.

## 1. Customer Inserter



The Customer Inserter Problem

Write an application that connects to the `CoffeeDB` database, and allows the user to insert a new row into the `Customer` table.

## 2. Customer Updater

Write an application that connects to the `CoffeeDB` database, and allows the user to select a customer, then change any of that customer's information. (You should not attempt to change the customer number, because it is referenced by the `UnpaidOrder` table.)

## 3. Unpaid Order Sum

Write an application that connects to the `CoffeeDB` database, then calculates and displays the total amount owed in unpaid orders. This will be the sum of each row's `Cost` column.

## 4. Population Database

Make sure you have downloaded the book's source code from the companion Web site at [www.pearson.com/gaddis](http://www.pearson.com/gaddis). In this chapter's source code folder you will find a program named `CreateCityDB.java`. Compile and run the program. The program will create a Java DB or Apache Derby database named `CityDB`. The `CityDB` database will have a table named `City`, with the following columns:

| Column Name | Data Type |
|-------------|-----------|
|-------------|-----------|

`CityName`      `CHAR (50)`

*Primary key*

`Population`    `DOUBLE`

The `CityName` column stores the name of a city and the `Population` column stores the population of that city. After you run the `CreateCityDB.java` program, the `City` table will contain 20 rows with various cities and their populations.

Next, write a program that connects to the `CityDB` database, and allows the user to select any of the following operations:

- Sort the list of cities by population, in ascending order.
- Sort the list of cities by population, in descending order.
- Sort the list of cities by name.
- Get the total population of all the cities.
- Get the average population of all the cities.
- Get the highest population.
- Get the lowest population.

## 5. Personnel Database Creator

Write an application that creates a database named `Personnel`. The database should have a table named `Employee`, with columns for employee ID, name, position, and hourly pay rate. The employee ID should be the primary key. Insert at least five sample rows of data into the `Employee` table.

## 6. Employee Insertter

Write a GUI application that allows the user to add new employees to the `Personnel` database you created in Programming Challenge 5.

## 7. Employee Updater

Write a GUI application that allows the user to look up an employee in the `Personnel` database you created in Programming Challenge 5. The user should be able to change any of the employee's information except employee ID, which is the primary key.



## 8. PhoneBook Database

Write an application that creates a database named `PhoneBook`. The database should have a table named `Entries`, with columns for a person's name and phone number.

Next, write an application that lets the user add rows to the `Entries` table, look up a person's phone number, change a person's phone number, and delete specified rows.

# Chapter 23 A First Look at GUI Applications with Swing

## TOPICS

1. [23.1 Introduction](#)
2. [23.2 Creating Windows](#)
3. [23.3 Layout Managers](#)
4. [23.4 Radio Buttons and Check Boxes](#)
5. [23.5 Borders](#)
6. [23.6 Focus on Problem Solving: Extending Classes from `JPanel`](#)
7. [23.7 Splash Screens](#)
8. [23.8 Using Console Output to Debug a GUI Application](#)
9. [23.9 Common Errors to Avoid](#)



This chapter discusses GUI development using the Swing classes. Oracle has announced that JavaFX is replacing Swing as the standard GUI library for Java. Swing will remain part of the Java API for the foreseeable future, however, so we are providing this chapter for you to use as you make the transition from Swing to JavaFX. To learn about JavaFX, see Chapters 12, 13, and 14.

# 23.1 Introduction

CONCEPT: In Java, you use the Java Foundation Classes (JFC) to create a graphical user interface for your application. Within the JFC you use the Abstract Windowing Toolkit (AWT) or Swing classes to create a graphical user interface.

In this chapter, we discuss the basics of creating a Java application with a *graphical user interface* or *GUI* (pronounced “gooey”). A GUI is a graphical window or a system of graphical windows that is presented by an application for interaction with the user. In addition to accepting input from the keyboard, GUIs typically accept input from a mouse as well.

A window in a GUI commonly consists of several *components* that present data to the user and/or allow interaction with the application. Some of the common GUI components are buttons, labels, text fields, check boxes, and radio buttons. [Figure 23-1](#) shows an example of a window with a variety of components. [Table 23-1](#) describes the components that appear in the window.



Figure 23-1 Various GUI components

(Oracle Corporate Counsel)

[Figure 23-1 Full Alternative Text](#)

Table 23-1 Some GUI components

| Component | Description |
|-----------|-------------|
|-----------|-------------|

|              |                                                                                                                                                                                                                                                 |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Label        | An area that can display text.                                                                                                                                                                                                                  |
| Text field   | An area in which the user may type a single line of input from the keyboard.                                                                                                                                                                    |
| Combo box    | A component that displays a drop-down list of items from which the user may select. A combo box also provides a text field in which the user may type input. It is called a combo box because it is the combination of a list and a text field. |
| Check box    | A component that has a box that may be checked or unchecked.                                                                                                                                                                                    |
| List         | A list from which the user may select an item.                                                                                                                                                                                                  |
| Radio button | A component that can be either selected or deselected. Radio buttons usually appear in groups and allow the user to select one of several options.                                                                                              |
| Slider       | A component that allows the user to select a value by moving a slider along a track.                                                                                                                                                            |
| Button       | A button that can cause an action to occur when it is clicked.                                                                                                                                                                                  |

## The JFC, AWT, and Swing

Java programmers use the *Java Foundation Classes (JFC)* to create GUI applications. The JFC consists of several sets of classes, many of which are beyond the scope of this book. The two sets of JFC classes that we focus on are the AWT and Swing classes. First, we discuss the differences between them.

Java has been equipped, since its earliest version, with a set of classes for drawing graphics and creating GUIs. These classes are part of the *Abstract Windowing Toolkit (AWT)*. The AWT allows programmers to create applications and applets that interact with the user via windows and other GUI components.

Programmers are limited in what they can do with the AWT classes, however. This is because the AWT classes do not actually draw user interface components on the screen. Instead, the AWT classes communicate with another layer of software, known as the *peer classes*, which directs the underlying

operating system to draw its own built-in components. Each version of Java that is developed for a particular operating system has its own set of peer classes. Although this means that Java programs have a look that is consistent with other applications on the same system, it also leads to some problems.

One problem is that not all operating systems offer the same set of GUI components. For example, one operating system might provide a sophisticated slider bar component that is not found on any other platform. Other operating systems might have their own unique components as well. In order for the AWT to retain its portability, it has to offer only those components that are common to all the operating systems that support Java.

Another problem is in the behavior of components across various operating systems. A component on one operating system might have slightly different behavior than the same component on a different operating system. In addition, the peer classes for some operating systems reportedly have bugs. As a result, programmers cannot be completely sure how their AWT programs will behave on different operating systems until they test each one.

A third problem is that programmers cannot easily customize the AWT components. Because these components rely on the appearance and behavior of the underlying operating system components, there is little that can be done by the programmer to change their properties.

To remedy these problems, Swing was introduced with the release of Java 2. *Swing* is a library of classes that do not replace the AWT, but provide an improved alternative for creating GUI applications and applets. Very few of the Swing classes rely on an underlying system of peer classes. Instead, Swing draws most of its own components on the screen. This means that Swing components can have a consistent look and predictable behavior on any operating system.



## NOTE:

Swing applications can have the look of a specific operating system. The programmer may choose from a variety of “look and feel” themes.

Swing components can also be easily customized. The Swing library provides many sophisticated components that are not found in the AWT. In this chapter and in [Chapter 24](#), we primarily use Swing to develop GUI applications. In [Chapter 25](#), we use AWT to develop applets.



## NOTE:

AWT components are commonly called heavyweight components because they are coupled with their underlying peer classes. Very few of the Swing components are coupled with peer classes, so they are referred to as lightweight components.

# Event-Driven Programming

Programs that operate in a GUI environment must be *event-driven*. An *event* is an action that takes place within a program, such as the clicking of a button. Part of writing a GUI application is creating event listeners. An *event listener* is an object that automatically executes one of its methods when a specific event occurs. If you wish for an application to perform an operation when a particular event occurs, you must create an event listener object that responds when that event takes place.

## The `javax.swing` and `java.awt` Packages

In this chapter, we use the Swing classes for all of the graphical components that we create in our GUIs. The Swing classes are part of the `javax.swing` package. (Take note of the letter `x` that appears after the word `java`.) The following `import` statement will be used in every applicaton:

```
import javax.swing.*;
```

We also use some of the AWT classes to determine when events, such as the clicking of a mouse, take place in our applications. The AWT classes are part of the `java.awt` package. (Note that there is no `x` after `java` in this package name.) Programs that use the AWT classes will have the following `import` statement:

```
import java.awt.*;
```

## 23.2 Creating Windows

CONCEPT: You can use Swing classes to create windows containing various GUI components.

The `JOptionPane` dialog boxes that you learned about in [Chapter 2](#) allow you to easily display messages and gather input. If an application is to provide a full graphical user interface, however, much more is needed. Often, applications need one or more windows with various components that allow the user to enter and/or select data and interact with the application. For example, the window that is displayed in [Figure 23-1](#) has several different components within it.



### Creating a Simple GUI Application

A window is a component, but because a window contains other components, it is more appropriately considered a container. A *container* is simply a component that holds other components. In GUI terminology, a container that can be displayed as a window is known as a *frame*. A frame appears as a basic window that has a border around it, a title bar, and a set of buttons for minimizing, maximizing, and closing the window. In a Swing application, you create a frame object from the `JFrame` class.

There are a number of steps involved in creating a window, so let's look at an example. The program in [Code Listing 23-1](#) displays the window shown in [Figure 23-2](#).




**Figure 23-2 Window displayed by  
`ShowWindow.java`**

# Code Listing 23-1 (ShowWindow.java)

```
1 import javax.swing.*; // Needed for Swing classes
2
3 /**
4 * This program displays a simple window with a title. The
5 * application exits when the user clicks the close button.
6 */
7
8 public class ShowWindow
9 {
10 public static void main(String[] args)
11 {
12 final int WINDOW_WIDTH = 350; // Window width in pixels
13 final int WINDOW_HEIGHT = 250; // Window height in pixels
14
15 // Create a window.
16 JFrame window = new JFrame();
17
18 // Set the title.
19 window.setTitle("A Simple Window");
20
21 // Set the size of the window.
22 window.setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
23
24 // Specify what happens when the close button is clicked.
25 window.setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
26
27 // Display the window.
28 window.setVisible(true);
29 }
30 }
```

The window shown in [Figure 23-2](#) was produced on a system running Microsoft Windows. Notice that the window has a border and a title bar with “A Simple Window” displayed in it. In addition, it has the standard Microsoft Windows buttons in the upper-right corner: a minimize button, a maximize button, and a close button. These standard features are sometimes referred to as *decorations*. If you run this program, you will see the window displayed on your screen. When you click on the close button, the window disappears and the program terminates.

Let’s take a closer look at the code. First, notice that the following `import` statement is used in line 1:

```
import javax.swing.*; // Needed for Swing classes
```

Any program that uses a Swing class, such as `JFrame`, must have this `import` statement. In lines 12 and 13 the two constants `WINDOW_WIDTH` and `WINDOW_HEIGHT` are declared as follows:

```
final int WINDOW_WIDTH = 350; // Window width in pixels
final int WINDOW_HEIGHT = 250; // Window height in pixels
```

We use these constants later in the program to set the size of the window. The window’s size is measured in pixels. A *pixel* is one of the small dots that make up a screen display; the resolution of



your monitor is measured in pixels. For example, if your monitor's resolution is 1024 by 768, that means the width of your screen is 1024 pixels, and the height of your screen is 768 pixels.

Next, we create an instance of the `JFrame` class with the following statement in line 16:

```
JFrame window = new JFrame();
```

This statement creates a `JFrame` object in memory and assigns its address to the `window` variable. This statement does not display the window on the screen, however. A `JFrame` is initially invisible.

In line 19 we call the `JFrame` object's `setTitle` method as follows:

```
window.setTitle("A Simple Window");
```

The string that is passed as an argument to `setTitle` will appear in the window's title bar when it is displayed. In line 22 we call the `JFrame` object's `setSize` method to set the window's size as follows:

```
window.setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
```

The two arguments passed to `setSize` specify the window's width and height in pixels. In this program we pass the constants `WINDOW_WIDTH` and `WINDOW_HEIGHT`, which we declared earlier, to set the size of the window to 350 pixels by 250 pixels.

In line 25 we specify the action that we wish to take place when the user clicks on the close button, which appears in the upper-right corner of the window as follows:

```
window.setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
```

There are a number of actions that can take place when the user clicks on the close button. The `setDefaultCloseOperation` method takes an `int` argument, which specifies the action. In this statement, we pass the constant `JFrame.EXIT_ON_CLOSE`, which causes the application to end with a `System.exit` method call. If we had passed `JFrame.HIDE_ON_CLOSE`, the window would be hidden from view, but the application would not end. The default action is `JFrame.HIDE_ON_CLOSE`.

Last, in line 28, we use the following code to display the window:

```
window.setVisible(true);
```

The `setVisible` method takes a `boolean` argument. If the argument is `true`, the window is made visible. If the argument is `false`, the window is hidden.

## Using Inheritance to Extend the `JFrame` Class

The program in [Code Listing 23-1](#) performs a very simple operation: It creates an instance of the `JFrame` class and displays it. Most of the time, your GUI applications will be much more involved

than this. As you progress through this chapter, you will add numerous components and capabilities to the windows that you create.

Instead of simply creating an instance of the `JFrame` class, as shown in [Code Listing 23-1](#), a more common technique is to use inheritance to create a new class that extends the `JFrame` class.

## IF YOU'VE SKIPPED AHEAD TO THIS CHAPTER:

This chapter is written so that you can skip ahead to it any time after [Chapter 6](#). Reading about inheritance and interfaces in [Chapter 10](#) would be helpful; but if you have not read that material yet, the following summarizes what you need to know for this chapter.

When a new class *extends* an existing class, it inherits many of the existing class's members just as if they were part of the new class. For example, you saw how the program in [Code Listing 23-1](#) created a `JFrame` object and then called four of its methods: `setTitle`, `setSize`, `setDefaultCloseOperation`, and `setVisible`. These methods are all members of the `JFrame` class. If you create a new class that extends the `JFrame` class, the new class will automatically inherit these methods. Then these methods can be called from an instance of the new class just as if they were written into its declaration. You can add your own custom code to the new class, making it a specialized, or extended, version of the `JFrame` class. Programs can then create instances of your new specialized class instead of the more generic `JFrame` class.

Let's look at the `SimpleWindow` class in [Code Listing 23-2](#). This is an example of a class that extends the `JFrame` class.

### Code Listing 23-2 (`SimpleWindow.java`)

```
1 import javax.swing.*; // Needed for Swing classes
2
3 /**
4 This class extends the JFrame class. Its constructor displays
5 a simple window with a title. The application exits when the
6 user clicks the close button.
7 */
8
9 public class SimpleWindow extends JFrame
10 {
11 /**
12 Constructor
13 */
14
15 public SimpleWindow()
16 {
```

```

17 final int WINDOW_WIDTH = 350; // Window width in pixels
18 final int WINDOW_HEIGHT = 250; // Window height in pixels
19
20 // Set this window's title.
21 setTitle("A Simple Window");
22
23 // Set the size of this window.
24 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
25
26 // Specify what happens when the close button is clicked.
27 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
28
29 // Display the window.
30 setVisible(true);
31 }
32 }

```

Notice the class header in line 9 as follows:

```
public class SimpleWindow extends JFrame
```

The words `extends JFrame` indicate that the `SimpleWindow` class extends the `JFrame` class. This means that the `SimpleWindow` class inherits members of the `JFrame` class, such as the `setTitle`, `setSize`, `setDefaultCloseOperation`, and `setVisible` methods, just as if they were written into the `SimpleWindow` class declaration. Now look at the constructor. In lines 17 and 18 we declare the `WINDOW_WIDTH` and `WINDOW_HEIGHT` constants, which will be used to establish the size of the window as follows:

```
final int WINDOW_WIDTH = 350; // Window width in pixels
final int WINDOW_HEIGHT = 250; // Window height in pixels
```

In line 21 we call the `setTitle` method to set the text for the window's title bar as follows:

```
setTitle("A Simple Window");
```

Notice that we are calling the method without an object reference and a dot preceding it. This is because the method was inherited from the `JFrame` class, and we can call it just as if it were written into the `SimpleWindow` class declaration.

The rest of the constructor calls the `setSize`, `setDefaultCloseOperation`, and `setVisible` methods. All that is necessary to display the window is to create an instance of the `SimpleWindow` class, as shown in the program in [Code Listing 23-3](#). When this program runs, the window that was previously shown in [Figure 23-2](#) is displayed. Remember, the `SimpleWindow` class is an extended version of the `JFrame` class. When we create an instance of the `SimpleWindow` class, we are really creating an instance of the `JFrame` class, with some customized code added to its constructor.

## Code Listing 23-3 (SimpleWindowDemo.java)

```
1 /**
```

```

2 This program creates an instance of the
3 SimpleWindow class.
4 */
5
6 public class SimpleWindowDemo
7 {
8 public static void main(String[] args)
9 {
10 SimpleWindow myWindow = new SimpleWindow();
11 }
12 }

```

# Equipping GUI Classes with a main Method

You know that a Java application always starts execution with a static method named `main`. The previous example consists of two separate files:

- `SimpleWindow.java`: This file contains the `SimpleWindow` class, which defines a GUI window.
- `SimpleWindowDemo.java`: This file contains a static `main` method that creates an object of the GUI window class, thus displaying it.

The purpose of the `SimpleWindowDemo.java` file is simply to create an instance of the `SimpleWindow` class. It is possible to eliminate the second file, `SimpleWindowDemo.java`, by writing the static `main` method directly into the `SimpleWindow.java` file. The `EmbeddedMain` class in [Code Listing 23-4](#) shows an example.

## Code Listing 23-4 (EmbeddedMain.java)

```

1 import javax.swing.*; // Needed for Swing classes
2
3 /**
4 * This class defines a GUI window and has its own
5 * main method.
6 */
7
8 public class EmbeddedMain extends JFrame
9 {
10 final int WINDOW_WIDTH = 350; // Window width in pixels
11 final int WINDOW_HEIGHT = 250; // Window height in pixels
12
13 /**
14 * Constructor
15 */
16
17 public EmbeddedMain()
18 {
19 // Set this window's title.
20 setTitle("A Simple Window");


```

```

21
22 // Set the size of this window.
23 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
24
25 // Specify what happens when the close button is clicked.
26 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
27
28 // Display the window.
29 setVisible(true);
30 }
31
32 /**
33 The main method creates an instance of the EmbeddedMain
34 class, which causes it to display its window.
35 */
36
37 public static void main(String[] args)
38 {
39 EmbeddedMain em = new EmbeddedMain();
40 }
41 }

```

The `EmbeddedMain` class contains its own static `main` method (in lines 37 through 40), which creates an instance of the class. Notice that the `main` method has exactly the same header as any other static `main` method that we have written. We can compile the *EmbeddedMain.java* file and then run the resulting *.class* file. When we do, we see the window shown in [Figure 23-3](#).



## Figure 23-3 Window displayed by the `EmbeddedMain` class

(Oracle Corporate Counsel)

Notice that in line 39 the `main` method declares a variable named `em` to reference the instance of the class. Once the instance is created, however, the variable is not used again. Because we do not need the variable, we can instantiate the class *anonymously* as shown here:

```
public static void main(String[] args)
```

```
{
 new EmbeddedMain();
}
```

In this version of the method, an instance of the `EmbeddedMain` class is created in memory, but its address is not assigned to any reference variable.

# Adding Components to a Window

Swing provides numerous GUI components that can be added to a window. Three fundamental components are the label, the text field, and the button. These are summarized in [Table 23-2](#).


## Table 23-2 Label, text field, and button controls

| Component  |  | Swing Class             | Description                                                                 |
|------------|--|-------------------------|-----------------------------------------------------------------------------|
| Label      |  | <code>JLabel</code>     | An area that can display text                                               |
| Text field |  | <code>JTextField</code> | An area in which the user may type a single line of input from the keyboard |
| Button     |  | <code>JButton</code>    | A button that can cause an action to occur when it is clicked               |

In Swing, labels are created with the `JLabel` class, text fields are created with the `JTextField` class, and buttons are created with the `JButton` class. To demonstrate these components, we will build a simple GUI application: The Kilometer Converter. This application will present a window in which the user will be able to enter a distance in kilometers, and then click a button to see that distance converted to miles. The conversion formula is as follows:

$$Miles = Kilometers \times 0.6214$$

When designing a GUI application, it is usually helpful to draw a sketch showing the window you are creating. [Figure 23-4](#) shows a sketch of what the Kilometer Converter application’s window will look like. As you can see from the sketch, the window will have a label, a text field, and a button. When the user clicks the button, the distance in miles will be displayed in a separate `JOptionPane` dialog box.



## Figure 23-4 Sketch of the Kilometer Converter window


(Oracle Corporate Counsel)

[Figure 23-4 Full Alternative Text](#)

## Content Panes and Panels

Before we start writing code, you should be familiar with content panes and panels. A *content pane* is a container that is part of every `JFrame` object. You cannot see the content pane and it does not have a border, but any component that is to be displayed in a `JFrame` must be added to its content pane.

A *panel* is also a container that can hold GUI components. Unlike `JFrame` objects, panels cannot be displayed by themselves; however, they are commonly used to hold and organize collections of related components. With Swing, you create panels with the `JPanel` class. In our Kilometer Converter application, we will create a panel to hold the label, text field, and button. Then we will add the panel to the `JFrame` object's content pane. This is illustrated in [Figure 23-5](#).



# Figure 23-5 A panel is added to the content pane

(Oracle Corporate Counsel)

[Figure 23-5 Full Alternative Text](#)

[Code Listing 23-5](#) shows the initial code for the `KiloConverter` class. We will be adding to this code as we develop the application. This version of the class is stored in the source code folder [Chapter 23\KiloConverter Phase 1](#).

## Code Listing 23-5 (`KiloConverter.java`)

```
1 import javax.swing.*;
2
3 /**
4 * The KiloConverter class displays a JFrame that
5 * lets the user enter a distance in kilometers. When
6 * the Calculate button is clicked, a dialog box is
7 * displayed with the distance converted to miles.
8 */
9
10 public class KiloConverter extends JFrame
11 {
12 private JPanel panel; // To reference a panel
13 private JLabel messageLabel; // To reference a label
14 private JTextField kiloTextField; // To reference a text field
15 private JButton calcButton; // To reference a button
16 private final int WINDOW_WIDTH = 310; // Window width
17 private final int WINDOW_HEIGHT = 100; // Window height
18
19 /**
20 * Constructor
21 */
22
23 public KiloConverter()
24 {
25 // Set the window title.
26 setTitle("Kilometer Converter");
27
28 // Set the size of the window.
29 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
30
31 // Specify what happens when the close button is clicked.
32 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
33
34 // Build the panel and add it to the frame.
35 buildPanel();
36 }
37 }
```



```

37 // Add the panel to the frame's content pane.
38 add(panel);
39
40 // Display the window.
41 setVisible(true);
42 }
43
44 /**
45 * The buildPanel method adds a label, a text field,
46 * and a button to a panel.
47 */
48
49 private void buildPanel()
50 {
51 // Create a label to display instructions.
52 messageLabel = new JLabel("Enter a distance " +
53 "in kilometers");
54
55 // Create a text field 10 characters wide.
56 kiloTextField = new JTextField(10);
57
58 // Create a button with the caption "Calculate".
59 calcButton = new JButton("Calculate");
60
61 // Create a JPanel object and let the panel
62 // field reference it.
63 panel = new JPanel();
64
65 // Add the label, text field, and button
66 // components to the panel.
67 panel.add(messageLabel);
68 panel.add(kiloTextField);
69 panel.add(calcButton);
70 }
71
72 /**
73 * main method
74 */
75
76 public static void main(String[] args)
77 {
78 new KiloConverter();
79 }
80 }

```

Let's take a closer look at this class. First, notice in line 10 that the `KiloConverter` class extends the `JFrame` class as follows:

```
public class KiloConverter extends JFrame
```

Next, in lines 12 through 17, notice in the following that the class declares a number of fields, and according to good class design principles, the fields are private:

```

private JPanel panel; // To reference a panel
private JLabel messageLabel; // To reference a label
private JTextField kiloTextField; // To reference a text field

```

```
private JButton calcButton; // To reference a button
private final int WINDOW_WIDTH = 310; // Window width
private final int WINDOW_HEIGHT = 100; // Window height
```

The statement in line 12 declares a `JPanel` reference variable named `panel`, which we will use to reference the panel that will hold the other components. The `messageLabel` variable, declared in line 13, will reference a `JLabel` object that displays a message instructing the user to enter a distance in kilometers. The `kiloTextField` variable, declared in line 14, will reference a `JTextField` object that will hold a value typed by the user. The `calcButton` variable, declared in line 15, will reference a `JButton` object that will calculate and display the kilometers converted to miles when clicked. The `WINDOW_WIDTH` and `WINDOW_HEIGHT` fields, declared in lines 16 and 17, are constants that hold the width and height of the window.

Now let's look at the constructor. In line 26 the `setTitle` method, which was inherited from the `JFrame` class, is called to set the text for the window's title bar. Next, in line 29, the inherited `setSize` method is called to establish the size of the window. In line 32, the inherited `setDefaultCloseOperation` method is called to establish the action that should occur when the window's close button is clicked.

Line 35 calls the `buildPanel` method. The `buildPanel` method is defined in this class, in lines 49 through 70. The purpose of the `buildPanel` method is to create a label, a text field, and a button, and then add those components to a panel. Let's look at the method.

First, look at the method header in line 49 and notice that it is declared private. When a method is private, only other methods in the same class can call it. This method is not meant to be called by code outside the class, so it is declared private. In lines 52 and 53, the method uses the following statement to create a `JLabel` object and assign its address to the `message` field:

```
messageLabel = new JLabel("Enter a distance " +
 "in kilometers");
```

The string that is passed to the `JLabel` constructor is the text that will be displayed in the label. The following statement appears in line 56. It creates a `JTextField` object, and assigns its address to the `kiloTextField` field:

```
kiloTextField = new JTextField(10);
```

The argument that is passed to the `JTextField` constructor is the width of the text field in columns. One column is enough space to hold the letter “m,” which is the widest letter in the alphabet.

The following statement appears in line 59; it creates a `JButton` object, and assigns its address to the `calcButton` field:

```
calcButton = new JButton("Calculate");
```

The string that is passed as an argument to the `JButton` constructor is the text that will be displayed on the button.

Next, in line 63, the method uses the following statement to create a `JPanel` object and assign its address to the `panel` field, which is a private field in the class:

```
panel = new JPanel();
```

A `JPanel` object is used to hold other components. You add a component to a `JPanel` object with the `add` method. The following code, in lines 67 through 69, adds the objects referenced by the `messageLabel`, `kiloTextField`, and `calcButton` variables to the `JPanel` object:

```
panel.add(messageLabel);
panel.add(kiloTextField);
panel.add(calcButton);
```

At this point, the panel is fully constructed in memory. The `buildPanel` method ends, and control returns to the class constructor. Here's the next statement in the constructor, which appears in line 38:


```
add(panel);
```

This statement calls the `add` method, which was inherited from the `JFrame` class. The purpose of the `add` method is to add an object to the content pane. This statement adds the object referenced by `panel` to the content pane.

The constructor's last statement, in line 41, calls the inherited `setVisible` method to display the window on the screen as follows:

```
setVisible(true);
```


The class has a static `main` method, which appears in lines 76 through 79. Line 78 creates an instance of the `KiloConverter` class. When this program is executed, the window shown in [Figure 23-6](#) is displayed on the screen.



## Figure 23-6 Kilometer Converter window

(Oracle Corporate Counsel)

[Figure 23-7](#) shows the window again, this time pointing out each of the components.



## Figure 23-7 Components in the Kilometer Converter window

(Oracle Corporate Counsel)

[Figure 23-7 Full Alternative Text](#)

Although you can type input into the text field, the application does nothing when you click the Calculate button because we have not written an event handler that will execute when the button is clicked. That's the next step.

### NOTE:

Recall that the size of the window in the `KiloConverter` class is set to 310 pixels wide by 100 pixels high. This is set with the `WINDOW_WIDTH` and `WINDOW_HEIGHT` constants. [Figures 23-6](#) and [23-7](#) show the window as it appears on a system set at a video resolution of 1024 by 768 pixels. If your video resolution is lower, the window might not appear exactly as shown in the figures. If this is the case, you can increase the values of the `WINDOW_WIDTH` and `WINDOW_HEIGHT` constants and recompile the program. This is true for other applications in this chapter as well.

## Handling Events with Action Listeners



### Handling Events

An *event* is an action that takes place within a program, such as the clicking of a button. When an event takes place, the component that is responsible for the event creates an event object in memory. The *event object* contains information about the event. The component that generated the event object is known as the *event source*. For example, when the user clicks a button, the `JButton` component

generates an event object. The `JButton` component that generated the event object is the event source.

But what happens to the event object once it is generated by a source component? It is possible that the source component is connected to one or more event listeners. An *event listener* is an object that responds to events. If the source component is connected to an event listener, then the event object is automatically passed, as an argument, to a specific method in the event listener. The method then performs any actions that it was programmed to perform in response to the event. This process is sometimes referred to as *event firing*.

When you are writing a GUI application, it is your responsibility to write the classes for the event listeners that your application needs. For example, if you write an application with a `JButton` component, an event will be generated each time the user clicks the button. Therefore, you should write an event listener class that can handle the event. In your application you would create an instance of the event listener class and connect it to the `JButton` component. Before looking at a specific example, we must discuss two important topics that arise when writing event listeners: private inner classes and interfaces.

## Writing Event Listener Classes as Private Inner Classes

Java allows you to write a class definition inside of another class definition. A class that is defined inside of another class is known as an *inner class*. [Figure 23-8](#) illustrates a class definition inside of another class definition.

```
public class Outer
{
 Fields and methods of the outer
 class appear here.

 private class Inner
 {
 Fields and methods of the Inner
 class appear here.
 }
}
```

### Figure 23-8 A class with an inner class

[Figure 23-8 Full Alternative Text](#)

When an inner class is private, as shown in the figure, it is accessible only to code in the class that contains it. For example, the `Inner` class shown in the figure would be accessible only to methods

that belong to the `Outer` class. Code outside the `Outer` class would not be able to access the `Inner` class. A common technique for writing an event listener class is to write it as a private inner class, inside the class that creates the GUI. Although this is not the only way to write event listener classes, it is the approach we take in this book.

# Event Listeners Must Implement an Interface

There is a special requirement that all event listener classes must meet: They must *implement an interface*.

We discussed interfaces in detail in [Chapter 10](#), but in case you haven't read that material, you can think of an interface as something like a class, containing one or more method headers. Interfaces do not have actual methods, however, only their headers. When you write a class that implements an interface, you are agreeing that the class will have all of the methods that are specified in the interface.

Java provides numerous interfaces that you can use with event listener classes. There are several different types of events that can occur within a GUI application, and the specific interface that you use depends on the type of event you want to handle. `JButton` components generate *action events*, and an event listener class that can handle action events is also known as an *action listener* class. When you write an action listener class for a `JButton` component, it must implement an interface known as `ActionListener`. In case you are curious, this is what the code for the `ActionListener` interface looks like:

```
public interface ActionListener
{
 public void actionPerformed(ActionEvent e);
}
```

As you can see, the `ActionListener` interface contains the header for only one method: `actionPerformed`. Notice that the method has `public` access, is `void`, and has a parameter of the `ActionEvent` type. When you write a class that implements this interface, it *must* have a method named `actionPerformed`, with a header exactly like the one in the interface.

## NOTE:

The `ActionListener` interface, as well as other event listener interfaces, is in the `java.awt.event` package. We will use the following `import` statement in order to use those interfaces:

```
import java.awt.event.*;
```

You use the `implements` key word in a class header to indicate that it implements an interface. Here is an example of a class named `MyButtonListener` that implements the `ActionListener` interface:

```
private class MyButtonListener implements ActionListener
{
 public void actionPerformed(ActionEvent e)
 {
 Write code here to handle the event.
 }
}
```

Remember, when you write a class that implements an interface, you are “promising” that the class will have the methods specified in the interface. Notice that this class lives up to its promise. It has a method named `actionPerformed`, with a header that matches the `actionPerformed` header in the `ActionListener` interface exactly.

## NOTE:

In your action listener class, the only part of the `actionPerformed` method header that does not have to match that which is shown in the `ActionListener` interface exactly is the name of the parameter variable. Instead of using the name `e`, you can use any legal variable name that you wish.

## Registering an Event Listener Object

Once you have written an event listener class, you can create an object of that class, and then connect the object with a GUI component. The process of connecting an event listener object to a GUI component is known as *registering* the event listener.

When a `JButton` component generates an event, it automatically executes the `actionPerformed` method of the event listener object that is registered with it, passing the event object as an argument. This is illustrated in [Figure 23-9](#).




**Figure 23-9 A `JButton` component firing**

# an action event

(Oracle Corporate Counsel)

[Figure 23-9 Full Alternative Text](#)

## Writing an Event Listener for the KiloConverter Class

Now that we've gone over the basics of event listeners, let's continue to develop the `KiloConverter` class. [Code Listing 23-6](#) shows the class with an action listener added to it. This version of the class is stored in the source code folder [Chapter 23\KiloConverter Phase 2](#). The action listener is a private inner class named `CalcButtonListener`. The new code is shown in bold.

### Code Listing 23-6 (`KiloConverter.java`)

```
1 import javax.swing.*; // Needed for Swing classes
2 import java.awt.event.*; // Needed for ActionListener Interface
3
4 /**
5 The KiloConverter class displays a JFrame that
6 lets the user enter a distance in kilometers. When
7 the Calculate button is clicked, a dialog box is
8 displayed with the distance converted to miles.
9 */
10
11 public class KiloConverter extends JFrame
12 {
13 private JPanel panel; // To reference a panel
14 private JLabel messageLabel; // To reference a label
15 private JTextField kiloTextField; // To reference a text field
16 private JButton calcButton; // To reference a button
17 private final int WINDOW_WIDTH = 310; // Window width
18 private final int WINDOW_HEIGHT = 100; // Window height
19
20 /**
21 Constructor
22 */
23
24 public KiloConverter()
25 {
26 // Set the window title.
27 setTitle("Kilometer Converter");
28
29 // Set the size of the window.
30 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
31 }
32 }
```



```

32 // Specify what happens when the close button is clicked.
33 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
34
35 // Build the panel and add it to the frame.
36 buildPanel();
37
38 // Add the panel to the frame's content pane.
39 add(panel);
40
41 // Display the window.
42 setVisible(true);
43 }
44
45 /**
46 The buildPanel method adds a label, a text field,
47 and a button to a panel.
48 */
49
50 private void buildPanel()
51 {
52 // Create a label to display instructions.
53 messageLabel = new JLabel("Enter a distance " +
54 "in kilometers");
55
56 // Create a text field 10 characters wide.
57 kiloTextField = new JTextField(10);
58
59 // Create a button with the caption "Calculate".
60 calcButton = new JButton("Calculate");
61
62 // Add an action listener to the button.
63 calcButton.addActionListener(new CalcButtonListener());
64
65 // Create a JPanel object and let the panel
66 // field reference it.
67 panel = new JPanel();
68
69 // Add the label, text field, and button
70 // components to the panel.
71 panel.add(messageLabel);
72 panel.add(kiloTextField);
73 panel.add(calcButton);
74 }
75
76 /**
77 CalcButtonListener is an action listener class for
78 the Calculate button.
79 */
80
81 private class CalcButtonListener implements ActionListener
82 {
83 /**
84 The actionPerformed method executes when the user
85 clicks on the Calculate button.
86 @param e The event object.
87 */
88

```

```

89 public void actionPerformed(ActionEvent e)
90 {
91 final double CONVERSION = 0.6214;
92 String input; // To hold the user's input
93 double miles; // The number of miles
94
95 // Get the text entered by the user into the
96 // text field.
97 input = kiloTextField.getText();
98
99 // Convert the input to miles.
100 miles = Double.parseDouble(input) * CONVERSION;
101
102 // Display the result.
103 JOptionPane.showMessageDialog(null, input +
104 " kilometers is " + miles + " miles.");
105 }
106 }
107
108 /**
109 main method
110 */
111
112 public static void main(String[] args)
113 {
114 new KiloConverter();
115 }
116 }

```

First, notice that we've added the `import java.awt.event.*;` statement in line 2. This is necessary for our program to use the `ActionListener` interface. Next, look at the following code in line 81:

```
private class CalcButtonListener implements ActionListener
```

This is the header for an inner class that we will use to create event listener objects. The name of this class is `CalcButtonListener` and it implements the `ActionListener` interface. We could have named the class anything we wanted to, but because it will handle the `JButton` component's action events, it must implement the `ActionListener` interface. The class has one method, `actionPerformed`, which is required by the `ActionListener` interface. The header for the `actionPerformed` method appears in line 89 as follows:

```
public void actionPerformed(ActionEvent e)
```

This method will be executed when the user clicks the `JButton` component. It has one parameter, `e`, which is an `ActionEvent` object. This parameter receives the event object that is passed to the method when it is called. Although we do not actually use the `e` parameter in this method, we still have to list it inside the method header's parentheses because it is required by the `ActionListener` interface.

The `actionPerformed` method declares a constant for the conversion factor in line 91, and two local variables in lines 92 and 93: `input`, a reference to a `String` object; and `miles`, a `double`. The following statement appears in line 97:

```
input = kiloTextField.getText();
```

All `JTextField` objects have a `getText` method that returns the text contained in the text field. This will be any value entered into the text field by the user. The value is returned as a string. So, this statement retrieves any value entered by the user into the text field and assigns it to `input`.

The following statement appears in line 100:

```
miles = Double.parseDouble(input) * CONVERSION;
```

This statement converts the value in `input` to a double, and then multiplies it by the constant `CONVERSION`, which is set to 0.6214. This will convert the number of kilometers entered by the user to miles. The result is stored in the `miles` variable. The method's last statement, in lines 103 and 104, uses `JOptionPane` to display a dialog box showing the distance converted to miles as follows:

```
JOptionPane.showMessageDialog(null, input +
 " kilometers is " + miles + " miles.");
```

Writing an action listener class is only part of the process of handling a `JButton` component's action events. We must also create an object from the class and then register the object with the `JButton` component. When we register the action listener object with the `JButton` component, we are creating a connection between the two objects.

`JButton` components have a method named `addActionListener`, which is used for registering action event listeners. In line 63, which is in the `buildPanel` method, the following statement creates a `CalcButtonListener` object and registers that object with the `calcButton` object:

```
calcButton.addActionListener(new CalcButtonListener());
```

You pass the address of an action listener object as the argument to the `addActionListener` method. This statement uses the expression `new CalcButtonListener()` to create an instance of the `CalcButtonListener` class. The address of that instance is then passed to the `addActionListener` method. Now, when the user clicks the Calculate button, the `CalcButtonListener` object's `actionPerformed` method will be executed.



Instead of the one statement in line 63, we could have written the following two statements:

```
CalcButtonListener listener = new CalcButtonListener();
calcButton.addActionListener(listener);
```

The first statement shown here declares a `CalcButtonListener` variable named `listener`, creates a new `CalcButtonListener` object, and assigns the object's address to the `listener` variable. The second statement passes the address in `listener` to the `addActionListener` method. These two statements accomplish the same thing as the one statement in line 63, but they declare a variable,


listener, that we will not use again in the program. A better way is to use the one statement that appears in line 63 as follows:

```
calcButton.addActionListener(new CalcButtonListener());
```


Recall that the `new` key word creates an object and returns the object's address. This statement uses the `new` key word to create a `CalcButtonListener` object, and passes the object's address directly to the `addActionListener` method. Because we do not need to refer to the object again in the program, we do not assign the object's address to a variable. It is known as an *anonymous object*.

When this program is executed, the first window shown in [Figure 23-10](#) is displayed on the screen. If the user enters 2 in the text field and clicks the Calculate button, the second window shown in the figure (a dialog box) appears. To exit the application, the user clicks the OK button on the dialog box, and then clicks the close button in the upper-right corner of the main window.

This window appears first. The user enters 2 in the text field and then clicks the Calculate button.



This dialog box appears next.



## Figure 23-10 Windows displayed by the KiloConverter class

(Oracle Corporate Counsel)

[Figure 23-10 Full Alternative Text](#)

## Background and Foreground Colors

Many of the Swing component classes have methods named `setBackground` and `setForeground`. You call these methods to change a component's color. The background color is the color of the component itself, and the foreground color is the color of text that might be displayed on the component.

The argument that you pass to the `setBackground` and `setForeground` methods is a color code. [Table 23-3](#) lists several predefined constants that you can use for colors. To use these constants, you must have the `import java.awt.*;` statement in your code.

# Table 23-3 Color constants

(Oracle Corporate Counsel)

|                               |                              |
|-------------------------------|------------------------------|
| <code>Color.BLACK</code>      | <code>Color.BLUE</code>      |
| <code>Color.CYAN</code>       | <code>Color.DARK_GRAY</code> |
| <code>Color.GRAY</code>       | <code>Color.GREEN</code>     |
| <code>Color.LIGHT_GRAY</code> | <code>Color.MAGENTA</code>   |
| <code>Color.ORANGE</code>     | <code>Color.PINK</code>      |
| <code>Color.RED</code>        | <code>Color.WHITE</code>     |
| <code>Color.YELLOW</code>     |                              |

For example, the following code creates a button with the text “OK” displayed on it. The `setBackground` and `setForeground` methods are called to make the button blue and the text yellow.

```
JButton okButton = new JButton("OK");
okButton.setBackground(Color.BLUE);
okButton.setForeground(Color.YELLOW);
```

The `ColorWindow` class in [Code Listing 23-7](#) displays a window with a label and three buttons. When the user clicks a button, it changes the background color of the panel that contains the components and the foreground color of the label.

## Code Listing 23-7 (`ColorWindow.java`)

```
1 import javax.swing.*; // Needed for Swing classes
2 import java.awt.*; // Needed for Color class
3 import java.awt.event.*; // Needed for event listener interface
4
5 /**
6 * This class demonstrates how to set the background color of
7 * a panel and the foreground color of a label.
8 */
```

```

9
10 public class ColorWindow extends JFrame
11 {
12 private JLabel messageLabel; // To display a message
13 private JButton redButton; // Changes color to red
14 private JButton blueButton; // Changes color to blue
15 private JButton yellowButton; // Changes color to yellow
16 private JPanel panel; // A panel to hold components
17 private final int WINDOW_WIDTH = 200; // Window width
18 private final int WINDOW_HEIGHT = 125; // Window height
19
20 /**
21 * Constructor
22 */
23
24 public ColorWindow()
25 {
26 // Set the title bar text.
27 setTitle("Colors");
28
29 // Set the size of the window.
30 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
31
32 // Specify an action for the close button.
33 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
34
35 // Create a label.
36 messageLabel = new JLabel("Click a button to " +
37 "select a color.");
38
39 // Create the three buttons.
40 redButton = new JButton("Red");
41 blueButton = new JButton("Blue");
42 yellowButton = new JButton("Yellow");
43
44 // Register an event listener with all 3 buttons.
45 redButton.addActionListener(new RedButtonListener());
46 blueButton.addActionListener(new BlueButtonListener());
47 yellowButton.addActionListener(new YellowButtonListener());
48
49 // Create a panel and add the components to it.
50 panel = new JPanel();
51 panel.add(messageLabel);
52 panel.add(redButton);
53 panel.add(blueButton);
54 panel.add(yellowButton);
55
56 // Add the panel to the content pane.
57 add(panel);
58
59 // Display the window.
60 setVisible(true);
61 }
62
63 /**
64 * Private inner class that handles the event when
65 * the user clicks the Red button.

```

```

66 */
67
68 private class RedButtonListener implements ActionListener
69 {
70 public void actionPerformed(ActionEvent e)
71 {
72 // Set the panel's background to red.
73 panel.setBackground(Color.RED);
74
75 // Set the label's text to blue.
76 messageLabel.setForeground(Color.BLUE);
77 }
78 }
79
80 /**
81 Private inner class that handles the event when
82 the user clicks the Blue button.
83 */
84
85 private class BlueButtonListener implements ActionListener
86 {
87 public void actionPerformed(ActionEvent e)
88 {
89 // Set the panel's background to blue.
90 panel.setBackground(Color.BLUE);
91
92 // Set the label's text to yellow.
93 messageLabel.setForeground(Color.YELLOW);
94 }
95 }
96
97 /**
98 Private inner class that handles the event when
99 the user clicks the Yellow button.
100 */
101
102 private class YellowButtonListener implements ActionListener
103 {
104 public void actionPerformed(ActionEvent e)
105 {
106 // Set the panel's background to yellow.
107 panel.setBackground(Color.YELLOW);
108
109 // Set the label's text to black.
110 messageLabel.setForeground(Color.BLACK);
111 }
112 }
113
114 /**
115 main method
116 */
117
118 public static void main(String[] args)
119 {
120 new ColorWindow();
121 }
122 }

```

Notice that this class has three action listener classes, one for each button. The action listener classes are `RedButtonListener`, `BlueButtonListener`, and `YellowButtonListener`. The following statements, in lines 45 through 47, register instances of these classes with the appropriate button components:

```
redButton.addActionListener(new RedButtonListener());
blueButton.addActionListener(new BlueButtonListener());
yellowButton.addActionListener(new YellowButtonListener());
```

When you run the program, the window shown in [Figure 23-11](#) appears.

The window components first appear in their default colors.



When the user clicks on the Red button, the panel turns red and the label turns blue.



When the user clicks on the Blue button, the panel turns blue and the label turns yellow.



When the user clicks on the Yellow button, the panel turns yellow and the label turns black.



**Figure 23-11 The window produced by the `ColorWindow` class**

(Oracle Corporate Counsel)

[Figure 23-11 Full Alternative Text](#)

## Changing the Background Color of a `JFrame` Object's Content Pane

Recall that a `JFrame` object has a content pane, which is a container for all the components that are added to the `JFrame`. When you add a component to a `JFrame` object, you are actually adding it to the



object's content pane. In the example shown in this section, we added a label and some buttons to a panel, and then added the panel to the `JFrame` object's content pane. When we changed the background color, we changed the background color of the panel. In this example, the color of the content pane does not matter because it is completely filled up by the panel. The color of the panel covers up the color of the content pane.

In some cases, where you have not filled up the `JFrame` object's content pane with a panel, you might want to change the background color of the content pane. If you wish to change the background color of a `JFrame` object's content pane, you must call the content pane's `setBackground` method, not the `JFrame` object's `setBackground` method. For example, in a class that extends the `JFrame` class, the following statement can be used to change the content pane's background to blue:

```
getContentPane().setBackground(Color.BLUE);
```

In this statement, the `getContentPane` method is called to get a reference to the `JFrame` object's content pane. This reference is then used to call the content pane's `setBackground` method. As a result, the content pane's background color will change to blue.

# The `ActionEvent` Object

The action listener's `actionPerformed` method has a parameter variable named `e` that is declared as follows:

```
ActionEvent e
```

`ActionEvent` is a class that is defined in the Java API. When an action event occurs, an object of the `ActionEvent` class is created, the action listener's `actionPerformed` method is called, and a reference to the `ActionEvent` object is passed into the `e` parameter variable. So, when the `actionPerformed` method executes, the `e` parameter references the event object that was generated in response to the event.

Earlier it was mentioned that the event object contains information about the event. If you wish, you can retrieve certain information about the event by calling one of the event object's methods. Two of the `ActionEvent` methods are listed in [Table 23-4](#).

## Table 23-4 `ActionEvent` methods

| Method Name                     | Description                                                        |
|---------------------------------|--------------------------------------------------------------------|
| <code>getActionCommand()</code> | Returns the action command for this event as a <code>String</code> |

# The `getActionCommand` Method

The first method listed in [Table 23-4](#), `getActionCommand`, returns the *action command* that is associated with the event. When a `JButton` component generates an event, the action command is the text that appears on the button. The `getActionCommand` returns this text as a `String`. You can use the `getActionCommand` method to determine which button was clicked when several buttons share the same action listener class.

To demonstrate, look at the `EventObjectWindow` class in [Code Listing 23-8](#). It produces a window with three buttons. The buttons have the text “Button 1”, “Button 2”, and “Button 3”. The action listener class displays the contents of the event object’s action command when any of these buttons are clicked.

## Code Listing 23-8 (`EventObjectWindow.java`)

```

1 import javax.swing.*; // Needed for Swing classes
2 import java.awt.event.*; // Needed for event listener interface
3
4 /**
5 This class demonstrates how to retrieve the action command
6 from an event object.
7 */
8
9 public class EventObject extends JFrame
10 {
11 private JButton button1; // Button 1
12 private JButton button2; // Button 2
13 private JButton button3; // Button 3
14 private JPanel panel; // A panel to hold components
15 private final int WINDOW_WIDTH = 300; // Window width
16 private final int WINDOW_HEIGHT = 70; // Window height
17
18 /**
19 Constructor
20 */
21
22 public EventObject()
23 {
24 // Set the title bar text.
25 setTitle("Event Object Demonstration");
26
27 // Set the size of the window.
28 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
29
30 // Specify what happens when the close button is clicked.
31 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);

```

```

32
33 // Create the three buttons.
34 button1 = new JButton("Button 1");
35 button2 = new JButton("Button 2");
36 button3 = new JButton("Button 3");
37
38 // Register an event listener with all 3 buttons.
39 button1.addActionListener(new ButtonListener());
40 button2.addActionListener(new ButtonListener());
41 button3.addActionListener(new ButtonListener());
42
43 // Create a panel and add the buttons to it.
44 panel = new JPanel();
45 panel.add(button1);
46 panel.add(button2);
47 panel.add(button3);
48
49 // Add the panel to the content pane.
50 add(panel);
51
52 // Display the window.
53 setVisible(true);
54 }
55
56 /**
57 Private inner class that handles the event when
58 the user clicks a button.
59 */
60
61 private class ButtonListener implements ActionListener
62 {
63 public void actionPerformed(ActionEvent e)
64 {
65 // Get the action command.
66 String actionCommand = e.getActionCommand();
67
68 // Determine which button was clicked and display
69 // a message.
70 if (actionCommand.equals("Button 1"))
71 {
72 JOptionPane.showMessageDialog(null, "You clicked " +
73 "the first button.");
74 }
75 else if (actionCommand.equals("Button 2"))
76 {
77 JOptionPane.showMessageDialog(null, "You clicked " +
78 "the second button.");
79 }
80 else if (actionCommand.equals("Button 3"))
81 {
82 JOptionPane.showMessageDialog(null, "You clicked " +
83 "the third button.");
84 }
85 }
86 }
87
88 /**

```

```

89 main method
90 */
91
92 public static void main(String[] args)
93 {
94 new EventObject();
95 }
96 }

```

Previously you saw the `ColorWindow` class, in [Code Listing 23-7](#), which had three buttons and three different action listener classes. The `EventObjectWindow` class also has three buttons, but only one action listener class. In lines 39 through 41, we create and register three separate instances of the class with the three buttons as follows:

```

button1.addActionListener(new ButtonListener());
button2.addActionListener(new ButtonListener());
button3.addActionListener(new ButtonListener());

```

[Figure 23-12](#) shows the output of the application when the user clicks Button 1, Button 2, and Button 3.




**Figure 23-12 Output of `EventObjectWindow`**

# class

(Oracle Corporate Counsel)

[Figure 23-12 Full Alternative Text](#)

## TIP:

The text that is displayed on a button is the default action command. You can change the action command by calling the `JButton` class's `setActionCommand` method. For example, assuming that `myButton` references a `JButton` component, the following statement would change the component's action command to "The button was clicked":

```
myButton.setActionCommand("The button was clicked");
```

## NOTE:

Changing a `JButton` component's action command does not change the text that is displayed on the button. For a demonstration of how to change the action command, see the *ActionCommand.java* file in this chapter's source code folder.

## The getSource Method

The second `ActionEvent` method listed in [Table 23-4](#), `getSource`, returns a reference to the component that is the source of the event. As with the `getActionCommand` method, if you have several buttons and use objects of the same action listener class to respond to their events, you can use the `getSource` method to determine which button was clicked. For example, the `ButtonListener` class's `actionPerformed` method in [Code Listing 23-8](#) could have been written as follows, to achieve the same result:

```
public void actionPerformed(ActionEvent e)
{
 // Determine which button was clicked and display
 // a message.
 if (e.getSource() == button1)
 {
 JOptionPane.showMessageDialog(null, "You clicked " +
 "the first button.");
 }
 else if (e.getSource() == button2)
 {
 JOptionPane.showMessageDialog(null, "You clicked " +
```

```

 "the second button.");
 }
 else if (e.getSource() == button3)
 {
 JOptionPane.showMessageDialog(null, "You clicked " +
 "the third button.");
 }
}

```

See the *EventObjectWindow2.java* file in this chapter's source code folder for a demonstration of this code.



## Checkpoint

MyProgrammingLab™ [www.myprogramminglab.com](http://www.myprogramminglab.com)

1. 23.1 What is a frame? How do you create a frame with Swing?
2. 23.2 How do you set a frame's size?
3. 23.3 How do you display a frame on the screen?
4. 23.4 What is a content pane?
5. 23.5 What is the difference between a frame and a panel?
6. 23.6 What is an event listener?
7. 23.7 If you are writing an event listener class for a `JButton` component, what interface must the class implement? What method must the class have? When is this method executed?
8. 23.8 How do you register an event listener with a `JButton` component?
9. 23.9 How do you change the background color of a component? How do you change the color of text displayed by a label or a button?

# 23.3 Layout Managers

CONCEPT: A layout manager is an object that governs the positions and sizes of components in a container. The layout manager automatically repositions and, in some cases, resizes the components when the container is resized.

An important part of designing a GUI application is determining the layout of the components that are displayed in the application’s windows. The term *layout* refers to the positioning and sizing of components. In Java, you do not normally specify the exact location of a component within a window. Instead, you let a layout manager control the positions of components for you. A *layout manager* is an object that has its own rules about how components are to be positioned and sized, and it makes adjustments when necessary. For example, when the user resizes a window, the layout manager determines where the components should be moved to.

In order to use a layout manager with a group of components, you must place the components in a container, and then create a layout manager object. The layout manager object and the container work together. In this chapter we discuss the three layout managers described in [Table 23-5](#). To use any of these classes, your code should have the following `import` statement: `import java.awt.*;`

Table 23-5 Layout managers

| Layout Manager | Description                                                                                                                                                   |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FlowLayout     | Arranges components in rows; this is the default layout manager for <code>JPanel</code> objects                                                               |
| BorderLayout   | Arranges components in five regions: north, south, east, west, and center; this is the default layout manager for a <code>JFrame</code> object’s content pane |
| GridLayout     | Arranges components in a grid with rows and columns                                                                                                           |

## Adding a Layout Manager to a Container

You add a layout manager to a container, such as a content pane or a panel, by calling the `setLayout`

method and passing a reference to a layout manager object as the argument. For example, the following code creates a `JPanel` object, then sets a `BorderLayout` object as its layout manager:

```
JPanel panel = new JPanel();
panel.setLayout(new BorderLayout());
```

Likewise, the following code might appear in the constructor of a class that extends the `JFrame` class. It sets a `FlowLayout` object as the layout manager for the content pane:

```
setLayout(new FlowLayout());
```

Once you establish a layout manager for a container, the layout manager governs the positions and sizes of the components that are added to the container.

# The `FlowLayout` Manager

The `FlowLayout` manager arranges components in rows. This is the default layout manager for `JPanel` objects. Here are some rules that the `FlowLayout` manager follows:

- You can add multiple components to a container that uses a `FlowLayout` manager.
- When you add components to a container that uses a `FlowLayout` manager, the components appear horizontally, from left to right, in the order that they were added to the component.
- When there is no more room in a row but more components are added, the new components “flow” to the next row.

For example, the `FlowWindow` class shown in [Code Listing 23-9](#) extends `JFrame`. This class creates a 200 pixel wide by 105 pixel high window. In the constructor, the `setLayout` method is called to give the content pane a `FlowLayout` manager. Then, three buttons are created and added to the content pane. The `main` method creates an instance of the `FlowWindow` class, which displays the window.

## Code Listing 23-9 (`FlowWindow.java`)

```
1 import javax.swing.*; // Needed for Swing classes
2 import java.awt.*; // Needed for FlowLayout class
3
4 /**
5 * This class demonstrates how to use a FlowLayout manager
6 * with the content pane.
7 */
8
9 public class FlowWindow extends JFrame
10 {
11 private final int WINDOW_WIDTH = 200; // Window width
12 private final int WINDOW_HEIGHT = 105; // Window height
13 }
```



```

14 /**
15 Constructor
16 */
17
18 public FlowWindow()
19 {
20 // Set the title bar text.
21 setTitle("Flow Layout");
22
23 // Set the size of the window.
24 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
25
26 // Specify an action for the close button.
27 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
28
29 // Add a FlowLayout manager to the content pane.
30 setLayout(new FlowLayout());
31
32 // Create three buttons.
33 JButton button1 = new JButton("Button 1");
34 JButton button2 = new JButton("Button 2");
35 JButton button3 = new JButton("Button 3");
36
37 // Add the three buttons to the content pane.
38 add(button1);
39 add(button2);
40 add(button3);
41
42 // Display the window.
43 setVisible(true);
44 }
45
46 /**
47 The main method creates an instance of the FlowWindow
48 class, causing it to display its window.
49 */
50
51 public static void main(String[] args)
52 {
53 new FlowWindow();
54 }
55 }

```

[Figure 23-13](#) shows the window that is displayed by this class. Notice that the buttons appear from left to right in the order they were added to the content pane. Because there is only enough room for the first two buttons in the first row, the third button is positioned in the second row. By default, the content of each row is centered and there is a five pixel gap between the components.



# Figure 23-13 The window displayed by the `FlowWindow` class

(Oracle Corporate Counsel)

If the user resizes the window, the layout manager repositions the components according to its rules. [Figure 23-14](#) shows the appearance of the window in three different sizes.



## Figure 23-14 The arrangements of the buttons after resizing

[Figure 23-14 Full Alternative Text](#)

## Adjusting the `FlowLayout` Alignment

The `FlowLayout` manager allows you to align components in the center of each row or along the left or right edge of each row. An overloaded constructor allows you to pass one of the following constants as an argument to set an alignment: `FlowLayout.CENTER`, `FlowLayout.LEFT`, or `FlowLayout.RIGHT`. Here is an example that sets left alignment:

```
setLayout(new FlowLayout(FlowLayout.LEFT));
```

[Figure 23-15](#) shows examples of windows that use a `FlowLayout` manager with left, center, and right alignment.




**Figure 23-15 Left, center, and right alignment**

(Oracle Corporate Counsel)

[Figure 23-15 Full Alternative Text](#)

## Adjusting the `FlowLayout` Component Gaps

By default, the `FlowLayout` manager inserts a gap of five pixels between components, both horizontally and vertically. You can adjust this gap by passing values for the horizontal and vertical gaps as arguments to an overloaded `FlowLayout` constructor. The constructor has the following format:

```
FlowLayout(int alignment, int horizontalGap, int verticalGap)
```


You pass one of the alignment constants discussed in the previous section to the `alignment` parameter. The `horizontalGap` parameter is the number of pixels to separate components horizontally, and the `verticalGap` parameter is the number of pixels to separate components vertically. Here is an example of the constructor call:

```
setLayout(new FlowLayout(FlowLayout.LEFT, 10, 7));
```

This statement causes components to be left aligned with a horizontal gap of 10 pixels and a vertical gap of seven pixels.

# The BorderLayout Manager

The `BorderLayout` manager divides a container into five regions. The regions are known as north, south, east, west, and center. The arrangement of these regions is shown in [Figure 23-16](#).



## Figure 23-16 The regions of a BorderLayout manager

(Oracle Corporate Counsel)

### [Figure 23-16 Full Alternative Text](#)

When a component is placed into a container that is managed by a `BorderLayout` manager, the component must be placed into one of these five regions. Only one component at a time may be placed into a region. When adding a component to the container, you specify the region by passing one of the following constants as a second argument to the container's `add` method: `BorderLayout.NORTH`, `BorderLayout.SOUTH`, `BorderLayout.EAST`, `BorderLayout.WEST`, or `BorderLayout.CENTER`.

For example, look at the following code:

```
JPanel panel = new JPanel();
JButton button = new JButton("Click Me");
panel.setLayout(new BorderLayout());
panel.add(button, BorderLayout.NORTH);
```

The first statement creates a `JPanel` object, referenced by the `panel` variable. The second statement creates a `JButton` object, referenced by the `button` variable. The third statement sets the `JPanel` object's layout manager to a `BorderLayout` object. The fourth statement adds the `JButton` object to the `JPanel` object's north region.

If you do not pass a second argument to the `add` method, the component will be added to the center region. Here are some rules that the `BorderLayout` manager follows:

- Each region can hold only one component at a time.
- When a component is added to a region, the component is stretched so it fills up the entire region.

Look at the `BorderWindow` class shown in [Code Listing 23-10](#), which extends `JFrame`. This class creates a 400 pixel wide by 300 pixel high window. In the constructor, the `setLayout` method is called to give the content pane a `BorderLayout` manager. Then, five buttons are created and each is added to a different region.

## Code Listing 23-10 (`BorderWindow.java`)

```
1 import javax.swing.*; // Needed for Swing classes
2 import java.awt.*; // Needed for BorderLayout class
3
4 /**
5 * This class demonstrates the BorderLayout manager.
6 */
7
8 public class BorderWindow extends JFrame
9 {
10 private final int WINDOW_WIDTH = 400; // Window width
11 private final int WINDOW_HEIGHT = 300; // Window height
12
13 /**
14 * Constructor
15 */
16
17 public BorderWindow()
18 {
19 // Set the title bar text.
20 setTitle("Border Layout");
21
22 // Set the size of the window.
23 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
24
25 // Specify an action for the close button.
26 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
27
28 // Add a BorderLayout manager to the content pane.
29 setLayout(new BorderLayout());
30
31 // Create five buttons.
```

```

32 JButton button1 = new JButton("North Button");
33 JButton button2 = new JButton("South Button");
34 JButton button3 = new JButton("East Button");
35 JButton button4 = new JButton("West Button");
36 JButton button5 = new JButton("Center Button");
37
38 // Add the five buttons to the content pane.
39 add(button1, BorderLayout.NORTH);
40 add(button2, BorderLayout.SOUTH);
41 add(button3, BorderLayout.EAST);
42 add(button4, BorderLayout.WEST);
43 add(button5, BorderLayout.CENTER);
44
45 // Display the window.
46 setVisible(true);
47 }
48
49 /**
50 * The main method creates an instance of the BorderWindow
51 * class, causing it to display its window.
52 */
53
54 public static void main(String[] args)
55 {
56 new BorderWindow();
57 }
58 }

```

## NOTE:

A `JFrame` object's content pane is automatically given a `BorderLayout` manager. We have explicitly added it in [Code Listing 23-10](#) so it is clear that we are using a `BorderLayout` manager.

[Figure 23-17](#) shows the window that is displayed. Normally the size of a button is just large enough to accommodate the text that is displayed on the button. Notice that the buttons displayed in this window did not retain their normal size. Instead, they were stretched to fill all of the space in their regions. If the user resizes the window, the sizes of the components will be changed as well. This is shown in [Figure 23-18](#).




**Figure 23-17 The window displayed by the `BorderWindow` class**

(Oracle Corporate Counsel)

[Figure 23-17 Full Alternative Text](#)



**Figure 23-18 The window resized**

(Oracle Corporate Counsel)

[Figure 23-18 Full Alternative Text](#)

Here are the rules that govern how a `BorderLayout` manager resizes components:

- A component that is placed in the north or south regions may be resized horizontally so it fills up the entire region.
- A component that is placed in the east or west regions may be resized vertically so it fills up the entire region.
- A component that is placed in the center region may be resized both horizontally and vertically so it fills up the entire region.



You do not have to place a component in every region of a border layout. To achieve the desired positioning, you might want to place components in only a few of the layout regions. In [Chapter 24](#), you will see examples of applications that do this.

By default there is no gap between the regions. You can use an overloaded version of the `BorderLayout` constructor to specify horizontal and vertical gaps, however. Here is the constructor's format:

```
BorderLayout(int horizontalGap, int verticalGap)
```

The `horizontalGap` parameter is the number of pixels to separate the regions horizontally, and the `verticalGap` parameter is the number of pixels to separate the regions vertically. Here is an example of the constructor call:

```
setLayout(new BorderLayout(5, 10));
```

This statement causes the regions to appear with a horizontal gap of five pixels and a vertical gap of 10 pixels.

## Nesting Panels Inside a Container's Regions


You might think that the `BorderLayout` manager is limiting because it allows only one component per region, and the components that are placed in its regions are automatically resized to fill up any extra space. These limitations are easy to overcome, however, by adding components to panels and then nesting the panels inside the regions.

For example, suppose we wish to modify the `BorderWindow` class in [Code Listing 23-10](#) so the buttons retain their original size. We can accomplish this by placing each button in a separate `JPanel` object and then adding the `JPanel` objects to the content pane's five regions. This is illustrated in [Figure 23-19](#). As a result, the `BorderLayout` manager resizes the `JPanel` objects to fill up the space




in the regions, not the buttons contained within the `JPanel` objects.

1. Five `JPanel` objects are created and a `JButton` object is added to each one.




2. The `JPanel` objects are then added to the content pane, one to each region.



## Figure 23-19 Nesting `JPanel` objects inside each region

[Figure 23-19 Full Alternative Text](#)

The `BorderPanelWindow` class in [Code Listing 23-11](#) demonstrates this technique. This class also introduces a new way of sizing windows. Notice that the constructor does not explicitly set the size of the window with the `setSize` method. Instead, it calls the `pack` method just before calling the `setVisible` method. The `pack` method, which is inherited from `JFrame`, automatically sizes the window to accommodate the components contained within it. [Figure 23-20](#) shows the window that the class displays.



## Figure 23-20 Window displayed by the `BorderPanelWindow` class

(Oracle Corporate Counsel)

## Code Listing 23-11 (`BorderPanelWindow.java`)

```

1 import java.awt.*; // Needed for BorderLayout class
2 import javax.swing.*; // Needed for Swing classes
3
4 /**
5 This class demonstrates how JPanels can be nested
6 inside each region of a content pane governed by
7 a BorderLayout manager.
8 */
9
10 public class BorderPanelWindow extends JFrame
11 {
12 /**
13 Constructor
14 */
15
16 public BorderPanelWindow()
17 {
18 // Set the title bar text.
19 setTitle("Border Layout");
20
21 // Specify an action for the close button.
22 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
23
24 // Add a BorderLayout manager to the content pane.
25 setLayout(new BorderLayout());
26
27 // Create five panels.
28 JPanel panel1 = new JPanel();
29 JPanel panel2 = new JPanel();
30 JPanel panel3 = new JPanel();
31 JPanel panel4 = new JPanel();
32 JPanel panel5 = new JPanel();
33
34 // Create five buttons.
35 JButton button1 = new JButton("North Button");
36 JButton button2 = new JButton("South Button");
37 JButton button3 = new JButton("East Button");
38 JButton button4 = new JButton("West Button");
39 JButton button5 = new JButton("Center Button");
40
41 // Add the buttons to the panels.
42 panel1.add(button1);
43 panel2.add(button2);
44 panel3.add(button3);
45 panel4.add(button4);
46 panel5.add(button5);
47
48 // Add the five panels to the content pane.
49 add(panel1, BorderLayout.NORTH);
50 add(panel2, BorderLayout.SOUTH);
51 add(panel3, BorderLayout.EAST);
52 add(panel4, BorderLayout.WEST);
53 add(panel5, BorderLayout.CENTER);
54
55 // Pack and display the window.
56 pack();
57 setVisible(true);
58 }

```

```

59
60 /**
61 The main method creates an instance of the
62 BorderLayoutWindow class, causing it to display
63 its window.
64 */
65
66 public static void main(String[] args)
67 {
68 new BorderLayoutWindow();
69 }
70 }

```

## NOTE:

There are multiple layout managers at work in the `BorderLayoutWindow` class. The content pane uses a `BorderLayout` manager, and each of the `JPanel` objects use a `FlowLayout` manager.

## The GridLayout Manager

The `GridLayout` manager creates a grid with rows and columns, much like a spreadsheet. As a result, the container that is managed by a `GridLayout` object is divided into equally sized cells. [Figure 23-21](#) illustrates a container with three rows and five columns. This means that the container is divided into 15 cells.




**Figure 23-21 The `GridLayout` manager divides a container into cells**

Here are some rules that the `GridLayout` manager follows:

- Each cell can hold only one component.
- All of the cells are the same size. This is the size of the largest component placed within the layout.
- A component that is placed in a cell is automatically resized to fill up any extra space.

You pass the number of rows and columns that a container should have as arguments to the `GridLayout` constructor. Here is the general format of the constructor:

```
GridLayout(int rows, int columns)
```

Here is an example of the constructor call:

```
setLayout(new GridLayout(2, 3));
```

This statement gives the container two rows and three columns, for a total of six cells. You can pass 0 as an argument for the rows or the columns, but not both. Passing 0 for both arguments will cause an error.

When adding components to a container that is governed by the `GridLayout` manager, you cannot specify a cell. Instead, the components are assigned to cells in the order they are added. The first component added to the container is assigned to the first cell, which is in the upper-left corner. As other components are added, they are assigned to the remaining cells in the first row, from left to right. When the first row is filled up, components are assigned to the cells in the second row, and so forth.

The `GridWindow` class shown in [Code Listing 23-12](#) demonstrates. It creates a 400 pixel wide by 200 pixel high window, governed by a `GridLayout` manager. The content pane is divided into two rows and three columns, and a button is added to each cell. [Figure 23-22](#) shows the window displayed by the class.



**Figure 23-22 Window displayed by the**

# GridWindow class

(Oracle Corporate Counsel)

## Code Listing 23-12 (GridWindow.java)

```
1 import javax.swing.*; // Needed for Swing classes
2 import java.awt.*; // Needed for GridLayout class
3
4 /**
5 * This class demonstrates the GridLayout manager.
6 */
7
8 public class GridWindow extends JFrame
9 {
10 private final int WINDOW_WIDTH = 400; // Window width
11 private final int WINDOW_HEIGHT = 200; // Window height
12
13 /**
14 * Constructor
15 */
16
17 public GridWindow()
18 {
19 // Set the title bar text.
20 setTitle("Grid Layout");
21
22 // Set the size of the window.
23 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
24
25 // Specify an action for the close button.
26 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
27
28 // Add a GridLayout manager to the content pane.
29 setLayout(new GridLayout(2, 3));
30
31 // Create six buttons.
32 JButton button1 = new JButton("Button 1");
33 JButton button2 = new JButton("Button 2");
34 JButton button3 = new JButton("Button 3");
35 JButton button4 = new JButton("Button 4");
36 JButton button5 = new JButton("Button 5");
37 JButton button6 = new JButton("Button 6");
38
39 // Add the six buttons to the content pane.
40 add(button1); // Goes into row 1, column 1
41 add(button2); // Goes into row 1, column 2
42 add(button3); // Goes into row 1, column 3
43 add(button4); // Goes into row 2, column 1
44 add(button5); // Goes into row 2, column 2
45 add(button6); // Goes into row 2, column 3
46
```

```

47 // Display the window.
48 setVisible(true);
49 }
50
51 /**
52 * The main method creates an instance of the GridWindow
53 * class, causing it to display its window.
54 */
55
56 public static void main(String[] args)
57 {
58 new GridWindow();
59 }
60 }

```

As previously mentioned, the `GridLayout` manager limits each cell to only one component and resizes components to fill up all of the space in a cell. To get around these limitations you can nest panels inside the cells and add other components to the panels. For example, the `GridPanelWindow` class shown in [Code Listing 23-13](#) is a modification of the `GridWindow` class. It creates six panels and adds a button and a label to each panel. These panels are then added to the content pane's cells. [Figure 23-23](#) shows the window displayed by this class.




**Figure 23-23 Window displayed by the `GridPanelWindow` class**

(Oracle Corporate Counsel)

[Figure 23-23 Full Alternative Text](#)

## Code Listing 23-13 (`GridPanelWindow.java`)

```

1 import javax.swing.*; // Needed for Swing classes
2 import java.awt.*; // Needed for GridLayout class
3
4 /**

```

```

5 This class demonstrates how panels may be added to
6 the cells created by a GridLayout manager.
7 */
8
9 public class GridPanelWindow extends JFrame
10 {
11 private final int WINDOW_WIDTH = 400; // Window width
12 private final int WINDOW_HEIGHT = 200; // Window height
13
14 /**
15 * Constructor
16 */
17
18 public GridPanelWindow()
19 {
20 // Set the title bar text.
21 setTitle("Grid Layout");
22
23 // Set the size of the window.
24 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
25
26 // Specify an action for the close button.
27 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
28
29 // Add a GridLayout manager to the content pane.
30 setLayout(new GridLayout(2, 3));
31
32 // Create six buttons.
33 JButton button1 = new JButton("Button 1");
34 JButton button2 = new JButton("Button 2");
35 JButton button3 = new JButton("Button 3");
36 JButton button4 = new JButton("Button 4");
37 JButton button5 = new JButton("Button 5");
38 JButton button6 = new JButton("Button 6");
39
40 // Create six labels.
41 JLabel label1 = new JLabel("This is cell 1.");
42 JLabel label2 = new JLabel("This is cell 2.");
43 JLabel label3 = new JLabel("This is cell 3.");
44 JLabel label4 = new JLabel("This is cell 4.");
45 JLabel label5 = new JLabel("This is cell 5.");
46 JLabel label6 = new JLabel("This is cell 6.");
47
48 // Create six panels.
49 JPanel panel1 = new JPanel();
50 JPanel panel2 = new JPanel();
51 JPanel panel3 = new JPanel();
52 JPanel panel4 = new JPanel();
53 JPanel panel5 = new JPanel();
54 JPanel panel6 = new JPanel();
55
56 // Add the labels to the panels.
57 panel1.add(label1);
58 panel2.add(label2);
59 panel3.add(label3);
60 panel4.add(label4);
61 panel5.add(label5);

```

```

62 panel6.add(label6);
63
64 // Add the buttons to the panels.
65 panel1.add(button1);
66 panel2.add(button2);
67 panel3.add(button3);
68 panel4.add(button4);
69 panel5.add(button5);
70 panel6.add(button6);
71
72 // Add the panels to the content pane.
73 add(panel1); // Goes into row 1, column 1
74 add(panel2); // Goes into row 1, column 2
75 add(panel3); // Goes into row 1, column 3
76 add(panel4); // Goes into row 2, column 1
77 add(panel5); // Goes into row 2, column 2
78 add(panel6); // Goes into row 2, column 3
79
80 // Display the window.
81 setVisible(true);
82 }
83
84 /**
85 * The main method creates an instance of the
86 * GridPanelWindow class, displaying its window.
87 */
88
89 public static void main(String[] args)
90 {
91 new GridPanelWindow();
92 }
93 }

```

Because we have containers nested inside the content pane, there are multiple layout managers at work in the `GridPanelWindow` class. The content pane uses a `GridLayout` manager, and each of the `JPanel` objects uses a `FlowLayout` manager.



MyProgrammingLab™ [www.myprogramminglab.com](http://www.myprogramminglab.com)

1. 23.10 How do you add a layout manager to a container?
2. 23.11 Which layout manager divides a container into regions known as north, south, east, west, and center?
3. 23.12 Which layout manager arranges components in a row, from left to right, in the order they were added to the container?
4. 23.13 Which layout manager arranges components in rows and columns?



5. 23.14 How many components can you have at one time in a `BorderLayout` region? In a `GridLayout` cell?
6. 23.15 How do you prevent the `BorderLayout` manager from resizing a component that has been placed in its region?
7. 23.16 How can you cause a content pane to be automatically sized to accommodate the components contained within it?
8. 23.17 What is the default layout manager for a `JFrame` object's content pane? For a `JPanel` object?

# 23.4 Radio Buttons and Check Boxes

CONCEPT: Radio buttons normally appear in groups of two or more and allow the user to select one of several possible options. Check boxes, which may appear alone or in groups, allow the user to make yes/no or on/off selections.

## Radio Buttons

*Radio buttons* are useful when you want the user to select one choice from several possible options. [Figure 23-24](#) shows a group of radio buttons.




**Figure 23-24 Radio buttons**

(Oracle Corporate Counsel)

A radio button may be selected or deselected. Each radio button has a small circle that appears filled in when the radio button is selected and appears empty when the radio button is deselected. You use the `JRadioButton` class to create radio buttons. Here are the general formats of two `JRadioButton` constructors:

```
JRadioButton(String text)
JRadioButton(String text, boolean selected)
```

The first constructor shown creates a deselected radio button. The argument passed to the `text` parameter is the string that is displayed next to the radio button. For example, the following statement creates a radio button with the text “Choice 1” displayed next to it. The radio button initially appears deselected.

```
JRadioButton radio1 = new JRadioButton("Choice 1");
```

The second constructor takes an additional `boolean` argument, which is passed to the `selected` parameter. If `true` is passed as the `selected` argument, the radio button initially appears selected. If `false` is passed, the radio button initially appears deselected. For example, the following statement creates a radio button with the text “Choice 1” displayed next to it. The radio button initially appears selected.

```
JRadioButton radio1 = new JRadioButton("Choice 1", true);
```

Radio buttons are normally grouped together. When a set of radio buttons are grouped together, only one of the radio buttons in the group may be selected at any time. Clicking a radio button selects it and automatically deselects any other radio button in the same group. Because only one radio button in a group can be selected at any given time, the buttons are said to be *mutually exclusive*.



The name “radio button” refers to the old car radios that had push buttons for selecting stations. Only one of the buttons could be pushed in at a time. When you pushed a button in, it automatically popped out any other button that was pushed in.

## Grouping with the ButtonGroup class

Once you have created the `JRadioButton` objects that you wish to appear in a group, you must create an instance of the `ButtonGroup` class, and then add the `JRadioButton` objects to it. The `ButtonGroup` object creates the mutually exclusive relationship among the radio buttons that it contains. The following code shows an example:

```
// Create three radio buttons.
JRadioButton radio1 = new JRadioButton("Choice 1", true);
JRadioButton radio2 = new JRadioButton("Choice 2");
JRadioButton radio3 = new JRadioButton("Choice 3");

// Create a ButtonGroup object.
ButtonGroup group = new ButtonGroup();

// Add the radio buttons to the ButtonGroup object.
group.add(radio1);
group.add(radio2);
group.add(radio3);
```

Although you add radio buttons to a `ButtonGroup` object, `ButtonGroup` objects are not containers like `JPanel` objects, or content frames. The function of a `ButtonGroup` object is to deselect all the other radio buttons when one of them is selected. If you wish to add the radio buttons to a panel or a content frame, you must add them individually, as shown here:


```
// Add the radio buttons to the JPanel referenced by panel.
panel.add(radio1);
panel.add(radio2);
panel.add(radio3);
```

## Responding to Radio Button Events

Just like `JButton` objects, `JRadioButton` objects generate an action event when they are clicked. To respond to a radio button action event, you must write an action listener class and then register an instance of that class with the `JRadioButton` object. To demonstrate, we will look at the `MetricConverter` class, which is similar to the `KiloConverter` class shown earlier. The `MetricConverter` class presents a window in which the user can enter a distance in kilometers, and then click radio buttons to see that distance converted to miles, feet, or inches. The conversion formulas are as follows:

$$\text{Miles} = \text{Kilometers} \times 0.6214 \quad \text{Feet} = \text{Kilometers} \times 3281.0 \quad \text{Inches} = \text{Kilometers} \times 39370.0$$

[Figure 23-25](#) shows a sketch of what the window will look like. As you can see from the sketch, the window will have a label, a text field, and three radio buttons. When the user clicks on one of the radio buttons, the distance will be converted to the selected units and displayed in a separate `JOptionPane` dialog box.




## Figure 23-25 Metric Converter window

(Oracle Corporate Counsel)


[Figure 23-25 Full Alternative Text](#)

The `MetricConverter` class is shown in [Code Listing 23-14](#). The class initially displays the window shown at the top of [Figure 23-26](#). The figure also shows the dialog boxes that are displayed when the user clicks any of the radio buttons.


This window appears first. The user enters 2 into the text field.




This dialog box appears when the user clicks the "Convert to miles" radio button.



This dialog box appears when the user clicks the "Convert to feet" radio button.



This dialog box appears when the user clicks the "Convert to inches" radio button.



## Figure 23-26 Window and dialog boxes displayed by the `MetricConverter` class

(Oracle Corporate Counsel)

[Figure 23-26 Full Alternative Text](#)

## Code Listing 23-14 (`MetricConverter.java`)

```
1 import javax.swing.*;
2 import java.awt.event.*;
3 import java.awt.*;
4
5 /**
6 * The MetricConverter class lets the user enter a
7 * distance in kilometers. Radio buttons can be selected to
8 * convert the kilometers to miles, feet, or inches.
9 */
10
```

```

11 public class MetricConverter extends JFrame
12 {
13 private JPanel panel; // A holding panel
14 private JLabel messageLabel; // A message to the user
15 private JTextField kiloTextField; // To hold user input
16 private JRadioButton milesButton; // To convert to miles
17 private JRadioButton feetButton; // To convert to feet
18 private JRadioButton inchesButton; // To convert to inches
19 private ButtonGroup radioButtonGroup; // To group radio buttons
20 private final int WINDOW_WIDTH = 400; // Window width
21 private final int WINDOW_HEIGHT = 100; // Window height
22
23 /**
24 * Constructor
25 */
26
27 public MetricConverter()
28 {
29 // Set the title.
30 setTitle("Metric Converter");
31
32 // Set the size of the window.
33 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
34
35 // Specify an action for the close button.
36 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
37
38 // Build the panel and add it to the frame.
39 buildPanel();
40
41 // Add the panel to the frame's content pane.
42 add(panel);
43
44 // Display the window.
45 setVisible(true);
46 }
47
48 /**
49 * The buildPanel method adds a label, text field, and
50 * and three buttons to a panel.
51 */
52
53 private void buildPanel()
54 {
55 // Create the label, text field, and radio buttons.
56 messageLabel = new JLabel("Enter a distance in kilometers");
57 kiloTextField = new JTextField(10);
58 milesButton = new JRadioButton("Convert to miles");
59 feetButton = new JRadioButton("Convert to feet");
60 inchesButton = new JRadioButton("Convert to inches");
61
62 // Group the radio buttons.
63 radioButtonGroup = new ButtonGroup();
64 radioButtonGroup.add(milesButton);
65 radioButtonGroup.add(feetButton);
66 radioButtonGroup.add(inchesButton);
67

```

```

68 // Add action listeners to the radio buttons.
69 milesButton.addActionListener(new RadioButtonListener());
70 feetButton.addActionListener(new RadioButtonListener());
71 inchesButton.addActionListener(new RadioButtonListener());
72
73 // Create a panel and add the components to it.
74 panel = new JPanel();
75 panel.add(messageLabel);
76 panel.add(kiloTextField);
77 panel.add(milesButton);
78 panel.add(feetButton);
79 panel.add(inchesButton);
80 }
81
82 /**
83 Private inner class that handles the event when
84 the user clicks one of the radio buttons.
85 */
86
87 private class RadioButtonListener implements ActionListener
88 {
89 public void actionPerformed(ActionEvent e)
90 {
91 String input; // To hold the user's input
92 String convertTo = ""; // The units we're converting to
93 double result = 0.0; // To hold the conversion
94
95 // Get the kilometers entered.
96 input = kiloTextField.getText();
97
98 // Determine which radio button was clicked.
99 if (e.getSource() == milesButton)
100 {
101 // Convert to miles.
102 convertTo = " miles.";
103 result = Double.parseDouble(input) * 0.6214;
104 }
105 else if (e.getSource() == feetButton)
106 {
107 // Convert to feet.
108 convertTo = " feet.";
109 result = Double.parseDouble(input) * 3281.0;
110 }
111 else if (e.getSource() == inchesButton)
112 {
113 // Convert to inches.
114 convertTo = " inches.";
115 result = Double.parseDouble(input) * 39370.0;
116 }
117
118 // Display the conversion.
119 JOptionPane.showMessageDialog(null, input +
120 " kilometers is " + result + convertTo);
121 }
122 }
123
124 /**

```

```

125 The main method creates an instance of the
126 MetricConverter class, displaying its window.
127 */
128
129 public static void main(String[] args)
130 {
131 new MetricConverter();
132 }
133 }

```

# Determining in Code Whether a Radio Button Is Selected

In many applications you will merely want to know whether a radio button is selected. The `JRadioButton` class's `isSelected` method returns a boolean value indicating whether the radio button is selected. If the radio button is selected, the method returns `true`. Otherwise, it returns `false`. In the following code, the `radio` variable references a radio button. The `if` statement calls the `isSelected` method to determine whether the radio button is selected.

```

if (radio.isSelected())
{
 // Code here executes if the radio
 // button is selected.
}

```

# Selecting a Radio Button in Code

It is also possible to select a radio button in code with the `JRadioButton` class's `doClick` method. When the method is called, the radio button is selected just as if the user had clicked on it. As a result, an action event is generated. In the following statement, the `radio` variable references a radio button. When this statement executes, the radio button will be selected.

```

radio.doClick();

```

# Check Boxes

A *check box* appears as a small box with a label appearing next to it. The window shown in [Figure 23-27](#) has three check boxes.





## Figure 23-27 Check boxes

(Oracle Corporate Counsel)

Like radio buttons, check boxes may be selected or deselected at run time. When a check box is selected, a small check mark appears inside the box. Although check boxes are often displayed in groups, they are not usually grouped in a `ButtonGroup` like radio buttons. This is because check boxes are not normally used to make mutually exclusive selections. Instead, the user is allowed to select any or all of the check boxes that are displayed in a group.

You create a check box with the `JCheckBox` class. Here are the general formats of two `JCheckBox` constructors:

```
JCheckBox(String text)
JCheckBox(String text, boolean selected)
```

The first constructor shown creates a deselected check box. The argument passed to the `text` parameter is the string that is displayed next to the check box. For example, the following statement creates a check box with the text “Macaroni” displayed next to it. The check box initially appears deselected.

```
JCheckBox check1 = new JCheckBox("Macaroni");
```

The second constructor takes an additional `boolean` argument, which is passed to the `selected` parameter. If `true` is passed as the `selected` argument, the radio check box initially appears selected. If `false` is passed, the check box initially appears deselected. For example, the following statement creates a check box with the text “Macaroni” displayed next to it. The radio check box initially appears selected.

```
JCheckBox check1 = new JCheckBox("Macaroni", true);
```

## Responding to Check Box Events

When a `JCheckBox` object is selected or deselected, it generates an *item event*. You handle item events in a manner similar to the way you handle the action events that are generated by `JButton` and `JRadioButton` objects. First, you write an *item listener* class, which must meet the following requirements:

- It must implement the `ItemListener` interface.
- It must have a method named `itemStateChanged` with the following header:

```
public void itemStateChanged(ItemEvent e)
```



## NOTE:

When implementing the `ItemListener` interface, your code must have the following import statement: `import java.awt.event.*;`

Once you have written an item listener class, you create an object of that class, and then register the item listener object with the `JCheckBox` component. When a `JCheckBox` component generates an event, it automatically executes the `itemStateChanged` method of the item listener object that is registered to it, passing the event object as an argument.

# Determining in Code Whether a Check Box Is Selected

As with `JRadioButton`, you use the `isSelected` method to determine whether a `JCheckBox` component is selected. The method returns a `boolean` value. If the check box is selected, the method returns `true`. Otherwise, it returns `false`. In the following code, the `checkBox` variable references a `JCheckBox` component. The `if` statement calls the `isSelected` method to determine whether the check box is selected.

```
if (checkBox.isSelected())
{
 // Code here executes if the check
 // box is selected.
}
```

The `ColorCheckBoxWindow` class, shown in [Code Listing 23-15](#), demonstrates how check boxes are used. It displays the window shown in [Figure 23-28](#). When the “Yellow background” check box is selected, the background color of the content pane, the label, and the check boxes turns yellow. When this check box is deselected, the background colors go back to light gray. When the “Red foreground” check box is selected, the color of the text displayed in the label and the check boxes turns red. When this check box is deselected, the foreground colors go back to black.




**Figure 23-28 Window displayed by the `ColorCheckBoxWindow` class**

(Oracle Corporate Counsel)

[Figure 23-28 Full Alternative Text](#)

## Code Listing 23-15

(`ColorCheckBoxWindow.java`)

```
1 import javax.swing.*;
2 import java.awt.*;
3 import java.awt.event.*;
4
5 /**
6 * The ColorCheckBoxWindow class demonstrates how check boxes
7 * can be used.
8 */
9
10 public class ColorCheckBoxWindow extends JFrame
11 {
12 private JLabel messageLabel; // A message to the user
13 private JCheckBox yellowCheckBox; // To select yellow background
14 private JCheckBox redCheckBox; // To select red foreground
15 private final int WINDOW_WIDTH = 300; // Window width
16 private final int WINDOW_HEIGHT = 100; // Window height
17
18 /**
19 * Constructor
20 */
21
22 public ColorCheckBoxWindow()
23 {
24 // Set the text for the title bar.
25 setTitle("Color Check Boxes");
26
27 // Set the size of the window.
28 setSize(WINDOW_WIDTH, WINDOW_HEIGHT);
29
30 // Specify an action for the close button.
31 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
```

```

32
33 // Create a label.
34 messageLabel = new JLabel("Select the check " +
35 "boxes to change colors.");
36
37 // Create the check boxes.
38 yellowCheckBox = new JCheckBox("Yellow background");
39 redCheckBox = new JCheckBox("Red foreground");
40
41 // Add an item listener to the check boxes.
42 yellowCheckBox.addItemListener(new CheckBoxListener());
43 redCheckBox.addItemListener(new CheckBoxListener());
44
45 // Add a FlowLayout manager to the content pane.
46 setLayout(new FlowLayout());
47
48 // Add the label and check boxes to the content pane.
49 add(messageLabel);
50 add(yellowCheckBox);
51 add(redCheckBox);
52
53 // Display the window.
54 setVisible(true);
55 }
56
57 /**
58 Private inner class that handles the event when
59 the user clicks one of the check boxes.
60 */
61
62 private class CheckBoxListener implements ItemListener
63 {
64 public void itemStateChanged(ItemEvent e)
65 {
66 // Determine which check box was clicked.
67 if (e.getSource() == yellowCheckBox)
68 {
69 // Is the yellow check box selected? If so, we
70 // want to set the background color to yellow.
71 if (yellowCheckBox.isSelected())
72 {
73 // The yellow check box was selected. Set
74 // the background color for the content
75 // pane and the two check boxes to yellow.
76 getContentPane().setBackground(Color.YELLOW);
77 yellowCheckBox.setBackground(Color.YELLOW);
78 redCheckBox.setBackground(Color.YELLOW);
79 }
80 else
81 {
82 // The yellow check box was deselected. Set
83 // the background color for the content
84 // pane and the two check boxes to light gray.
85 getContentPane().setBackground(Color.LIGHT_GRAY);
86 yellowCheckBox.setBackground(Color.LIGHT_GRAY);
87 redCheckBox.setBackground(Color.LIGHT_GRAY);
88 }
89 }
90 }
91 }

```

```

89 }
90 else if (e.getSource() == redCheckBox)
91 {
92 // Is the red check box selected? If so, we want
93 // to set the foreground color to red.
94 if (redCheckBox.isSelected())
95 {
96 // The red check box was selected. Set the
97 // foreground color for the label and the
98 // two check boxes to red.
99 messageLabel.setForeground(Color.RED);
100 yellowCheckBox.setForeground(Color.RED);
101 redCheckBox.setForeground(Color.RED);
102 }
103 else
104 {
105 // The red check box was deselected. Set the
106 // foreground color for the label and the
107 // two check boxes to black.
108 messageLabel.setForeground(Color.BLACK);
109 yellowCheckBox.setForeground(Color.BLACK);
110 redCheckBox.setForeground(Color.BLACK);
111 }
112 }
113 }
114 }
115
116 /**
117 The main method creates an instance of the
118 ColorCheckBoxWindow class, displaying its window.
119 */
120
121 public static void main(String[] args)
122 {
123 new ColorCheckBoxWindow();
124 }
125 }

```

## Selecting a Check Box in Code

As with radio buttons, it is possible to select check boxes in code with the `JCheckBox` class's `doClick` method. When the method is called, the radio check box is selected just as if the user had clicked on it. As a result, an item event is generated. In the following statement, the `checkBox` variable references a `JCheckBox` object. When this statement executes, the check box will be selected.

```
checkBox.doClick();
```



1. 23.18 You want the user to be able to select only one item from a group of items. Which type of component would you use for the items, radio buttons or check boxes?
2. 23.19 You want the user to be able to select any number of items from a group of items. Which type of component would you use for the items, radio buttons or check boxes?
3. 23.20 What is the purpose of a `ButtonGroup` object?
4. 23.21 Do you normally add radio buttons, check boxes, or both to a `ButtonGroup` object?
5. 23.22 What type of event does a radio button generate when the user clicks on it?
6. 23.23 What type of event does a check box generate when the user clicks on it?
7. 23.24 How do you determine in code whether a radio button is selected?
8. 23.25 How do you determine in code whether a check box is selected?

# 23.5 Borders

CONCEPT: A component can appear with several different styles of borders around it. A `Border` object specifies the details of a border. You use the `BorderFactory` class to create `Border` objects.

Sometimes it is helpful to place a border around a component or a group of components on a panel. You can give windows a more organized look by grouping related components inside borders. For example, [Figure 23-29](#) shows a group of check boxes that are enclosed in a border. In addition, notice that the border has a title.




**Figure 23-29 A group of check boxes with a titled border**

`JPanel` components have a method named `setBorder`, which is used to add a border to the panel. The `setBorder` method accepts a `Border` object as its argument. A `Border` object contains detailed information describing the appearance of a border.

Rather than creating `Border` objects yourself, you should use the `BorderFactory` class to create them for you. The `BorderFactory` class has methods that return various types of borders. [Table 23-6](#) describes borders that can be created with the `BorderFactory` class. The table also lists the `BorderFactory` methods that can be called to create the borders. Note that there are several overloaded versions of each method.

**Table 23-6 Borders produced by the `BorderFactory` class**

| Border   | BorderFactory Method              | Description                                                                                   |
|----------|-----------------------------------|-----------------------------------------------------------------------------------------------|
| Compound | <code>createCompoundBorder</code> | A border that has two parts: an inside edge and an outside edge. The inside and outside edges |

|                      |                                       |                                                                                                                                           |
|----------------------|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| border               |                                       | can be any of the other borders.                                                                                                          |
| Empty border         | <code>createEmptyBorder</code>        | A border that contains only empty space.                                                                                                  |
| Etched border        | <code>createEtchedBorder</code>       | A border with a 3-D appearance that looks “etched” into the background.                                                                   |
| Line border          | <code>createLineBorder</code>         | A border that appears as a line.                                                                                                          |
| Lowered bevel border | <code>createLoweredBevelBorder</code> | A border that looks like beveled edges. It has a 3-D appearance that gives the illusion of being sunken into the surrounding background.  |
| Matte border         | <code>createMatteBorder</code>        | A line border that can have edges of different thicknesses.                                                                               |
| Raised bevel border  | <code>createRaisedBevelBorder</code>  | A border that looks like beveled edges. It has a 3-D appearance that gives the illusion of being raised above the surrounding background. |
| Titled border        | <code>createTitledBorder</code>       | An etched border with a title.                                                                                                            |

## NOTE:

If you use the `BorderFactory` class in your code, you should have the following `import` statement:  
`import javax.swing.*;`

In this chapter, we will concentrate on empty borders, line borders, and titled borders.

## Empty Borders



An empty border is simply empty space around the edges of a component. To create an empty border, call the `BorderFactory` class's `createEmptyBorder` method. Here is the method's general format:

```
BorderFactory.createEmptyBorder(int top, int left,
 int bottom, int right);
```

The arguments passed into `top`, `left`, `bottom`, and `right` specify in pixels the size of the border's top, left, bottom, and right edges. The method returns a reference to a `Border` object. The following is an example of a statement that uses the method. Assume that the `panel` variable references a `JPanel` object.

```
panel.setBorder(BorderFactory.createEmptyBorder(5, 5, 5, 5));
```

After this statement executes, the `JPanel` referenced by `panel` will have an empty border of five pixels around each edge.



## NOTE:

In case you've skipped ahead to this chapter, the `BorderFactory` methods are static, which means that you call them without creating an instance of the `BorderFactory` class. (You simply write *`BorderFactory`*. before the method name to call the method.) This is similar to the way the `Math` class and wrapper class methods we have discussed are called. Static methods are covered in [Chapter 8](#).

## Line Borders

A line border is a line of a specified color and thickness that appears around the edges of a component. To create a line border, call the `BorderFactory` class's `createLineBorder` method. Here is the method's general format:

```
BorderFactory.createLineBorder(Color color, int thickness);
```

The arguments passed into `color` and `thickness` specify the color of the line and the size of the line in pixels. The method returns a reference to a `Border` object. The following is an example of a statement that uses the method. Assume that the `panel` variable references a `JPanel` object.

```
panel.setBorder(BorderFactory.createLineBorder(Color.RED, 1));
```

After this statement executes, the `JPanel` referenced by `panel` will have a red line border that is one pixel thick around its edges.

## Titled Borders

A titled border is an etched border with a title displayed on it. To create a titled border, call the `BorderFactory` class's `createTitledBorder` method. Here is the method's general format:

```
BorderFactory.createTitledBorder(String title);
```

The argument passed into `title` is the text that is to be displayed as the border's title. The method returns a reference to a `Border` object. The following is an example of a statement that uses the method. Assume that the `panel` variable references a `JPanel` object.

```
panel.setBorder(BorderFactory.createTitledBorder("Choices"));
```

After this statement executes, the `JPanel` referenced by `panel` will have an etched border with the title "Choices" displayed on it.

## Checkpoint

MyProgrammingLab™ [www.myprogramminglab.com](http://www.myprogramminglab.com)

1. 23.26 What method do you use to set a border around a component?
2. 23.27 What is the preferred way of creating a `Border` object?

## 23.6 Focus on Problem Solving: Extending Classes from JPanel

**CONCEPT:** By writing a class that is extended from the `JPanel` class, you can create a custom panel component that can hold other components and their related code.

In the applications that you have studied so far in this chapter, we have used the `extends JFrame` clause in the class header to extend the class from the `JFrame` class. Recall that the extended class is then a specialized version of the `JFrame` class, and we use its constructor to create the panels, buttons, and all of the other components needed. This approach works well for simple applications. But for applications that use many components, this approach can be cumbersome. Bundling all of the code and event listeners for a large number of components into a single class can lead to a large and complex class. A better approach is to encapsulate smaller groups of related components and their event listeners into their own classes.

A commonly used technique is to extend a class from the `JPanel` class. This allows you to create your own specialized panel component, which can contain other components and related code such as event listeners. A complex application that uses numerous components can be constructed from several specialized panel components. In this section we will examine such an application.

### The Brandi's Bagel House Application

Brandi's Bagel House has a bagel and coffee delivery service for the businesses in her neighborhood. Customers may call in and order white and whole wheat bagels with a variety of toppings. In addition, customers may order three different types of coffee. (Delivery for coffee alone is not available, however.) Here is a complete price list:


*Bagels: White bagel \$1.25, whole wheat bagel \$1.50*

*Toppings: Cream cheese \$0.50, butter \$0.25, peach jelly \$0.75, blueberry jam \$0.75*

*Coffee: Regular coffee \$1.25, decaf coffee \$1.25, cappuccino \$2.00*

Brandi, the owner, needs an “order calculator” application that her staff can use to calculate the price of an order as it is called in. The application should display the subtotal, the amount of a 6 percent sales tax, and the total of the order. [Figure 23-30](#) shows a sketch of the application's window. The user selects the type of bagel, toppings, and coffee, then clicks the Calculate button. A dialog box

appears displaying the subtotal, amount of sales tax, and total. The user can exit the application by clicking either the Exit button or the standard close button in the upper-right corner.



## Figure 23-30 Sketch of the Order Calculator window

(Oracle Corporate Counsel)

[Figure 23-30 Full Alternative Text](#)

The layout shown in the sketch can be achieved using a `BorderLayout` manager with the window's content pane. The label that displays "Welcome to Brandi's Bagel House" is in the north region, the radio buttons for the bagel types are in the west region, the check boxes for the toppings are in the center region, the radio buttons for the coffee selection are in the east region, and the Calculate and Exit buttons are in the south region. To construct this window, we create the following specialized panel classes that are extended from `JPanel`:

- `GreetingsPanel`. This panel contains the label that appears in the window's north region.
- `BagelPanel`. This panel contains the radio buttons for the types of bagels.
- `ToppingPanel`. This panel contains the check boxes for the types of bagels.
- `CoffeePanel`. This panel contains the radio buttons for the coffee selections.

(We will not create a specialized panel for the Calculate and Exit buttons. The reason is explained later.) After these classes have been created, we can create objects from them and add the objects to the correct regions of the window's content pane. Let's take a closer look at each of these classes.

## The GreetingPanel Class

The `GreetingPanel` class holds the label displaying the text “Welcome to Brandi’s Bagel House”. [Code Listing 23-16](#) shows the class, which extends `JPanel`.

## Code Listing 23-16 (`GreetingPanel.java`)

```
1 import javax.swing.*;
2
3 /**
4 The GreetingPanel class displays a greeting in a panel.
5 */
6
7 public class GreetingPanel extends JPanel
8 {
9 private JLabel greeting; // To display a greeting
10
11 /**
12 Constructor
13 */
14
15 public GreetingPanel()
16 {
17 // Create the label.
18 greeting = new JLabel("Welcome to Brandi's Bagel House");
19
20 // Add the label to this panel.
21 add(greeting);
22 }
23 }
```

In line 21 the `add` method is called to add the `JLabel` component referenced by `greeting`. Notice that we are calling the method without an object reference and a dot preceding it. This is because the method was inherited from the `JPanel` class, and we can call it just as if it were written into the `GreetingPanel` class declaration.

When we create an instance of this class, we are creating a `JPanel` component that displays a label with the text “Welcome to Brandi’s Bagel House”. [Figure 23-31](#) shows how the component will appear when it is placed in the window’s north region.



Welcome to Brandi's Bagel House

**Figure 23-31 Appearance of the `GreetingPanel` component**

## The `BagelPanel` Class

The `BagelPanel` class holds the radio buttons for the types of bagels. Notice that this panel uses a `GridLayout` manager with two rows and one column. [Code Listing 23-17](#) shows the class, which is extended from `JPanel`.

## Code Listing 23-17 (`BagelPanel.java`)

```
1 import javax.swing.*;
2 import java.awt.*;
3
4 /**
5 * The BagelPanel class allows the user to select either
6 * a white or whole wheat bagel.
7 */
8
9 public class BagelPanel extends JPanel
10 {
11 // The following constants are used to indicate
12 // the cost of each type of bagel.
13 public final double WHITE_BAGEL = 1.25;
14 public final double WHEAT_BAGEL = 1.50;
15
16 private JRadioButton whiteBagel; // To select white
17 private JRadioButton wheatBagel; // To select wheat
18 private ButtonGroup bg; // Radio button group
19
20 /**
21 * Constructor
22 */
23
24 public BagelPanel()
25 {
26 // Create a GridLayout manager with
27 // two rows and one column.
28 setLayout(new GridLayout(2, 1));
29
30 // Create the radio buttons.
31 whiteBagel = new JRadioButton("White", true);
32 wheatBagel = new JRadioButton("Wheat");
33
34 // Group the radio buttons.
35 bg = new ButtonGroup();
36 bg.add(whiteBagel);
37 bg.add(wheatBagel);
38
39 // Add a border around the panel.
40 setBorder(BorderFactory.createTitledBorder("Bagel"));
41
42 // Add the radio buttons to the panel.
43 add(whiteBagel);
44 add(wheatBagel);
45 }
46
47 /**
48 * getBagelCost method
```

```

49 @return The cost of the selected bagel.
50 */
51
52 public double getBagelCost()
53 {
54 double bagelCost = 0.0;
55
56 if (whiteBagel.isSelected())
57 bagelCost = WHITE_BAGEL;
58 else
59 bagelCost = WHEAT_BAGEL;
60
61 return bagelCost;
62 }
63 }

```

Notice that the `whiteBagel` radio button is automatically selected when it is created. This is the default choice. This class does not have an inner event listener class because we do not want to execute any code when the user selects a bagel. Instead, we want this class to be able to report the cost of the selected bagel. That is the purpose of the `getBagelCost` method, which returns the cost of the selected bagel as a `double`. (This method will be called by the Calculate button's event listener.) [Figure 23-32](#) shows how the component appears when it is placed in the window's west region.



## Figure 23-32 Appearance of the BagelPanel component

(Oracle Corporate Counsel)

## The ToppingPanel Class

The `ToppingPanel` class holds the check boxes for the available toppings. [Code Listing 23-18](#) shows the class, which is also extended from `JPanel`.

## Code Listing 23-18 (ToppingPanel.java)

```

1 import javax.swing.*;
2 import java.awt.*;

```

```

3
4 /**
5 The ToppingPanel class allows the user to select
6 the toppings for the bagel.
7 */
8
9 public class ToppingPanel extends JPanel
10 {
11 // The following constants are used to indicate
12 // the cost of toppings.
13 public final double CREAM_CHEESE = 0.50;
14 public final double BUTTER = 0.25;
15 public final double PEACH_JELLY = 0.75;
16 public final double BLUEBERRY_JAM = 0.75;
17
18 private JCheckBox creamCheese; // To select cream cheese
19 private JCheckBox butter; // To select butter
20 private JCheckBox peachJelly; // To select peach jelly
21 private JCheckBox blueberryJam; // To select blueberry jam
22
23 /**
24 Constructor
25 */
26
27 public ToppingPanel()
28 {
29 // Create a GridLayout manager with
30 // four rows and one column.
31 setLayout(new GridLayout(4, 1));
32
33 // Create the check boxes.
34 creamCheese = new JCheckBox("Cream cheese");
35 butter = new JCheckBox("Butter");
36 peachJelly = new JCheckBox("Peach jelly");
37 blueberryJam = new JCheckBox("Blueberry jam");
38
39 // Add a border around the panel.
40 setBorder(BorderFactory.createTitledBorder("Toppings"));
41
42 // Add the check boxes to the panel.
43 add(creamCheese);
44 add(butter);
45 add(peachJelly);
46 add(blueberryJam);
47 }
48
49 /**
50 getToppingCost method
51 @return The cost of the selected toppings.
52 */
53
54 public double getToppingCost()
55 {
56 double toppingCost = 0.0;
57
58 if (creamCheese.isSelected())
59 toppingCost += CREAM_CHEESE;
60 if (butter.isSelected())

```



```

61 toppingCost += BUTTER;
62 if (peachJelly.isSelected())
63 toppingCost += PEACH_JELLY;
64 if (blueberryJam.isSelected())
65 toppingCost += BLUEBERRY_JAM;
66
67 return toppingCost;
68 }
69 }

```

As with the `BagelPanel` class, this class does not have an inner event listener class because we do not want to execute any code when the user selects a topping. Instead, we want this class to be able to report the total cost of all the selected toppings. That is the purpose of the `getToppingCost` method, which returns the cost of all the selected toppings as a `double`. (This method will be called by the Calculate button's event listener.) [Figure 23-33](#) shows how the component appears when it is placed in the window's center region.



## Figure 23-33 Appearance of the ToppingPanel component

(Oracle Corporate Counsel)

## The CoffeePanel Class

The `CoffeePanel` class holds the radio buttons for the available coffee selections. [Code Listing 23-19](#) shows the class, which extends `JPanel`.

### Code Listing 23-19 (`CoffeePanel.java`)

```

1 import javax.swing.*;
2 import java.awt.*;
3
4 /**
5 * The CoffeePanel class allows the user to select coffee.
6 */
7
8 public class CoffeePanel extends JPanel

```

```

9 {
10 // The following constants are used to indicate
11 // the cost of coffee.
12 public final double NO_COFFEE = 0.0;
13 public final double REGULAR_COFFEE = 1.25;
14 public final double DECAF_COFFEE = 1.25;
15 public final double CAPPUCCINO = 2.00;
16
17 private JRadioButton noCoffee; // To select no coffee
18 private JRadioButton regularCoffee; // To select regular coffee
19 private JRadioButton decafCoffee; // To select decaf
20 private JRadioButton cappuccino; // To select cappuccino
21 private ButtonGroup bg; // Radio button group
22
23 /**
24 * Constructor
25 */
26
27 public CoffeePanel()
28 {
29 // Create a GridLayout manager with
30 // four rows and one column.
31 setLayout(new GridLayout(4, 1));
32
33 // Create the radio buttons.
34 noCoffee = new JRadioButton("None");
35 regularCoffee = new JRadioButton("Regular coffee", true);
36 decafCoffee = new JRadioButton("Decaf coffee");
37 cappuccino = new JRadioButton("Cappuccino");
38
39 // Group the radio buttons.
40 bg = new ButtonGroup();
41 bg.add(noCoffee);
42 bg.add(regularCoffee);
43 bg.add(decafCoffee);
44 bg.add(cappuccino);
45
46 // Add a border around the panel.
47 setBorder(BorderFactory.createTitledBorder("Coffee"));
48
49 // Add the radio buttons to the panel.
50 add(noCoffee);
51 add(regularCoffee);
52 add(decafCoffee);
53 add(cappuccino);
54 }
55
56 /**
57 * getCoffeeCost method
58 * @return The cost of the selected coffee.
59 */
60
61 public double getCoffeeCost()
62 {
63 double coffeeCost = 0.0;
64
65 if (noCoffee.isSelected())
66 coffeeCost = NO_COFFEE;

```

```

67 else if (regularCoffee.isSelected())
68 coffeeCost = REGULAR_COFFEE;
69 else if (decafCoffee.isSelected())
70 coffeeCost = DECAF_COFFEE;
71 else if (cappuccino.isSelected())
72 coffeeCost = CAPPUCCINO;
73
74 return coffeeCost;
75 }
76 }

```

As with the `BagelPanel` and `ToppingPanel` classes, this class does not have an inner event listener class because we do not want to execute any code when the user selects coffee. Instead, we want this class to be able to report the cost of the selected coffee. The `getCoffeeCost` method returns the cost of the selected coffee as a `double`. (This method will be called by the Calculate button's event listener.) [Figure 23-34](#) shows how the component appears when it is placed in the window's east region.





**Figure 23-34 Appearance of the `CoffeePanel` component**

## Putting It All Together

The last step in creating this application is to write a class that builds the application's window and adds the Calculate and Exit buttons. This class, which we name `OrderCalculatorGUI`, is extended from `JFrame` and uses a `BorderLayout` manager with its content pane. [Figure 23-35](#) shows how instances of the `GreetingPanel`, `BagelPanel`, `ToppingPanel`, and `CoffeePanel` classes are placed in the content pane.



## Figure 23-35 Placement of the custom panels

### [Figure 23-35 Full Alternative Text](#)

We have not created a custom panel class to hold the Calculate and Exit buttons. The reason is that the Calculate button's event listener must call the `getBagelCost`, `getToppingCost`, and `getCoffeeCost` methods. In order to call those methods, the event listener must have access to the `BagelPanel`, `ToppingPanel`, and `CoffeePanel` objects that are created in the `OrderCalculatorGUI` class. The approach taken in this example is to have the `OrderCalculatorGUI` class itself create the buttons. The code for the `OrderCalculatorGUI` class is shown in [Code Listing 23-20](#).

## Code Listing 23-20 (OrderCalculatorGUI.java)

```

1 import javax.swing.*;
2 import java.awt.*;
3 import java.awt.event.*;
4
5 /**
6 The OrderCalculatorGUI class creates the GUI for the
7 Brandi's Bagel House application.
8 */
9
10 public class OrderCalculatorGUI extends JFrame
11 {

```

```

12 private BagelPanel bagels; // Bagel panel
13 private ToppingPanel toppings; // Topping panel
14 private CoffeePanel coffee; // Coffee panel
15 private GreetingPanel banner; // To display a greeting
16 private JPanel buttonPanel; // To hold the buttons
17 private JButton calcButton; // To calculate the cost
18 private JButton exitButton; // To exit the application
19 private final double TAX_RATE = 0.06; // Sales tax rate
20
21 /**
22 Constructor
23 */
24
25 public OrderCalculatorGUI()
26 {
27 // Display a title.
28 setTitle("Order Calculator");
29
30 // Specify an action for the close button.
31 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
32
33 // Create a BorderLayout manager.
34 setLayout(new BorderLayout());
35
36 // Create the custom panels.
37 banner = new GreetingPanel();
38 bagels = new BagelPanel();
39 toppings = new ToppingPanel();
40 coffee = new CoffeePanel();
41
42 // Create the button panel.
43 buildButtonPanel();
44
45 // Add the components to the content pane.
46 add(banner, BorderLayout.NORTH);
47 add(bagels, BorderLayout.WEST);
48 add(toppings, BorderLayout.CENTER);
49 add(coffee, BorderLayout.EAST);
50 add(buttonPanel, BorderLayout.SOUTH);
51
52 // Pack the contents of the window and display it.
53 pack();
54 setVisible(true);
55 }
56
57 /**
58 The buildButtonPanel method builds the button panel.
59 */
60
61 private void buildButtonPanel()
62 {
63 // Create a panel for the buttons.
64 buttonPanel = new JPanel();
65
66 // Create the buttons.
67 calcButton = new JButton("Calculate");
68 exitButton = new JButton("Exit");

```

```

69
70 // Register the action listeners.
71 calcButton.addActionListener(new CalcButtonListener());
72 exitButton.addActionListener(new ExitButtonListener());
73
74 // Add the buttons to the button panel.
75 buttonPanel.add(calcButton);
76 buttonPanel.add(exitButton);
77 }
78
79 /**
80 Private inner class that handles the event when
81 the user clicks the Calculate button.
82 */
83
84 private class CalcButtonListener implements ActionListener
85 {
86 public void actionPerformed(ActionEvent e)
87 {
88 // Variables to hold the subtotal, tax, and total
89 double subtotal, tax, total;
90
91 // Calculate the subtotal.
92 subtotal = bagels.getBagelCost() +
93 toppings.getToppingCost() +
94 coffee.getCoffeeCost();
95
96 // Calculate the sales tax.
97 tax = subtotal * TAX_RATE;
98
99 // Calculate the total.
100 total = subtotal + tax;
101
102 // Display the charges.
103 JOptionPane.showMessageDialog(null,
104 String.format("Subtotal: $%,.2f\n" +
105 "Tax: $%,.2f\n" +
106 "Total: $%,.2f",
107 subtotal, tax, total));
108 }
109 }
110
111 /**
112 Private inner class that handles the event when
113 the user clicks the Exit button.
114 */
115
116 private class ExitButtonListener implements ActionListener
117 {
118 public void actionPerformed(ActionEvent e)
119 {
120 System.exit(0);
121 }
122 }
123
124 /**
125 main method


```

```

126 */
127
128 public static void main(String[] args)
129 {
130 new OrderCalculatorGUI();
131 }
132 }


```

When the application runs, the window shown in [Figure 23-36](#) appears. [Figure 23-37](#) shows the `JOptionPane` dialog box that is displayed when the user selects a wheat bagel with butter, cream cheese, and decaf coffee.



## Figure 23-36 The Order Calculator window

(Oracle Corporate Counsel)



## Figure 23-37 The subtotal, tax, and total displayed

(Oracle Corporate Counsel)





# APPENDIX F More about the Math Class

The Java `Math` class is a collection of static methods for performing specific mathematical operations. This class is in the `java.lang` package, so there is no need for an `import` statement to use it.

In [Chapter 2](#), you were introduced to the `Math.pow` method, which returns the value of a number raised to a power. [Table F-1](#) describes several of the `Math` class's methods.

## Table F-1 Several `Math` class methods

| Method            | Example Usage                  | Description                                                                                                                                                                                                                |
|-------------------|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>abs</code>  | <code>y = Math.abs(x);</code>  | Returns the absolute value of the argument. This method can accept and return values of the <code>double</code> , <code>float</code> , <code>int</code> , and <code>long</code> data types.                                |
| <code>acos</code> | <code>y = Math.acos(x);</code> | Returns the arc-cosine of the argument. The argument should be the cosine of an angle. (The argument's value must be in the range from $-1.0$ through $1.0$ .) The return type and the argument are <code>doubles</code> . |
| <code>asin</code> | <code>y = Math.asin(x);</code> | Returns the arc-sine of the argument. The argument should be the sine of an angle. (The argument's value must be in the range from $-1.0$ through $1.0$ .) The return type and the argument are <code>doubles</code> .     |
| <code>atan</code> | <code>y = Math.atan(x);</code> | Returns the arc-tangent of the argument. The argument should be the tangent of an angle. The return type and the argument are <code>doubles</code> .                                                                       |
| <code>cbrt</code> | <code>y = Math.cbrt(x);</code> | Returns the cube root of the argument. The return type and argument are <code>doubles</code> . Note: This method was introduced in Java 5.                                                                                 |

|                    |                                    |                                                                                                                                                                                                                                                                                                                                         |
|--------------------|------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>ceil</code>  | <code>y = Math.ceil(x);</code>     | Returns the smallest number that is greater than or equal to the argument. The return type and the argument are <code>double</code> s.                                                                                                                                                                                                  |
| <code>cos</code>   | <code>y = Math.cos(x);</code>      | Returns the cosine of the argument. The argument should be an angle expressed in radians. The return type and the argument are <code>double</code> s.                                                                                                                                                                                   |
| <code>exp</code>   | <code>y = Math.exp(x);</code>      | Computes the exponential function of the argument, which is $e^x$ . The return type and the argument are <code>double</code> s.                                                                                                                                                                                                         |
| <code>floor</code> | <code>y = Math.floor(x);</code>    | Returns the largest number that is less than or equal to the argument. The return type and the argument are <code>double</code> s.                                                                                                                                                                                                      |
| <code>hypot</code> | <code>c = Math.hypot(a, b);</code> | Returns the length of the hypotenuse of a right triangle. The arguments <code>a</code> and <code>b</code> are the lengths of the other two sides of the triangle. The return type and the argument are <code>double</code> s. Note: This method was introduced in Java 5.                                                               |
| <code>log</code>   | <code>y = Math.log(x);</code>      | Returns the natural logarithm of the argument. The return type and the argument are <code>double</code> s.                                                                                                                                                                                                                              |
| <code>pow</code>   | <code>y = Math.pow(x, z);</code>   | Returns the value of the first argument raised to the power of the second argument. The return type and the argument are <code>double</code> s.                                                                                                                                                                                         |
| <code>round</code> | <code>y = Math.round(x);</code>    | Returns the value of the argument, as an integer, rounded to the nearest whole number. The argument is expected to be a <code>double</code> or a <code>float</code> . If the argument is a <code>double</code> , the return type is <code>long</code> . If the argument is a <code>float</code> , the return type is <code>int</code> . |
| <code>sin</code>   | <code>y = Math.sin(x);</code>      | Returns the sine of the argument. The argument should be an angle expressed in radians. The return type and                                                                                                                                                                                                                             |

the argument are `doubles`.

|                   |                                |                                                                                                  |
|-------------------|--------------------------------|--------------------------------------------------------------------------------------------------|
| <code>sqrt</code> | <code>y = Math.sqrt(x);</code> | Returns the square root of the argument. The return type and argument are <code>doubles</code> . |
|-------------------|--------------------------------|--------------------------------------------------------------------------------------------------|

|                  |                               |                                                                                                                                                        |
|------------------|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>tan</code> | <code>y = Math.tan(x);</code> | Returns the tangent of the argument. The argument should be an angle expressed in radians. The return type and the argument are <code>doubles</code> . |
|------------------|-------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|

|                        |                                         |                                                                                                                                                        |
|------------------------|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>toDegrees</code> | <code>y =<br/>Math.toDegrees(x);</code> | Accepts as an argument an angle in radians. The angle converted to degrees is returned. The argument and return values are both <code>doubles</code> . |
|------------------------|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|

|                        |                                         |                                                                                                                                                        |
|------------------------|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>toRadians</code> | <code>y =<br/>Math.toRadians(x);</code> | Accepts as an argument an angle in degrees. The angle converted to radians is returned. The argument and return values are both <code>doubles</code> . |
|------------------------|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|

We will not cover all of these methods in depth, but let's take a closer look at some of them. The `Math.abs` method returns the absolute value of its argument. The following code segment shows an example of how it is used.

```
double x = -4.2, y = 4.2, a, b;
a = Math.abs(x);
b = Math.abs(y);
```

After this code executes, both the variables `a` and `b` will contain the value 4.2. The argument to the `Math.abs` method can be of the `double`, `float`, `int`, or `long` data types. This is because there are several overloaded versions of the method. The method's return value will be of the same data type as the argument.

Here is a program segment that demonstrates the `Math.sqrt` method, which returns the square root of a number:

```
num = 25.0;
s = Math.sqrt(num);
System.out.println("The square root of " + num + " is " + s);
```

The output of this program segment is

```
The square root of 25.0 is 5.0
```

The `Math.round` method returns the value of its argument rounded to the nearest whole number. The following code segment shows an example.

```
double x = 4.2, y = 4.8;
long a, b;
a = Math.round(x);
b = Math.round(y);
```

After this code executes, `a` will contain 4 and `b` will contain 5.

In addition to these and other methods, the `Math` class defines two static `final` variables, which are listed in [Table F-2](#).

**Table F-2     The `E` and `PI` constants**

| Method          | Example Usage                                | Description                                                                                                             |
|-----------------|----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <code>E</code>  | <pre>y = x * Math.E;</pre>                   | This is the mathematical constant known as <code>e</code> (for Euler’s number). It is defined as 2.7182818284590452354. |
| <code>PI</code> | <pre>area = Math.PI * radius * radius;</pre> | This is the mathematical constant <code>Pi</code> , or $\pi$ . It is defined as 3.14159265358979323849.                 |

Both of these `final` variables are public, so you can access them directly from the class.

# APPENDIX G Packages

## NOTE:

To use this appendix, you must understand how your operating system uses directories, or folders. In addition, you must know how to set the value of an environment variable. The process of setting an environment variable differs among operating systems.

So far, you have been storing the classes you have created in the same folder or directory as the program that uses them. That is where the compiler looks for classes by default. In real-world application development, however, this approach is less than ideal. If you are developing more than one application that uses the same set of classes, you would have to make separate copies of the classes and store them in the same disk location as each application. A better approach is to have the classes stored in a central location available to all applications. Only one copy of the classes is needed, regardless of the number of applications you develop using them. In Java, this can be accomplished by using packages.

A *package*, which is also called a *library*, is a named group of related classes. Packages are stored in their own folder or directory. The compiler is informed of the package's location, so it can find it regardless of where the application that uses the package may be stored.

Let's look at a simple example. Suppose we create two classes: `Car` and `Truck`. Both classes are part of a package named `vehicles`. [Code Listing G-1](#) shows the listing for the `Car` class, and [Code Listing G-2](#) shows the listing for the `Truck` class.

## Code Listing G-1 (Car.java)

```
1 package vehicles;
2
3 /**
4 * This class is in the vehicles package.
5 */
6
7 public class Car
8 {
9 private int passengers; // Number of passengers
10 private double topSpeed; // Top speed in miles per hour
11
12 /**
13 * Constructor
14 */
15
16 public Car(int passengers, int topSpeed)
```

```

17 {
18 this.passengers = passengers;
19 this.topSpeed = topSpeed;
20 }
21
22 /**
23 * toString method
24 */
25
26 public String toString()
27 {
28 return "Passengers: " + passengers
29 + "\nTop Speed: " + topSpeed + " miles per hour";
30 }
31 }

```

## Code Listing G-2 (Truck.java)

```

1 package vehicles;
2
3 /**
4 * This class is in the vehicles package.
5 */
6
7 public class Truck
8 {
9 private double mpg; // Fuel economy in miles-per-gallon
10 private double tons; // Hauling capacity in tons
11
12 /**
13 * Constructor
14 */
15
16 public Truck(double mpg, double tons)
17 {
18 this.mpg = mpg;
19 this.tons = tons;
20 }
21
22 /**
23 * toString method
24 */
25
26 public String toString()
27 {
28 return "Fuel economy: " + mpg + " miles per gallon"
29 + "\nHauling capacity: " + tons + " tons";
30 }
31 }

```

Notice the first line of each file reads:

```
package vehicles;
```

The word `package` is a key word in Java, and `vehicles` is the name of the package to which the class belongs. Notice the statement is terminated by a semicolon. This statement informs the compiler that the contents of the file belong to the `vehicles` package.

## NOTE:

The compiler expects the package statement to be the first statement in a file. Only comments and blank lines may be placed before it.

When a class is part of a package, the compiler expects the class file to be stored in a directory with the same name as the package. For example, the `Car` and `Truck` classes are part of the `vehicles` package, so they must be stored in a directory named `vehicles`. Typically, the `vehicles` directory would be stored under another directory that acts as the base directory for packages. [Figure G-1](#) depicts an example directory structure on a Windows system. The `mypackages` directory is the base directory for packages. Under `mypackages`, the `vehicles` directory appears, which holds the `Car.java` and `Truck.java` files.




**Figure G-1 Example directory structure for storing a package on windows**

Once a package has been created, the Java compiler must be informed of its location. The exact process that you use depends on your Java compiler and your operating system. You should consult your instructor for precise details, but we will give a general overview of what needs to be done when using the Oracle JDK under Windows and UNIX or Linux.

The Oracle JDK uses an operating system environment variable named `CLASSPATH` to determine where the base directory for your packages is located. The `CLASSPATH` environment variable holds a string that consists of one or more pathnames, separated by a delimiting character. To inform the compiler of your base directory's location, simply add its pathname to the contents of the `CLASSPATH` environment variable.



# Setting CLASSPATH under Windows

The following is an example command that can be used at the command line in Windows to set the `CLASSPATH` environment variable:

```
set CLASSPATH=C: \mypackages;
```

There are no spaces before or after the `=` sign. In Windows, multiple pathnames stored in the variable are separated by a semicolon. This command stores two pathnames in the `CLASSPATH` environment variable: `C: \mypackages` and the `.` symbol, which represents the current directory. This tells Java to search two locations for packages: `C: \mypackages` and the current directory. Here is another example:

```
set CLASSPATH=C: \mypackages;C: \project;
```

This command stores three pathnames in the `CLASSPATH` environment variable: `C: \myPackages`, `C: \project`, and the `.` symbol, which represents the current directory. This tells Java to search these three locations when looking for packages.

# Setting CLASSPATH under UNIX or Linux

## NOTE:

The UNIX and Linux instructions in this appendix assume you are using the `bash` shell.

The following is an example command that can be used in UNIX or Linux to set the `CLASSPATH` environment variable.

```
export CLASSPATH=/home/tsmith/mypackages:
```

In UNIX and Linux, multiple pathnames stored in the variable are separated by a colon. This command stores two pathnames in the `CLASSPATH` environment variable: `/home/tsmith/mypackages` and the `.` symbol, which represents the current directory. This tells Java to search two locations for packages: `/home/tsmith/mypackages` and the current directory. Here is another example:

```
export CLASSPATH=/home/tsmith/mypackages:/home/tsmith/project;
```

This command stores three pathnames in the `CLASSPATH` environment variable: `/home/tsmith/mypackages`, `/home/tsmith/project`, and the `.` symbol, which represents the current directory. This tells Java to search these three locations when looking for packages.

 **NOTE:**

You should always include the dot symbol, which represents the current directory, as a pathname in the `CLASSPATH` variable. If you do not, your programs will not run correctly.

## Using the `import` Statement

When a package has been created and stored on the disk as previously described and the `CLASSPATH` variable has been set with the pathname of the base directory for your packages, you are ready to create programs that use your package. [Code Listing G-3](#) uses both the `Car` and `Truck` classes.

### Code Listing G-3 (`CarTruckDemo.java`)

```
1 import vehicles.*;
2
3 /**
4 * This program demonstrates the Car and Truck
5 * classes which are part of the vehicles package.
6 */
7
8 public class CarTruckDemo
9 {
10 public static void main(String[] args)
11 {
12 Car roadster = new Car(2, 155);
13 Truck pickUp = new Truck(18, 2);
14
15 System.out.println("Here's information " +
16 "about the car:");
17 System.out.println(roadster);
18 System.out.println();
19 System.out.println("Here's information " +
20 "about the truck:");
21 System.out.println(pickUp);
22 }
23 }
```

## Program Output

```
Here's information about the car:
Passengers: 2
Top speed: 155.0 miles per hour
```

```
Here's information about the truck:
Fuel economy: 18.0 miles per gallon
```

Hauling capacity: 2.0 tons

Notice line 1 contains an `import` statement. The word `import` is a key word in Java. The name that follows `import` is the name of a package the program intends to use. The `.*` that follows the package name means to import all the classes that are part of that package. This statement tells the compiler to make all the classes that are part of the `vehicles` package available to the program. If we only wanted to make the `Truck` class available, we could have used the following `import` statement:

```
import vehicles.Truck;
```

In this case, the compiler would only make the `Truck` class available. Any references to the `Car` class would cause an error. Likewise, we could use the following statement to make only the `Car` class available:

```
import vehicles.Car;
```

You will recall that you have previously used the following `import` statement in programs that use the `DecimalFormat` class:

```
import java.text.DecimalFormat;
```

This statement tells the compiler to use the `DecimalFormat` class, which is part of the `java.text` package. The `java.text` package is part of the Java standard class library. You have also used the following `import` statement in programs using the `Scanner` class:

```
import java.util.Scanner;
```

This statement tells the compiler to use the `Scanner` class, which is part of the `java.util` package. The `java.util` package is also part of the Java standard class library. At this point, it might be helpful to summarize the steps necessary to create and use a package:

1. Place an appropriate `package` statement in each class file that is to be part of the package. The `package` statement must be the first line of the file or preceded only by comments and/or blank lines.
2. You should have a base directory on your system for storing your packages. Under this directory, create a directory that bears the same name as the package. Store the package's class files in this folder.
3. Add the pathname of the package base directory to the `CLASSPATH` environment variable.
4. Place an appropriate `import` statement in each program that intends to use the package.

## Using Fully Qualified Class Names

It is not required that you use the `import` statement to access classes in a package. Without the

`import` statement, however, you must use the classes' fully qualified names. The fully qualified names of the `Car` and `Truck` classes in the `vehicle` package are `vehicles.Car` and `vehicles.Truck`. [Code Listing G-4](#) is a version of [Code Listing G-3](#), which uses fully qualified class names instead of the `import` statement.

## Code Listing G-4 (`CarTruckDemo2.java`)

```
1 /**
2 * This program demonstrates the Car and Truck
3 * classes by using their fully-qualified names
4 * instead of an import statement.
5 */
6
7 public class CarTruckDemo2
8 {
9 public static void main(String[] args)
10 {
11 vehicles.Car roadster = new vehicles.Car(2, 155);
12 vehicles.Truck pickUp = new vehicles.Truck(18, 2);
13
14 System.out.println("Here's information " +
15 "about the car:");
16 System.out.println(roadster);
17 System.out.println();
18 System.out.println("Here's information " +
19 "about the truck:");
20 System.out.println(pickUp);
21 }
22 }
```


## Program Output

Same as [Code Listing G-3](#).

Notice every occurrence of the `Car` and `Truck` class names must be written as `vehicles.Car` and `vehicles.Truck`. Obviously, it is more convenient to use the `import` statement.

## Using Package Subdirectories


You have learned that the Java compiler locates packages by searching the directories listed by the `CLASSPATH` environment variable. It is possible to create subdirectories in these locations and store packages there as well. For example, suppose we create a set of classes for calculating the pay of employees. The first step would be to create a folder to hold the packages. [Figure G-2](#) shows the `employees` directory has been added to the `mypackages` directory.



## Figure G-2 employees directory added to the mypackages directory

[Figure G-2 Full Alternative Text](#)

Under the `employees` directory are two other directories: `hourly` and `salaried`. [Figure G-3](#) shows the contents of the `hourly` directory.



# Figure G-3 Directory structure showing the contents of the `hourly` directory

## [Figure G-3 Full Alternative Text](#)

The files *Clerical.java* and *AssemblyLine.java* contain code for classes that calculate the hourly pay for clerical staff and assembly line workers. If we looked at the *Clerical.java* file, we would see something similar to the following:

```
package employees.hourly;
/**
 * This class is part of the employees.hourly package.
 */
public class Clerical
{
 Details of this class are omitted.
}
```

If we looked in the *AssemblyLine.java* file, we would see something similar to the following:

```
package employees.hourly;
/**
 * This class is part of the employees.hourly package.
 */
public class AssemblyLine
{
 Details of this class are omitted.
}
```

Notice the first line of both files is the following package statement:

```
package employees.hourly;
```

The files are stored in the `hourly` directory. Because `hourly` is a subdirectory of `employees`, the two names are separated by a period. This makes the package name `employees.hourly`. The following `import` statement could then be added to any program that needs to use the `Clerical` or `AssemblyLine` classes:

```
import employees.hourly.*;
```

This statement makes all of the classes that are part of the `employees.hourly` package accessible. As before, individual classes can be specifically named in the `import` statement. For example, the following statement will make only the `Clerical` class accessible:

```
import employees.hourly.Clerical;
```

## Access Specifiers and Packages

So far you have learned that class members may be declared as either `public` or `private`. Members that are declared as `public` may be accessed by statements outside the class as well as inside the class. Members that are declared as `private`, however, may only be accessed by statements inside the class.

In addition to `public` and `private`, Java also allows class members to be declared with no access specifier at all. When a class member is declared with no access specifier, it is accessible by any statement within the same package. For example, consider the following scenario. Class `A` and class `B` are in the same package. The variable `x` is a member of class `A` and is declared with no access specifier. Because `x` is declared with no access specifier, the methods in class `B` have access to it (as if it were `public`). Statements that are outside the package, however, cannot access `x`. [Table G-1](#) summarizes the effect of access specifiers, as they pertain to what you have learned so far.

**Table G-1 The effect of access specifiers**

| Access Attribute     | Description                                                                                                                                                                         |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| No                   | May be applied to classes and class members (variables or methods). This makes a class or class member accessible within the package.                                               |
| <code>public</code>  | May be applied to classes and class members (variables or methods). This makes a class or class member accessible to all statements in the program (inside or outside the package). |
| <code>private</code> | May only be applied to member variables and member methods. A private variable or method is only accessible by statements in the same class.                                        |

## The Standard Java Packages

The standard Java classes that make up the API are organized into packages. [Table G-2](#) lists a few of them.

**Table G-2 A few of the standard Java packages**

| Package                    | Description                                                                                                                           |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| <code>java.applet</code>   | Provides the classes necessary to create an applet.                                                                                   |
| <code>java.awt</code>      | Provides classes for the Abstract Windowing Toolkit. These classes are used in drawing images and creating graphical user interfaces. |
| <code>java.io</code>       | Provides classes that perform various types of input and output.                                                                      |
| <code>java.lang</code>     | Provides general classes for the Java language. This package is automatically imported.                                               |
| <code>java.net</code>      | Provides classes for network communications.                                                                                          |
| <code>java.security</code> | Provides classes that implement security features.                                                                                    |
| <code>java.sql</code>      | Provides classes for accessing databases using structured query language.                                                             |
| <code>java.text</code>     | Provides various classes for formatting text.                                                                                         |
| <code>java.util</code>     | Provides various utility classes.                                                                                                     |
| <code>javax.swing</code>   | Provides classes for creating graphical user interfaces.                                                                              |

To use a class from a Java package, you must have an appropriate `import` statement in your program. For example, you have used the `Scanner` class, which is in the `java.util` package. This class requires the following `import` statement:

```
import java.util.Scanner;
```

You have also used various classes from the `java.io` package to perform file input and output. To import all of the classes from the `java.io` package, you use the following `import` statement:

```
import java.io.*;
```



The `java.lang` package is the only package that does not require an `import` statement. This package contains general classes, such as `String` and `System`, that are fundamental to the Java programming language. The `java.lang` package is automatically imported by all Java programs.

# APPENDIX H Working with Records and Random Access Files\*

\*From *Starting Out with Java: From Control Structures through Objects*, Sixth Edition. Tony Gaddis. Copyright © 2016 by Pearson Education, Inc. All rights reserved.

## Structuring Data into Fields and Fixed-Length Records

Data that is written to a file is commonly structured as fields and records. In file terminology, a *field* is an individual piece of data, such as a person's name or telephone number. A *record* is a collection of fields pertaining to a single item. For example, a record might consist of a specific person's name, age, address, and telephone number.

Quite often you can save the contents of an object as a record in a file. You do this by writing each of the object's fields to the file, one after the other. When you have saved all of the object's fields, a complete record has been written. When the fields from multiple objects have been saved, then multiple records have been written to the file.

Random access files are particularly useful for storing and retrieving records. However, the sizes of the items stored in a random access file must be known in order to calculate the position of a specific item. Records that are stored in a random access file must be the same size and must have a *fixed length*. This means that the size of a record cannot change.

In Java, the sizes of the primitive data types are well documented and guaranteed to be the same on all systems. If an object's fields are all of the primitive data types, you can easily calculate the size of the record: it will be the sum of the sizes of all the fields. However, a problem arises if an object has a field that is a `String` because its contents can vary in length. You can get around this problem by making sure that a `String` field is always written as a specific number of characters. The example in this appendix shows one way to do this.

First we will introduce the `InventoryItem` class shown in [Code Listing H-1](#). An object of this class can represent an item that a company might have in its inventory. This class has two fields: `description`, a `String` that holds an item's description, and `units`, an `int` that holds the number of units on hand. The class also has the necessary accessor methods, mutator methods, and two constructors.

### Code Listing H-1 (`InventoryItem.java`)

```

1 /**
2 InventoryItem class
3 */
4
5 public class InventoryItem
6 {
7 private String description; // Item description
8 private int units; // Units on hand
9
10 /**
11 This constructor assigns an empty string
12 to description and 0 to units.
13 */
14
15 public InventoryItem()
16 {
17 description = "";
18 units = 0;
19 }
20
21 /**
22 This constructor assigns values
23 to the description and units fields.
24 @param d The description.
25 @param u The units on hand.
26 */
27
28 public InventoryItem(String d, int u)
29 {
30 description = d;
31 units = u;
32 }
33
34 /**
35 The setDescription method assigns a string
36 to the description field.
37 @param d The string to assign to description.
38 */
39
40 public void setDescription(String d)
41 {
42 description = d;
43 }
44
45 /**
46 The setUnits method assigns a value
47 to the units field.
48 @param u The value to assign to units.
49 */
50
51 public void setUnits(int u)
52 {
53 units = u;
54 }
55
56 /**
57 The getDescription method returns the item's
58 description.

```

```

59 @return The description field.
60 */
61
62 public String getDescription()
63 {
64 return description;
65 }
66
67 /**
68 * The getUnits method returns the number of
69 * units on hand.
70 * @return The units field.
71 */
72
73 public int getUnits()
74 {
75 return units;
76 }
77 }

```

The `InventoryItemFile` class shown in [Code Listing H-2](#) is designed to read and write `InventoryItem` objects as records in a random access file. The class can also move the file pointer to a specific record. To keep the code simple, none of the exceptions are caught.

## Code Listing H-2 (`InventoryItemFile.java`)

```

1 import java.io.*;
2
3 /**
4 * This class manages a random access file which contains
5 * InventoryItem records.
6 */
7
8 public class InventoryItemFile
9 {
10 private final int RECORD_SIZE = 44;
11 private RandomAccessFile inventoryFile;
12
13 /**
14 * The constructor opens a random access file
15 * for both reading and writing.
16 * @param filename The name of the file.
17 * @exception FileNotFoundException When the file
18 * is not found.
19 */
20
21 public InventoryItemFile(String filename)
22 throws FileNotFoundException
23 {
24 // Open the file for reading and writing.
25 inventoryFile =
26 new RandomAccessFile(filename, "rw");
27 }
28

```

```

29 /**
30 The writeInventoryItem method writes the contents
31 of an InventoryItem object to the file at the
32 current file pointer position.
33 @param item The InventoryItem object to write.
34 @exception IOException When a file error occurs.
35 */
36
37 public void writeInventoryItem(InventoryItem item)
38 throws IOException
39 {
40 // Get the item's description.
41 String str = item.getDescription();
42
43 // Write the description.
44 if (str.length() > 20)
45 {
46 // If there are more than 20 characters in the
47 // string, then write only the first 20.
48 for (int i = 0; i < 20; i++)
49 inventoryFile.writeChar(str.charAt(i));
50 }
51 else
52 {
53 // Write the description to the file.
54 inventoryFile.writeChars(str);
55 // Write enough spaces to pad it out
56 // to 20 characters.
57 for (int i = 0; i < (20 - str.length()); i++)
58 inventoryFile.writeChar(' ');
59 }
60
61 // Write the units to the file.
62 inventoryFile.writeInt(item.getUnits());
63 }
64
65 /**
66 The readInventoryItem method reads and returns
67 the record at the current file pointer position.
68 @return A reference to an InventoryItem object.
69 @exception IOException When a file error occurs.
70 */
71
72 public InventoryItem readInventoryItem()
73 throws IOException
74 {
75 char[] charArray = new char[20];
76
77 // Read the description, character by character,
78 // from the file into the char array.
79 for (int i = 0; i < 20; i++)
80 charArray[i] = inventoryFile.readChar();
81
82 // Store the char array in a String.
83 String desc = new String(charArray);
84
85 // Trim any trailing spaces from the string.
86 desc.trim();

```

```

87
88 // Read the units from the file.
89 int u = inventoryFile.readInt();
90
91 // Create an InventoryItem object and initialize
92 // it with these values.
93 InventoryItem item =
94 new InventoryItem(desc, u);
95
96 // Return the object.
97 return item;
98 }
99
100 /**
101 The getByteNum method returns a record's
102 starting byte number.
103 @param recordNum The record number of the
104 desired record.
105 */
106
107 private long getByteNum(long recordNum)
108 {
109 return RECORD_SIZE * recordNum;
110 }
111
112 /**
113 The moveFilePointer method moves the file
114 pointer to a specified record.
115 @param recordNum The number of the record to
116 move to.
117 @exception IOException When a file error occurs.
118 */
119
120 public void moveFilePointer(long recordNum)
121 throws IOException
122 {
123 inventoryFile.seek(getByteNum(recordNum));
124 }
125
126 /**
127 The getNumberOfRecords method returns the number
128 of records stored in the file.
129 @return The number of records in the file.
130 @exception IOException When a file error occurs.
131 */
132
133 public long getNumberOfRecords() throws IOException
134 {
135 return inventoryFile.length() / RECORD_SIZE;
136 }
137
138 /**
139 The close method closes the file.
140 @exception IOException When a file error occurs.
141 */
142
143 public void close() throws IOException

```

```

144 {
145 inventoryFile.close();
146 }
147 }

```

The `RECORD_SIZE` field, declared in line 10, is a `final int` variable initialized with the value 44. This is the size, in bytes, of a record. In a moment you will see how this number was determined. The `inventoryFile` field, declared in line 11, is a `RandomAccessFile` reference variable that will be used to open and work with a random access file. The constructor accepts a filename as a `String`. This filename is used to open a random access file, referenced by the `inventoryFile` variable, for reading and writing.

By looking at the `writeInventoryItem` method, in lines 37 through 63, we can see how the record size of 44 bytes was determined. The method accepts an `InventoryItem` object as an argument, the contents of which will be written as a record to the file. In line 41 the `description` field is retrieved and referenced by `str`, a local variable. Next, in lines 44 through 59, we write the `description` field to the file. To ensure that each record has the same fixed length, this method always writes the description as 20 characters. If the description has more than 20 characters, then only the first 20 are written. If the description has fewer than 20 characters, spaces are added to make up the difference. Next, in line 62, the method writes the `units` field, as an `int`, to the file.

Now we can see how the record size of 44 bytes was determined. When a character is written to the file, it is written as two bytes. The `description` field is written as 20 characters, so that's 40 bytes. The `units` field is written as an `int`, which uses 4 bytes. That makes a total record size of 44 bytes.

The `readInventoryItem` method in lines 72 through 98 reads a record from the file and returns an `InventoryItem` object containing the record's data. In line 75 the reference variable `charArray` is declared and a 20-element `char` array is created to hold the description. Then the code in lines 79 and 80 reads the 20 characters from the file and stores them in the array. Next, in line 83, a `String` object is created and the `char` array is passed as an argument. This copies the characters from the array to the `String` object.

If the description was less than 20 characters long, it will be padded with trailing spaces. The statement in line 86 trims any trailing spaces that might be in the string. Then the statement in line 89 reads the `units` field from the file and stores it in the `u` variable.

Now we can construct an `InventoryItem` object with the data we have read. This is done in lines 93 and 94. The last step, in line 97, is to return the object.

The class also has the ability to move the file pointer to a specific record. Two methods work together to perform this. First, `getBytesNum` (in lines 107 through 110) is a private method that accepts a record number as an argument, and returns the record's starting byte number. It calculates the starting byte number by multiplying the record size by the record number. (The first record in the file is considered record 0.) The `moveFilePointer` method (in lines 120 through 124) accepts a record number as its argument, and moves the file pointer to the specified record. This method calls the `getBytesNum` method to determine the record's starting location.

The `getNumberOfRecords` method appears in lines 133 through 136. This method returns the number of records in the file. It calculates the number of records by dividing the length of the file by the record size. The length of the file is returned by the `RandomAccessFile` class's `length` method.

The last method in the class is the `close` method, which closes the file. The program in [Code Listing H-3](#) shows a simple demonstration of this class. This program asks the user to enter data for five items, which are stored in an array of `InventoryItem` objects. The program then saves the contents of the array elements to a file.

## Code Listing H-3 (CreateInventoryFile.java)

```
1 import java.io.*;
2 import java.util.Scanner;
3
4 /**
5 * This program uses the InventoryFile class to create a
6 * file containing data from 5 InventoryItem objects.
7 */
8
9 public class CreateInventoryFile
10 {
11 public static void main(String[] args) throws IOException
12 {
13 final int NUM_ITEMS = 5; // Number of items
14 String description; // Item description
15 int units; // Units on hand
16
17 // Create a Scanner object for keyboard input.
18 Scanner keyboard = new Scanner(System.in);
19
20 // Create an array to hold InventoryItem objects.
21 InventoryItem[] items = new InventoryItem[NUM_ITEMS];
22
23 // Get data for the InventoryItem objects.
24 System.out.println("Enter data for " + NUM_ITEMS +
25 " inventory items.");
26
27 for (int i = 0; i < items.length; i++)
28 {
29 // Get the description.
30 System.out.print("Enter an item description: ");
31 description = keyboard.nextLine();
32
33 // Get the units on hand.
34 System.out.print("Enter the number of units: ");
35 units = keyboard.nextInt();
36
37 // Consume the remaining newline.
38 keyboard.nextLine();
39
40 // Create an InventoryItem object in the array.
41 items[i] = new InventoryItem(description, units);
42 }
43 }
44 }
```



```

43
44 // Create an InventoryFile object.
45 InventoryItemFile file =
46 new InventoryItemFile("Inventory.dat");
47
48 // Write the contents of the array to the file.
49 for (int i = 0; i < items.length; i++)
50 {
51 file.writeInventoryItem(items[i]);
52 }
53
54 // Close the file.
55 file.close();
56
57 System.out.println("The data was written to the " +
58 "Inventory.dat file.");
59 }
60 }

```

## Program Output with Example Input Shown in Bold

```

Enter data for 5 inventory items.
Enter an item description: Wrench [Enter]
Enter the number of units: 20 [Enter]
Enter an item description: Hammer [Enter]
Enter the number of units: 15 [Enter]
Enter an item description: Pliers [Enter]
Enter the number of units: 12 [Enter]
Enter an item description: Screwdriver [Enter]
Enter the number of units: 25 [Enter]
Enter an item description: Ratchet [Enter]
Enter the number of units: 10 [Enter]
The data was written to the Inventory.dat file.

```

The program in [Code Listing H-4](#) demonstrates how records can be randomly read from the file.

## Code Listing H-4 (ReadInventoryFile.java)

```

1 import java.io.*;
2 import java.util.Scanner;
3
4 /**
5 * This program displays specified records from
6 * the Inventory.dat file.
7 */
8
9 public class ReadInventoryFile
10 {
11 public static void main(String[] args) throws IOException

```

```

12 {
13 int recordNumber; // Record number
14 String again; // To get a Y or an N
15 InventoryItem item; // An object from the file
16
17 // Create a Scanner object for keyboard input.
18 Scanner keyboard = new Scanner(System.in);
19
20 // Open the file.
21 InventoryItemFile file =
22 new InventoryItemFile("Inventory.dat");
23
24 // Report the number of records in the file.
25 System.out.println("The Inventory.dat file has " +
26 file.getNumberOfRecords() + " records.");
27
28 // Get a record number from the user and
29 // display the record.
30 do
31 {
32 // Get the record number.
33 System.out.print("Enter the number of the record " +
34 "you wish to see: ");
35 recordNumber = keyboard.nextInt();
36
37 // Consume the remaining newline.
38 keyboard.nextLine();
39
40 // Move the file pointer to that record.
41 file.moveFilePointer(recordNumber);
42
43 // Read the record at that location.
44 item = file.readInventoryItem();
45
46 // Display the record.
47 System.out.println("\nDescription: " +
48 item.getDescription());
49 System.out.println("Units: " + item.getUnits());
50
51 // Ask the user whether to get another record.
52 System.out.print("\nDo you want to see another " +
53 "record? (Y/N): ");
54 again = keyboard.nextLine();
55 } while (again.charAt(0) == 'Y' || again.charAt(0) == 'y');
56
57 // Close the file.
58 file.close();
59 }
60 }

```

**Program Output with Example Input Shown in Bold**

```

The Inventory.dat file has 5 records.
Enter the number of the record you wish to see: 4 [Enter]
Description: Ratchet
Units: 10
Do you want to see another record? (Y/N): y [Enter]
Enter the number of the record you wish to see: 2 [Enter]
Description: Pliers
Units: 12
Do you want to see another record? (Y/N): y [Enter]
Enter the number of the record you wish to see: 0 [Enter]
Description: Wrench
Units: 20
Do you want to see another record? (Y/N): y [Enter]
Enter the number of the record you wish to see: 1 [Enter]
Description: Hammer
Units: 15
Do you want to see another record? (Y/N): y [Enter]
Enter the number of the record you wish to see: 3 [Enter]
Description: Screwdriver
Units: 25
Do you want to see another record? (Y/N): n [Enter]

```

As a last demonstration, the program in [Code Listing H-5](#) shows how an existing record in the file can be overwritten with a new record.

## Code Listing H-5 (ModifyRecord.java)

```

1 import java.io.*;
2 import java.util.Scanner;
3
4 /*
5 This program allows the user to modify records in the
6 Inventory.dat file.
7 */
8
9 public class ModifyRecord
10 {
11 public static void main(String[] args) throws IOException
12 {
13 int recordNumber; // Record number
14 int units; // Units on hand
15 String again; // Want to change another one?
16 String sure; // Is the user sure?
17 String description; // Item description
18 InventoryItem item; // To reference an item
19
20 // Create a Scanner object for keyboard input.
21 Scanner keyboard = new Scanner(System.in);
22
23 // Open the file.
24 InventoryItemFile file =
25 new InventoryItemFile("Inventory.dat");
26
27 // Report the number of records in the file.

```

```

28 System.out.println("The Inventory.dat file has " +
29 file.getNumberOfRecords() + " records.");
30
31 // Get a record number from the user and
32 // allow the user to modify it.
33 do
34 {
35 // Get the record number.
36 System.out.print("Enter the number of the record " +
37 "you wish to modify: ");
38 recordNumber = keyboard.nextInt();
39
40 // Consume the remaining newline.
41 keyboard.nextLine();
42
43 // Move the file pointer to that record number.
44 file.moveFilePointer(recordNumber);
45
46 // Read the record at that location.
47 item = file.readInventoryItem();
48
49 // Display the existing contents.
50 System.out.println("Existing data:");
51 System.out.println("\nDescription: " +
52 item.getDescription());
53 System.out.println("Units: " + item.getUnits());
54
55 // Get the new data.
56 System.out.print("\nEnter the new description: ");
57 description = keyboard.nextLine();
58 System.out.print("Enter the number of units: ");
59 units = keyboard.nextInt();
60 keyboard.nextLine(); // Consume the remaining newline.
61
62 // Store the new data in the object.
63 item.setDescription(description);
64 item.setUnits(units);
65
66 // Make sure the user wants to save this data.
67 System.out.print("Are you sure you want to save " +
68 "this data? (Y/N) ");
69 sure = keyboard.nextLine();
70 if (sure.charAt(0) == 'Y' || sure.charAt(0) == 'y')
71 {
72 // Move back to the record's starting position.
73 file.moveFilePointer(recordNumber);
74 // Save the new data.
75 file.writeInventoryItem(item);
76 }
77
78 // Ask the user whether to change another record.
79 System.out.print("\nDo you want to modify another " +
80 "record? (Y/N): ");
81 again = keyboard.nextLine();
82 } while (again.charAt(0) == 'Y' || again.charAt(0) == 'y');
83
84 // Close the file.
85 file.close();

```

```
86 }
87 }
```

# Program Output with Example Input Shown in Bold

```
The Inventory.dat file has 5 records.
Enter the number of the record you wish to modify: 3 [Enter]
Existing data:
Description: Screwdriver
Units: 25
Enter the new description: Duct Tape [Enter]
Enter the number of units: 30 [Enter]
Are you sure you want to save this data? (Y/N) y [Enter]
Do you want to modify another record? (Y/N): n [Enter]
```

In the example running of the program, record 3 was modified. We can run the `ReadInventoryFile` program in [Code Listing H-4](#) again to verify that the record was changed. Here is the output of that program if we run it again.

# Program Output with Example Input Shown in Bold (ReadInventoryFile.java)

```
The Inventory.dat file has 5 records.
Enter the number of the record you wish to see: 3 [Enter]
Description: Duct Tape
Units: 30
Do you want to see another record? (Y/N): n [Enter]
```

# **APPENDIX I Configuring Java DB and Installing Apache Derby**

Beginning with Java 7, the Java DB database engine was bundled with the Java JDK. When you installed either Java 7 or Java 8, Java DB was automatically installed. Beginning with Java 9, however, Java DB is no longer bundled with the JDK. This is because Java DB is simply a rebranded version of Derby, an open source DBMS that is provided by the Apache Software Foundation. If you are using Java 9, you will need to separately download Derby from the Apache Software Foundation's website and install it. This appendix can help you configure Java DB, or install and configure Apache Derby, depending on the version of Java you are using.

# Configuring Java DB with Java 7 or Java 8

In Java 7 and Java 8, Java DB is automatically installed on your system when you install the JDK. When you install the JDK, take note of the location on your system where it is installed. If you accept the default location for the JDK, it will be installed in a folder with a path such as this:

```
C:\Program Files\Java\jdk1.8.0_101
```

The name of the last folder in the path might be different on your system, because it's name indicates the version of the JDK that you are using. Nevertheless, the path will be similar to the one shown here.

If you look inside the folder where the JDK is installed, you will see another folder named `db`. Inside the `db` folder, you will see another folder named `lib`, and inside the `lib` folder, you will see a file named *derby.jar*. The *derby.jar* file contains the Java DB database engine and its associated JDBC embedded driver. Take a moment to verify that you can locate the *derby.jar* file, and make a note of its path. For example, if you are using the JDK 8, update 101, the path of the *derby.jar* file will be something like:

```
C:\Program Files\Java\jdk1.8.0_101\db\lib\derby.jar
```

Once you have verified the path of the *derby.jar* file, you need to set the `CLASSPATH` environment variable on your system.

## Setting the CLASSPATH Environment Variable

To use Java DB with this text, you will need to set your system's `CLASSPATH` environment variable to include the location of the *derby.jar* file. The `CLASSPATH` environment variable contains a list of class libraries that are needed by the Java Virtual Machine and other Java applications.

The following instructions describe the procedure for setting the `CLASSPATH` variable on Windows 10, Windows 8, and Windows 7.

## Windows 10

In the Windows search box, type “environment variables” then select **Edit the system environment variables** from the search results. This will display the *System Properties* window (with the *Advanced* tab selected at the top of the window). Click the Environment Variables . . . button. Look in the **System Variables** list to see if the `CLASSPATH` variable exists.

If the `CLASSPATH` variable already exists, perform the following:

- Select the `CLASSPATH` variable and click the **Edit** button. Add a semicolon to the end of the existing contents and then add the path of the `derby.jar` file to the end of the list.
- Click the OK buttons until all the dialog boxes are closed and exit the control panel.

If the `CLASSPATH` variable does not exist, perform the following:

- Click the New . . . button that appears below the list of system variables.
- In the Variable name box, type `CLASSPATH`.
- In the Variable value box, type a period, followed by a semicolon, followed by the path of the *derby.jar* file. For example, the value will be something similar to the following (depending on the version of the JDK that you are using):

```
.;C:\Program Files\Java\jdk1.8.0_101\db\lib\derby.jar
```

- Click the OK buttons until all the dialog boxes are closed and exit the control panel.

## Windows 8

In the right bottom corner of the screen, click on the **Search** icon and type “Control Panel”. Click on **Control Panel**, then click **System**, then click **Advanced system settings**. Click on the **Advanced** tab, then click **Environment Variables**. Look in the System Variables list to see if the `CLASSPATH` variable exists.

If the `CLASSPATH` variable already exists, perform the following:

- Select the `CLASSPATH` variable and click the **Edit** button. Add a semicolon to the end of the existing contents and then add the path of the *derby.jar* file to the end of the list.
- Click the OK buttons until all the dialog boxes are closed and exit the control panel.

If the `CLASSPATH` variable does not exist, perform the following:

- Click the New . . . button that appears below the list of system variables.
- In the Variable name box, type `CLASSPATH`.
- In the Variable value box, type a period, followed by a semicolon, followed by the path of the *derby.jar* file. For example, the value will be something similar to the following (depending on the version of the JDK that you are using):

```
.;C:\Program Files\Java\jdk1.8.0_101\db\lib\derby.jar
```

- Click the OK buttons until all the dialog boxes are closed and exit the control panel.



# Windows 7

Click the **Start button** and then right-click **Computer**. On the pop-up menu, select **Properties**. In the window that appears next, click **Advanced system settings**. This displays the System Properties window. Click the Environment Variables . . . button. Look in the System Variables list to see if the `CLASSPATH` variable exists.

If the `CLASSPATH` variable already exists, perform the following:

- Select the `CLASSPATH` variable and click the Edit button. Add a semicolon to the end of the existing contents and then add the path of the *derby.jar* file to the end of the list.
- Click the OK buttons until all the dialog boxes are closed and exit the control panel.

If the `CLASSPATH` variable does not exist, perform the following:

- Click the New . . . button that appears below the list of system variables.
- In the Variable name box, type `CLASSPATH`.
- In the Variable value box, type a period, followed by a semicolon, followed by the path of the *derby.jar* file. For example, the value will be something similar to the following (depending on the version of the JDK that you are using):

```
.;C:\Program Files\Java\jdk1.8.0_101\db\lib\derby.jar
```

- Click the OK buttons until all the dialog boxes are closed and exit the control panel.

# Configuring Derby with Java 9

Derby is available as a free download from the Apache Software Foundation. To download Derby, follow these instructions:

1. Open your browser and go the following address: [http://db.apache.org/derby/derby\\_downloads.html](http://db.apache.org/derby/derby_downloads.html)
2. The page that opens will look similar to [Figure I-1](#). Look at the section titled *For Java 8 and Higher*, and click on the most recent version number that is listed. (The most recent version number will most likely be different than the one shown in the figure.)




**Figure I-1 The Apache Derby Download Page**

## Figure I-1 Full Alternative Text

3. The page that opens will look similar to [Figure I-2](#). In the *Distributions* section you will see links to download various different distributions. As shown in the figure, click on the link with the most recent version number ending with `bin.zip`. (The most recent version number will most likely be different than the one shown in the figure.)




**Figure I-2 Select the latest bin**

# distribution

(Apache Software Foundation)

## [Figure I-2 Full Alternative Text](#)

4. The ZIP file that you downloaded contains a folder with a name such as `db-derby-10.14.1.0-bin`. (This name reflects the version number of the distribution. The name you see will be similar, but most likely it will be a different version number.) Extract this folder and its contents to a location on your system. Make a note of the location to which you extracted the folder. For example, on a Windows system, if you extracted the folder to the root folder of drive C: the path would be `C:\db-derby-10.14.1.0-bin`.

## Setting the CLASSPATH Environment Variable

The Derby distribution file that you downloaded and extracted contains a file named `derby.jar`. This file is the Derby database engine. The `derby.jar` file will be located in the `lib` folder, which is in the distribution folder that you extracted. For example, suppose you extracted the `db-derby-10.14.1.0-bin` distribution file to the root folder of drive C:. The path of the `derby.jar` file would be:

```
C:\db-derby-10.14.1.0-bin\lib\derby.jar
```

You will need to configure your system so the Java Virtual Machine can find the `derby.jar` file. You do this by setting your system's `CLASSPATH` environment variable to include the path of the file. The following instructions describe the procedure for setting the `CLASSPATH` variable on Windows 10, Windows 8, and Windows 7.

## Windows 10

In the Windows search box, type *environment variables* and then select *Edit the system environment variables* from the search results. This will display the System Properties window (with the Advanced tab selected at the top of the window). Click the *Environment Variables . . .* button. Look in the System Variables list to see if the `CLASSPATH` variable exists.

If the `CLASSPATH` variable already exists, perform the following:

- Select the `CLASSPATH` variable and click the *Edit* button. Add a semicolon to the end of the existing contents and then add the path of the `derby.jar` file to the end of the list.
- Click the *OK* buttons until all the dialog boxes are closed and exit the control panel.

If the `CLASSPATH` variable does not exist, perform the following:

- Click the *New . . .* button that appears below the list of system variables.
- In the Variable name box, type `CLASSPATH`.
- In the Variable value box, type a period, followed by a semicolon, followed by the path of the `derby.jar` file. For example, the value will be something similar to the following (depending on the version of the Derby distribution that you are using):

```
.; C:\db-derby-10.14.1.0-bin\lib\derby.jar
```

- Click the *OK* buttons until all the dialog boxes are closed and exit the control panel.

## Windows 8

In the right bottom corner of the screen, click on the Search icon and type *Control Panel*. Click on *Control Panel*, then click *System*, then click *Advanced system settings*. Click on the *Advanced* tab, then click *Environment Variables*. Look in the System Variables list to see if the `CLASSPATH` variable exists.

If the `CLASSPATH` variable already exists, perform the following:

- Select the `CLASSPATH` variable and click the *Edit* button. Add a semicolon to the end of the existing contents and then add the path of the `derby.jar` file to the end of the list.
- Click the *OK* buttons until all the dialog boxes are closed and exit the control panel.

If the `CLASSPATH` variable does not exist, perform the following:

- Click the *New . . .* button that appears below the list of system variables.
- In the Variable name box, type `CLASSPATH`.
- In the Variable value box, type a period, followed by a semicolon, followed by the path of the `derby.jar` file. For example, the value will be something similar to the following (depending on the version of the Derby distribution that you are using):

```
.; C:\db-derby-10.14.1.0-bin\lib\derby.jar
```

- Click the *OK* buttons until all the dialog boxes are closed and exit the control panel.

## Windows 7

Click the Windows Start button and then right-click *Computer*. On the pop-up menu, select *Properties*. In the window that appears next, click *Advanced system settings*. This displays the System Properties window. Click the *Environment Variables . . .* button. Look in the System Variables list to see if the `CLASSPATH` variable exists.

If the `CLASSPATH` variable already exists, perform the following:

- Select the `CLASSPATH` variable and click the *Edit* button. Add a semicolon to the end of the existing contents and then add the path of the `derby.jar` file to the end of the list.
- Click the *OK* buttons until all the dialog boxes are closed and exit the control panel.

If the `CLASSPATH` variable does not exist, perform the following:

- Click the *New . . .* button that appears below the list of system variables.
- In the Variable name box, type `CLASSPATH`.
- In the Variable value box, type a period, followed by a semicolon, followed by the path of the `derby.jar` file. For example, the value will be something similar to the following (depending on the version of the Derby distribution that you are using):

```
.; C:\db-derby-10.14.1.0-bin\lib\derby.jar
```

- Click the *OK* buttons until all the dialog boxes are closed and exit the control panel.

# APPENDIX J The QuickSort Algorithm

The QuickSort algorithm is a popular general-purpose sorting routine developed in 1960 by C. A. R. Hoare. It sorts an array by dividing it into two sub lists. Between the sub lists is a selected value known as the pivot. This is illustrated in [Figure J-1](#).




**Figure J-1 Sub lists and pivot**

Notice in the figure that sub list 1 is positioned to the left of (before) the pivot, and sub list 2 is positioned to the right of (after) the pivot. Once a pivot value has been selected, the algorithm swaps the other values in the array until all the elements in sub list 1 are less than the pivot, and all the elements in sub list 2 are greater than the pivot.

Once this is done, the algorithm repeats the procedure on sub list 1, then on sub list 2. The recursion stops when there is only one element in a sub list. At that point, the array is completely sorted.

The algorithm is coded primarily in two methods: `quickSort` and `partition`. The `quickSort` method is recursive. Its psuedocode is shown here:

- *quickSort:*
  - *If Starting Index < Ending Index*
    - *Partition the List around a Pivot*
    - *quickSort Sub list 1*
    - *quickSort Sub list 2*
  - *End If*

Here is the Java code for the `quickSort` method:

```
public static void quickSort(int array[], int start, int end)
{
 int pivotPoint;
```

```

if (start < end)
{
 // Get the pivot point.
 pivotPoint = partition(array, start, end);

 // Sort the first sub list.
 quickSort(array, start, pivotPoint - 1);

 // Sort the second sub list.
 quickSort(array, pivotPoint + 1, end);
}
}

```

This version of `quickSort` works with an array of integers. Its first argument is the array holding the list that is to be sorted. The second and third arguments are the starting and ending subscripts of the list.

The subscript of the pivot element is returned by the `partition` method. The `partition` method not only determines which element will be the pivot, but also controls the rearranging of the other values in the list. The method selects the element in the middle of the array as the pivot, and then scans the remainder of the array searching for values less than the pivot.

The code for the `partition` method is shown here:

```

private static int partition(int array[], int start, int end)
{
 int pivotValue, pivotIndex, mid;

 mid = (start + end) / 2;

 swap(array, start, mid);
 pivotIndex = start;
 pivotValue = array[start];
 for (int scan = start + 1; scan <= end; scan++)
 {
 if (array[scan] < pivotValue)
 {
 pivotIndex++;
 swap(array, pivotIndex, scan);
 }
 }
 swap(array, start, pivotIndex);

 return pivotIndex;
}

```

The `partition` method does not initially sort the values into their final order. Its job is only to move the values that are less than the pivot to the pivot's left, and move the values that are greater than the pivot to the pivot's right. As long as that condition is met, they may appear in any order. The ultimate sorting order of the entire array is achieved cumulatively, though the recursive calls to the `quickSort` method.



There are many different ways of partitioning the array. As previously stated, the technique shown in this `partition` method selects the middle value as the pivot. That value is then moved to the beginning of the array (by swapping it with the value stored there). This simplifies the next step, which is to scan the array.

A `for` loop scans the remainder of the array, and when an element is found whose value is less than the pivot, that value is moved to a location left of the pivot point.

A third method, `swap`, is used to swap the values found in any two elements of the array. The method is shown here:

```
private static void swap(int[] array, int a, int b)
{
 int temp;

 temp = array[a];
 array[a] = array[b];
 array[b] = temp;
}
```

The program in Code Listing J-1 demonstrates these methods.

## Code Listing J-1      (`QuickSortDemo.java`)

```
1 /**
2 * This program demonstrates the QuickSort algorithm.
3 */
4
5 public class QuickSortDemo
6 {
7 /**
8 * main method
9 */
10
11 public static void main(String[] args)
12 {
13 int[] array = { 7, 3, 9, 2, 0, 1, 8, 4, 6, 5 };
14
15 // Display the array as it is now.
16 System.out.println("Before the sort:");
17 for (int i = 0; i < 10; i++)
18 System.out.print(array[i] + " ");
19 System.out.println();
20
21 // Sort the array.
22 quickSort(array, 0, 9);
23
24 // Display the array again.
25 System.out.println("After the sort:");
26 for (int i = 0; i < 10; i++)
27 System.out.print(array[i] + " ");
28 System.out.println();
```

```

29 }
30
31 /**
32 * The quickSort method uses the QuickSort algorithm to
33 * sort array, from array[start] through array[end].
34 */
35
36 public static void quickSort(int array[], int start, int end)
37 {
38 int pivotPoint;
39
40 if (start < end)
41 {
42 // Get the pivot point.
43 pivotPoint = partition(array, start, end);
44
45 // Sort the first sub list.
46 quickSort(array, start, pivotPoint - 1);
47
48 // Sort the second sub list.
49 quickSort(array, pivotPoint + 1, end);
50 }
51 }
52
53 /**
54 * The partition method selects the value in the middle of
55 * the array as the pivot. The list is rearranged so all
56 * the values less than the pivot are on its left and all
57 * the values greater than pivot are on its right.
58 */
59
60 private static int partition(int array[], int start, int end)
61 {
62 int pivotValue, pivotIndex, mid;
63
64 mid = (start + end) / 2;
65
66 swap(array, start, mid);
67 pivotIndex = start;
68 pivotValue = array[start];
69 for (int scan = start + 1; scan <= end; scan++)
70 {
71 if (array[scan] < pivotValue)
72 {
73 pivotIndex++;
74 swap(array, pivotIndex, scan);
75 }
76 }
77 swap(array, start, pivotIndex);
78
79 return pivotIndex;
80 }
81
82 /**
83 * The swap method swaps the element at array[a] with the
84 * element at array[b].
85 */
86

```

```
87 private static void swap(int[] array, int a, int b)
88 {
89 int temp;
90
91 temp = array[a];
92 array[a] = array[b];
93 array[b] = temp;
94 }
95 }
```

# Program Output

```
Before the sort:
7 3 9 2 0 1 8 4 6 5
After the sort:
0 1 2 3 4 5 6 7 8 9
```

# APPENDIX K Named Colors

- aliceblue = #f0f8ff
- antiquewhite = #faebd7
- aqua = #00ffff
- aquamarine = #7fffd4
- azure = #f0ffff
- beige = #f5f5dc
- bisque = #ffe4c4
- black = #000000
- blanchedalmond = #ffeacd
- blue = #0000ff
- blueviolet = #8a2be2
- brown = #a52a2a
- burlywood = #deb887
- cadetblue = #5f9ea0
- chartreuse = #7fff00
- chocolate = #d2691e
- coral = #ff7f50
- cornflowerblue = #6495ed
- cornsilk = #fff8dc
- crimson = #dc143c
- cyan = #00ffff
- darkblue = #00008b

- darkcyan = #008b8b
- darkgoldenrod = #b8860b
- darkgray = #a9a9a9
- darkgreen = #006400
- darkgrey = #a9a9a9
- darkkhaki = #bdb76b
- darkmagenta = #8b008b
- darkolivegreen = #556b2f
- darkorange = #ff8c00
- darkorchid = #9932cc
- darkred = #8b0000
- darksalmon = #e9967a
- darkseagreen = #8fbc8f
- darkslateblue = #483d8b
- darkslategray = #2f4f4f
- darkslategrey = #2f4f4f
- darkturquoise = #00ced1
- darkviolet = #9400d3
- deeppink = #ff1493
- deepskyblue = #00bfff
- dimgray = #696969
- dimgrey = #696969
- dodgerblue = #1e90ff
- firebrick = #b22222

- floralwhite = #fffaf0
- forestgreen = #228b22
- fuchsia = #ff00ff
- gainsboro = #dcdcdc
- ghostwhite = #f8f8ff
- gold = #ffd700
- goldenrod = #daa520
- gray = #808080
- green = #008000
- greenyellow = #adff2f
- grey = #808080
- honeydew = #f0ffff
- hotpink = #ff69b4
- indianred = #cd5c5c
- indigo = #4b0082
- ivory = #fffff0
- khaki = #f0e68c
- lavender = #e6e6fa
- lavenderblush = #fff0f5
- lawngreen = #7cfc00
- lemonchiffon = #ffffacd
- lightblue = #add8e6
- lightcoral = #f08080
- lightcyan = #e0ffff

- lightgoldenrodyellow = #fafad2
- lightgray = #d3d3d3
- lightgreen = #90ee90
- lightgrey = #d3d3d3
- lightpink = #ffb6c1
- lightsalmon = #ffa07a
- lightseagreen = #20b2aa
- lightskyblue = #87cefa
- lightslategray = #778899
- lightslategrey = #778899
- lightsteelblue = #b0c4de
- lightyellow = #ffffe0
- lime = #00ff00
- limegreen = #32cd32
- linen = #faf0e6
- magenta = #ff00ff
- maroon = #800000
- mediumaquamarine = #66cdaa
- mediumblue = #0000cd
- mediumorchid = #ba55d3
- mediumpurple = #9370db
- mediumseagreen = #3cb371
- mediumslateblue = #7b68ee
- mediumspringgreen = #00fa9a

- mediumturquoise = #48d1cc
- mediumvioletred = #c71585
- midnightblue = #191970
- mintcream = #f5fffa
- mistyrose = #ffe4e1
- moccasin = #ffe4b5
- navajowhite = #ffdead
- navy = #000080
- oldlace = #fdf5e6
- olive = #808000
- olivedrab = #6b8e23
- orange = #ffa500
- orangered = #ff4500
- orchid = #da70d6
- palegoldenrod = #eee8aa
- palegreen = #98fb98
- paleturquoise = #afeeee
- palevioletred = #db7093
- papayawhip = #ffefd5
- peachpuff = #ffdab9
- peru = #cd853f
- pink = #ffc0cb
- plum = #dda0dd
- powderblue = #b0e0e6



- purple = #800080
- red = #ff0000
- rosybrown = #bc8f8f
- royalblue = #4169e1
- saddlebrown = #8b4513
- salmon = #fa8072
- sandybrown = #f4a460
- seagreen = #2e8b57
- seashell = #fff5ee
- sienna = #a0522d
- silver = #c0c0c0
- skyblue = #87ceeb
- slateblue = #6a5acd
- slategray = #708090
- slategrey = #708090
- snow = #ffaafa
- springgreen = #00ff7f
- steelblue = #4682b4
- tan = #d2b48c
- teal = #008080
- thistle = #d8bfd8
- tomato = #ff6347
- transparent = rgba(0,0,0,0)
- turquoise = #40e0d0

- violet = #ee82ee
- wheat = #f5deb3
- white = #ffffff
- whitesmoke = #f5f5f5
- yellow = #ffff00
- yellowgreen = #9acd32

# APPENDIX L Answers to Checkpoints\*

\*From *Starting Out with Java: From Control Structures through Data Structures*, Fourth Edition. Tony Gaddis. Copyright © 2019 by Pearson Education, Inc. All rights reserved.

# Chapter 1

1. [1.1](#) Because the computer can be programmed to do so many different tasks.
2. [1.2](#) The Central Processing Unit (CPU), main memory, secondary storage devices, input devices, output devices.
3. [1.3](#) Arithmetic and Logic Unit (ALU), and Control Unit
4. [1.4](#) Fetch: The CPU's control unit fetches the program's next instruction from main memory.

Decode: The control unit decodes the instruction, which is encoded in the form of a number. An electrical signal is generated.

Execute: The signal is routed to the appropriate component of the computer, which causes a device to perform an operation.

5. [1.5](#) A unique number assigned to each section of memory. The address is used to identify a location in memory.
6. [1.6](#) Program instructions and data are stored in main memory while the program is operating. Main memory is volatile, and loses its contents when power is removed from the computer. Secondary storage holds data for long periods of time—even when there is no power to the computer.
7. [1.7](#) It means that an operating system is capable of running multiple programs at once.
8. [1.8](#) A key word has a special purpose, and is defined as part of a programming language. A programmer-defined symbol is a word or name defined by the programmer.
9. [1.9](#) Operators perform operations on one or more operands. Punctuation symbols mark the beginning or ending of a statement, or separates items in a list.
10. [1.10](#) A line is a single line as it appears in the body of a program. A statement is a complete instruction that causes the computer to perform an action.
11. [1.11](#) Because their contents may be changed.
12. [1.12](#) The original value is overwritten.
13. [1.13](#) A compiler is a program that translates source code into an executable form.
14. [1.14](#) Syntax errors are mistakes that the programmer has made that violate the rules of the programming language. These errors must be corrected before the compiler can translate the source code.

15. [1.15](#) The Java compiler translates Java source code into byte code, which is an intermediate language. The Java Virtual Machine executes the byte code instructions.
16. [1.16](#) The Java Virtual Machine (JVM) is a program that reads Java byte code instructions and executes them as they are read. In other words, it interprets byte code instructions. You can think of the JVM as a program that simulates a computer whose machine language is Java byte code.
17. [1.17](#) The program's purpose, input, process, and output.
18. [1.18](#) Before you create a program on the computer, it is often a good idea to imagine what the computer screen will look like while the program is running. If it helps, draw pictures of the screen, with sample input and output, at various points in the program.
19. [1.19](#) A cross between human language and a programming language. Pseudocode is especially helpful when designing an algorithm. Although the computer can't understand pseudocode, programmers often find it helpful to write an algorithm in a language that's "almost" a programming language, but still very similar to natural language.
20. [1.20](#) A compiler translates source code into an executable form.
21. [1.21](#) A runtime error is an error that occurs while the program is running. These are usually logical errors, such as mathematical mistakes.
22. [1.22](#) Syntax errors are found by the compiler.
23. [1.23](#) You can provide sample data and predict what the output should be. If the program does not produce the correct output, a logical error is present in the program.
24. [1.24](#) Data and the code that operates on the data.
25. [1.25](#) The data contained in an object.
26. [1.26](#) The procedures, or behaviors, that an object performs.
27. [1.27](#) Encapsulation refers to the combining of data and code into a single object.
28. [1.28](#) Data hiding refers to an object's ability to hide its data from code that is outside the object. Only the object's methods may then directly access and make changes to the object's data. An object typically hides its data, but allows outside code to access the methods that operate on the data.

# Chapter 2

## 1. [2.1](#)

```
// A crazy mixed up program
public class Columbus
{
 public static void main(String[] args)
 {
 System.out.println("In 1492 Columbus sailed the ocean blue.");
 }
}
```

## 2. [2.2](#) Columbus.java

## 3. [2.3](#)

```
public class Hello
{
 public static void main(String[] args)
 {
 System.out.println("Hello World");
 }
}
```

## 4. [2.4](#)

```
// Example
// August 22, 2018
public class MyName
{
 public static void main(String[] args)
 {
 System.out.println("Herbert Dorfmann");
 }
}
```

## 5. [2.5](#)C

## 6. [2.6](#)A

## 7. [2.7](#)

```
// Its a mad, mad program
public class Success
{
 public static void main(String[] args)
 {
 System.out.print("Success\n");
 System.out.print("Success ");
 System.out.print("Success\n");
 }
}
```

```

 System.out.println("\nSuccess");
 }
}

```

## 8. [2.8](#)

The works of Wolfgang  
include the following  
The Turkish March and Symphony No. 40 in G minor.

## 9. [2.9](#)

```

// August 22, 2018
public class PersonalInfo
{
 public static void main(String[] args)
 {
 System.out.println("Herbert Dorfmann");
 System.out.println("123 Elm Street");
 System.out.println("My Town, NC 21111");
 System.out.println("919-555-1234");
 }
}

```

## 10. [2.10](#) *Variables:*

little

big

*Literals:*

2

2000

"The little number is "

"The big number is "

## 11. [2.11](#) The value is number

## 12. [2.12](#) 99bottles is illegal because it starts with a number.

r&d is illegal because the & character is illegal.

## 13. [2.13](#) They are not the same because one begins with an uppercase S while the other begins with a lowercase s. Variable names are case-sensitive.

## 14. [2.14](#)

1. short

2. `int`

3. 22.1 because it is stored as a `double`.

15. [2.15](#) `6.31E17`

16. [2.16](#) Append the `F` suffix to the numeric literal, such as:

```
number = 7.4F;
```

17. [2.17](#) `true` and `false`

18. [2.18](#)

```
1. char letter;
```

```
2. letter = 'A';
```

```
3. System.out.println(letter);
```

19. [2.19](#) The code for ‘C’ is 67.

The code for ‘F’ is 70.

The code for ‘W’ is 87.

20. [2.20](#) ‘B’ is a character literal.

21. [2.21](#) You cannot assign a string literal to a `char` variable. The statement should be:

```
char letter = 'Z';
```

22. [2.22](#)

| Expression | Value |
|------------|-------|
|------------|-------|

|                        |    |
|------------------------|----|
| <code>6 + 3 * 5</code> | 21 |
|------------------------|----|

|                         |   |
|-------------------------|---|
| <code>12 / 2 - 4</code> | 2 |
|-------------------------|---|

|                             |    |
|-----------------------------|----|
| <code>9 + 14 * 2 - 6</code> | 31 |
|-----------------------------|----|

|                             |   |
|-----------------------------|---|
| <code>5 + 19 % 3 - 1</code> | 5 |
|-----------------------------|---|



`(6 + 2) * 3`      `24`

`14 / (11 - 4)`      `2`

`9 + 12 * (8 - 3)` `69`

23. [2.23](#) Integer division. The value `23.0` will be stored in `portion`.

24. [2.24](#)

1. `x += 6;`

2. `amount -= 4;`

3. `y *= 4;`

4. `total /= 27;`

5. `x %= 7;`

25. [2.25](#)

1. No

2. Because the result of `basePay + bonus` results in an `int` value, which cannot be stored in the `short` variable `totalPay`. You can fix the problem by declaring `totalPay` as an `int`, or casting the result of the expression to a `short`.

26. [2.26](#) `a = (float)b;`

27. [2.27](#) `String city = "San Francisco";`

28. [2.28](#) `stringLength = city.length();`

29. [2.29](#) `oneChar = city.charAt(0);`

30. [2.30](#) `upperCity = city.toUpperCase();`

31. [2.31](#) `lowerCity = city.toLowerCase();`

32. [2.32](#) To write a single line comment you begin the comment with `//`. To write a multi-line comment you begin the comment with `/*` and end it with `*/`. To write a documentation comment you begin the comment with `/**` and end it with `*/`.

33. [2.33](#) Documentation comments can be read and processed by a program named `javadoc`. The

`javadoc` program can read Java source code files and generate attractively formatted HTML documentation files. If the source code files contain any documentation comments, the information in the comments becomes part of the HTML documentation.

34. [2.34](#) A message dialog box is used to display a message to the user. An input dialog box is used to gather keyboard input from the user.

35. [2.35](#)

```
1. JOptionPane.showMessageDialog(null, "Greetings Earthling.");
```

```
2. str = JOptionPane.showInputDialog("Enter a number.");
```

36. [2.36](#)

```
String str;
int age;
str = JOptionPane.showInputDialog("Enter your age.");
age = Integer.parseInt(str);
```

37. [2.37](#) `import javax.swing.JOptionPane;`

# Chapter 3

## 1. [3.1](#)

```
if (y == 20)
 x = 0;
```

## 2. [3.2](#)

```
if (hours > 40)
 payRate *= 1.5;
```

## 3. [3.3](#)

```
if (sales >= 10000)
 commission = 0.2;
```

## 4. [3.4](#)

```
if (max)
 fees = 50;
```

## 5. [3.5](#)

```
if (x > 100)
{
 y = 20;
 z = 40;
}
```

## 6. [3.6](#)

```
if (a < 10)
{
 b = 0;
 c = 1;
}
```

## 7. [3.7](#)

```
if (myCharacter == 'D')
 System.out.println("Goodbye");
```

## 8. [3.8](#)

```
if (x > 100)
 y = 20;
else
 y = 0;
```

9. [3.9](#)

```
if (y == 100)
 x = 1;
else
 x = 0;
```

10. [3.10](#)

```
if (sales >= 50000.0)
 commission = 0.2;
else
 commission = 0.1;
```

11. [3.11](#)

```
if (a < 10)
{
 b = 0;
 c = 1;
}
else
{
 b = -99;
 c = 0;
}
```

12. [3.12](#)<sub>1 1</sub>

13. [3.13](#)

If the customer purchases this many books . . . this many coupons are given.

|   |   |
|---|---|
| 1 | 1 |
| 2 | 1 |
| 3 | 2 |
| 4 | 2 |
| 5 | 3 |

14. [3.14](#)

```
if (amount1 > 10)
{
 if (amount2 < 100)
 {
 if (amount1 > amount2)
 System.out.println(amount1);
 else
 System.out.println(amount2);
 }
}
```

15. [3.15](#)

```
if (x > 0)
{
 if (y < 20)
 {
 z = 1;
 }
 else
 {
 z = 0;
 }
}
```

16. [3.16](#)**Logical Expression Result (true or false)**

true && false      false

true && true      true

false && true      false

false && false      false

true || false      true

true || true      true

false || true      true

false || false    false

!true                false

!false               true

17. [3.17](#) T, F, T, T, T

18. [3.18](#)

```
if (speed >= 0 && speed <= 200)
 System.out.println("The number is valid");
```

19. [3.19](#)

```
if (speed < 0 || speed > 200)
 System.out.println("The number is not valid");
```

20. [3.20](#)

```
if (name.equals("Timothy"))
 System.out.println("Do I know you?");
```

21. [3.21](#)

```
if (name1.compareTo(name2) < 0)
 System.out.println(name1 + " " + name2);
else
 System.out.println(name2 + " " + name1);
```

22. [3.22](#)

```
if (name.equalsIgnoreCase("Timothy"))
 System.out.println("Do I know you?");
```

23. [3.23](#)

1. `z = x > y ? 1 : 20;`

2. `population = temp > 45 ? base * 10 : base * 2;`

3. `wages = hours > 40 ? wages * 1.5 : wages * 1;`

4. `System.out.println(result >=0 ? "The result is positive" :  
"The result is negative");`

24. [3.24](#)

```
// Here is the switch statement.
switch(userNum)
{
 case 1 : System.out.println("One");
 break;
 case 2 : System.out.println("Two");
 break;
 case 3 : System.out.println("Three");
 break;
 default: System.out.println("Error: invalid number.");
}
```

25. [3.25](#)

```
switch(selection)
{
 case 'A' : System.out.println("You selected A.");
 break;
 case 'B' : System.out.println("You selected B.");
 break;
 case 'C' : System.out.println("You selected C.");
 break;
 case 'D' : System.out.println("You selected D.");
 break;
 default : System.out.println("Not good with letters, eh?");
}
```

26. [3.26](#) Because it uses greater-than and less-than operators in the comparisons.

27. [3.27](#) The case expressions must be a literal or a final variable, which must be of the char, byte, short, int, or String types. In this code, relational expressions are used.

28. [3.28](#) That is serious.

29. [3.29](#) `System.out.printf("%.2f", number);`

30. [3.30](#) `System.out.printf("%10.1f", number);`

31. [3.31](#) `System.out.printf("%08.1f", number);`

32. [3.32](#) `System.out.printf("%,10d", number);`

33. [3.33](#) `System.out.printf("%-,20.2f", number);`

34. [3.34](#) `System.out.printf("%20s", name);`

# Chapter 4

1. [4.1](#)
  1. 2
  2. 2
  3. 1
  4. 8
2. [4.2](#) 0 times
3. [4.3](#) This must be a trick question. The statement that prints “I love Java programming!” is not in the body of the loop. The only statement that is in the body of the loop is the one that prints “Hello World”. Because there is no code in the loop to change the contents of the `count` variable, the loop will execute infinitely, printing “Hello World”. So, “I love Java programming!” is never printed.

4. [4.4](#)

```
Scanner keyboard = new Scanner(System.in);
System.out.print("Enter a number in the range of 10 - 24: ");
number = keyboard.nextInt();
while (number < 10 || number > 24)
{
 System.out.println("That number is not in the range.");
 System.out.print("Enter a number in the range of 10 - 24: ");
 number = keyboard.nextInt();
}
```

5. [4.5](#)

```
String input;
Scanner keyboard = new Scanner(System.in);
System.out.print("Enter Y, y, N, or n: ");
input = keyboard.nextLine();
ch = input.charAt(0);
while (ch != 'Y' && ch != 'y' && ch != 'N' && ch != 'n')
{
 System.out.println("Try again.");
 System.out.print("Enter Y, y, N, or n: ");
 input = keyboard.nextLine();
 ch = input.charAt(0);
}
```

6. [4.6](#)



```
String str;
Scanner keyboard = new Scanner(System.in);
System.out.print("Enter Yes or No: ");
str = keyboard.nextLine();
while ((!str.equals("Yes")) && (!str.equals("No")))
{
 System.out.print("Please enter Yes or No: ");
 str = keyboard.nextLine();
}
```

7. [4.7](#) Initialization, test, and update.

8. [4.8](#)

1. count = 1
2. count <= 50
3. count++
4. for (int count = 1; count <= 50; count++)  
System.out.println("I love to program");

9. [4.9](#)

1.
 

|    |
|----|
| 0  |
| 2  |
| 4  |
| 6  |
| 8  |
| 10 |
2.
 

|    |
|----|
| -5 |
| -4 |
| -3 |
| -2 |
| -1 |
| 0  |
| 1  |
| 2  |
| 3  |
| 4  |
3.
 

|    |
|----|
| 5  |
| 8  |
| 11 |
| 14 |
| 17 |

10. [4.10](#)

```
// This is how Chloe Ashlyn would write the loop.
for (int i = 1; i <= 10; i++)
```

```
System.out.println("Chloe Ashlyn");
```

11. [4.11](#)

```
for (int i = 1; i <= 49; i += 2)
 System.out.println(i);
```

12. [4.12](#)

```
for (int i = 0; i <= 100; i += 5)
 System.out.println(i);
```

13. [4.13](#)

```
Scanner keyboard = new Scanner(System.in);
int number = 0, total = 0;
for (int i = 1; i <= 7; i++)
{
 System.out.print("Enter a number: ");
 number = keyboard.nextInt();
 total += number;
}
```

14. [4.14](#) The variable  $x$  is the loop control variable and  $y$  is the accumulator.

15. [4.15](#) You should be careful to choose a value that cannot be mistaken as a valid input value.

16. [4.16](#) Data is read from an input file, and data is written to an output file.

17. [4.17](#) `import java.io.*;`

18. [4.18](#) `PrintWriter`

19. [4.19](#)

```
// This is how Kathryn would write the code.
PrintWriter outputFile = new PrintWriter("MyName.txt");
outputFile.println("Kathryn");
outputFile.close();
```

20. [4.20](#) `File` and `Scanner`

21. [4.21](#)

```
File file = new File("MyName.txt");
Scanner inputFile = new Scanner(file);
if (inputFile.hasNext())
{
 String str = inputFile.nextLine();
 System.out.println(str);
}
inputFile.close();
```

22. [4.22](#) You create an instance of the `FileWriter` class to open the file. You pass the name of the file (a string) as the constructor's first argument, and the boolean value `true` as the second argument. Then, when you create an instance of the `PrintWriter` class, you pass a reference to the `FileWriter` object as an argument to the `PrintWriter` constructor. The file will not be erased if it already exists and new data will be written to the end of the file.
23. [4.23](#) throws `IOException`
24. [4.24](#) It assigns the variable `x` a random integer in the range `-2,147,483,648` to `+2,147,483,647`.
25. [4.25](#) It assigns the variable `x` a random integer in the range 0 through 99.
26. [4.26](#) It assigns the variable `x` a random integer in the range 1 through 9.
27. [4.27](#) It assigns the variable `x` a random number in the range 0.0 to 1.0.

# Chapter 5

1. [5.1](#) A value-returning method returns a value back to the code that called it, and a `void` method does not.
2. [5.2](#) Method call
3. [5.3](#) Method header
4. [5.4](#) If the user enters 5 the program will display *Able was I*. If the user enters 10 the program will display *I saw Elba*. If the user enters 100 the program will display *I saw Elba*.
5. [5.5](#)

```
// This is how Mary Catherine Jones would write it.
public static void myName()
{
 System.out.println("Mary Catherine Jones");
}
```

6. [5.6](#) An argument is a value that is passed into a method. A parameter is a special variable that holds the value being passed into the method.
7. [5.7](#) `b` and `c` will cause a compiler error because the values being sent as arguments (a `double` and a `long`) cannot be automatically converted to an `int`.
8. [5.8](#) Only `d` is written correctly.
9. [5.9](#) Only a copy of an argument's value is passed into a parameter variable. A method's parameter variables are separate and distinct from the arguments that are listed inside the parentheses of a method call. If a parameter variable is changed inside a method, it has no effect on the original argument.
10. [5.10](#)  


```
99 1.5
99 1.5
0 0.0
99 1.5
```
11. [5.11](#) `double`
12. [5.12](#) `public static int days(int years, int months, int weeks)`
13. [5.13](#) `public static double distance(double rate, double time)`

14. [5.14](#)public static long lightYears(long miles)

# Chapter 6

1. [6.1](#) To store data
2. [6.2](#) Operations that the object can perform
3. [6.3](#) It serves a purpose similar to that of the blueprint for a house. The blueprint itself is not a house, but rather a detailed description of a house. When we use the blueprint to build an actual house, we could say that we are building an instance of the house described by the blueprint. A class is not an object, but a description of an object. When a program is running, it can use the class to create, in memory, as many objects of a specific type as needed. Each object that is created from a class is called an instance of the class.
4. [6.4](#) In the Java API
5. [6.5](#) The `new` operator creates an object in memory, and returns that object's memory address.
6. [6.6](#) It holds an object's memory address.
7. [6.7](#) In order to fly a kite, you need a spool of string attached to it. When the kite is airborne, you use the spool of string to hold on to the kite and control it. This is similar to the relationship between an object and the variable that references the object. Think of the object as the kite, and the reference variable as the spool of string.
8. [6.8](#) Classes are the blueprints.
9. [6.9](#) The kite represents an object.
10. [6.10](#) The memory address of the object.
11. [6.11](#) A `String` object.
12. [6.12](#)
  1. `Car`
  2. `make` and `yearModel`
  3. `setMake`, `setYearModel`, `getMake`, and `getYearModel`
  4. `make` and `yearModel`
  5. `setMake`, `setYearModel`, `getMake`, and `getYearModel`
13. [6.13](#) `limo.setMake("Cadillac");`

14. [6.14](#) Creates an instance of an object in memory.
15. [6.15](#) An accessor is a method that gets a value from a class's field but does not change it. A mutator is a method that stores a value in a field or in some other way changes the value of a field.
16. [6.16](#) When the value of an item is dependent on other data and that item is not updated when the other data is changed, it is said that the item has become stale.
17. [6.17](#)



#### [L.6-1 Full Alternative Text](#)

18. [6.18](#) It has the same name as the class.
19. [6.19](#) It has no return type, not even `void`.
20. [6.20](#)
  1. `ClassAct`
  2. `ClassAct myact = new ClassAct(25);`
21. [6.21](#) Overloaded methods must have different parameter lists. Their return types do not matter.
22. [6.22](#) A method's signature consists of the method's name and the data types of the method's parameters, in the order that they appear.
23. [6.23](#) This is the second version of the method.  
`This is the first version of the method.`
24. [6.24](#) Only one.
25. [6.25](#) The problem domain is the set of real-world objects, parties, and major events related to a problem.

26. [6.26](#) Someone who has an adequate understanding of the problem. If you adequately understand the nature of the problem you are trying to solve, you can write a description of the problem domain yourself. If you do not thoroughly understand the nature of the problem, you should have an expert write the description for you.
27. [6.27](#) Identify all the nouns (including pronouns and noun phrases) in the problem domain description. Each of these is a potential class. Then, refine the list to include only the classes that are relevant to the problem.
28. [6.28](#) A class's responsibilities are the things that the class is responsible for knowing and the actions that the class is responsible for doing.
29. [6.29](#) It is often helpful to ask the questions "In the context of this problem, what must the class know? What must the class do?"
30. [6.30](#) No. Often responsibilities are discovered through brainstorming.



# Chapter 7

## 1. [7.1](#)

```
1. int[] employeeNumbers = new int[100];
```

```
2. double[] payRates = new double[25];
```

```
3. float[] miles = new float[14];
```

```
4. char[] letters = new char[1000];
```

2. [7.2](#) An array's size declarator must be a non-negative integer expression. The first statement is incorrect because the size declarator is negative. The second statement is incorrect because the size declarator is a floating-point number.

## 3. [7.3](#) 0 through 3

4. [7.4](#) The size declarator specifies the number of elements in the array. A subscript identifies a specific element in the array.

5. [7.5](#) The subscript is outside the range of valid subscripts for the array.

6. [7.6](#) When the statement executes, it crashes the program and displays a runtime error message.

## 7. [7.7](#)

```
1
2
3
4
5
```

8. [7.8](#) `double[] array = { 1.7, 6.4, 8.9, 3.1, 9.2 };`

There are five elements in the array.

9. [7.9](#) `result = numbers1[0] * numbers2[3];`

## 10. [7.10](#)

```
for (int i = 0; i < array.length; i++)
 array[i] = -1;
```

## 11. [7.11](#)

```
// Assume keyboard references a Scanner object.
```

```
int size;
System.out.print("Enter the size of the array: ");
size = keyboard.nextInt();
values = new double[size];
```

## 12. [7.12](#)

```
for (int i = 0; i < a.length; i++)
 b[i] = a[i];
```

## 13. [7.13](#) myMethod(numbers);

## 14. [7.14](#)

```
public static void zero(int[] array)
{
 for (int i = 0; i < array.length; i++)
 array[i] = 0;
}
```

## 15. [7.15](#)

1. String[] planets = { "Mercury", "Venus", "Earth", "Mars" };
2. for (int i = 0; i < planets.length; i++)
 System.out.println(planets[i]);
3. for (int i = 0; i < planets.length; i++)
 System.out.println(planets[i].charAt(0));

## 16. [7.16](#)

```
Rectangle[] rectArray = new Rectangle[5];
for (int i = 0; i <= rectArray.length; i++)
{
 // Initialize each rectangle with the values
 // i and i+1 for length and width.
 rectArray[i] = new Rectangle(i, i+1);
}
```

## 17. [7.17](#)

```
final int RACKS = 50;
final int SHELVES = 10;
final int VIDEOS = 25;
// Create an array to hold video numbers.
int videoNumbers[][][] = new int[RACKS][SHELVES][VIDEOS];
```

## 18. [7.18](#) The selection sort first looks for the smallest value in the array. When it finds it, it moves it to element 0.

## 19. [7.19](#) Only once.

20. [7.20](#) The sequential search steps through each element of the array, starting at element 0, looking for the search value. The binary search requires that the array be sorted in ascending order. It starts by looking at the middle element. If it is not the search value, and is greater than the search value, then the lower half of the array is searched next. If the middle element is not the search value, and is less than the search value, the upper half of the array is searched next. This same technique is repeated on the half of the array being searched until the element is either found or there are no more elements to search.
21. [7.21](#) 10,000 comparisons
22. [7.22](#) Move the items that are frequently searched for to the beginning of the array.
23. [7.23](#) `import java.util.ArrayList;`
24. [7.24](#) `ArrayList frogs = new ArrayList();`
25. [7.25](#) `ArrayList<String> lizards = new ArrayList<String>();`
26. [7.26](#) You use the `add` method.
27. [7.27](#) The `ArrayList` class has a `remove` method that removes an item at a specific index. You pass the index as an argument to the method.
28. [7.28](#) The `ArrayList` class has a `get` method that retrieves an item at a specific index. You pass the index as an argument to the method.
29. [7.29](#) Inserting means adding an item at a specific index. The `ArrayList` class has an overloaded version of the `add` method that allows you to add an item at a specific index.
30. [7.30](#) The `ArrayList` class has a `size` method that returns the number of items stored in the `ArrayList`.
31. [7.31](#) An `ArrayList`'s size is the number of items stored in the `ArrayList` object. An `ArrayList`'s capacity is the number of items the `ArrayList` object can hold without having to increase its size.

# Chapter 8

1. [8.1](#) Each instance of a class has its own copy of the class's instance fields. A static field does not belong to any instance of the class, and there is only one copy of a static field in memory, regardless of the number of instances of the class.
2. [8.2](#) It isn't necessary for an instance of the class to be created in order to execute the method.
3. [8.3](#) They cannot directly refer to non-static members of the class. This means that any method called from a static method must also be static. It also means that if the method uses any of the class's fields, they must be static as well.
4. [8.4](#) The `this` variable will reference the `stock2` object.
5. [8.5](#)
  1. `Flower`
  2. The ordinal value for `Rose` is 0. The ordinal value for `DAISY` is 1. The ordinal value for `PETUNIA` is 2.
  3. `Flower.Rose, Flower.DAISY, Flower.PETUNIA`
  4. `Flower flora = Flower.PETUNIA;`
6. [8.6](#) `HOBBIT ELF DRAGON`
7. [8.7](#) `z` is not greater than `x`.

# Chapter 9

1. [9.1](#) `little = Character.toLowerCase(big);`

2. [9.2](#)

```
if (Character.isDigit(ch))
 System.out.println("digit");
else
 System.out.println("Not a digit");
```

3. [9.3](#)<sub>A</sub>

4. [9.4](#)

```
String input;
char ch;
Scanner keyboard = new Scanner(System.in);
System.out.print("Do you want to repeat the " +
 "program or quit? (R/Q)");
input = keyboard.nextLine();
ch = input.charAt(0);
ch = Character.toUpperCase(ch);
while (ch != 'R' && ch != 'Q')
{
 System.out.print("Do you want to repeat the " +
 "program or quit? (R/Q)");
 input = keyboard.nextLine();
 ch = input.charAt(0);
 ch = Character.toUpperCase(ch);
}
```

5. [9.5](#) \$

6. [9.6](#)

```
int total = 0;
for (int i = 0; i < str.length(); i++)
{
 if (Character.isUpperCase(str.charAt(i)))
 total++;
}
```

7. [9.7](#)

```
public static boolean endsWithGer(String str)
{
 boolean status;

 if (str.endsWith("ger"))
 status = true;
```

```

else
 status = false;
return status;
}

```

## 8. [9.8](#)

```

public static boolean endsWithGer(String str)
{
 boolean status;
 String strUpper = str.toUpperCase();

 if (strUpper.endsWith("GER"))
 status = true;
 else
 status = false;
 return status;
}

```

## 9. [9.9](#) You would use the `substring` method.

10. [9.10](#) The `indexOf` method searches for a character or substring, starting at the beginning of a string. The `lastIndexOf` method searches for a character or substring, starting at the end of a string. This means that the `indexOf` method returns the index of the first occurrence of a character or substring, and the `lastIndexOf` method returns the index of the last occurrence of a character or substring.

11. [9.11](#) The `substring` method returns a reference to a substring. The `getChars` method stores a substring in a `char` array.

12. [9.12](#) The `concat` method.

13. [9.13](#) The `toCharArray` method returns all of the characters in the calling object as a `char` array.

14. [9.14](#) The fellow student is wrong. The `replace` method will return a reference to a `String` which is a copy of `str1`, in which all of the `o` characters have been replaced with `u` characters. The original string in `str1` will not be changed, however. The code will produce the following output:

```

To be, or not to be
Tu be, ur nut tu be

```

15. [9.15](#) `WilliamtheConqueror`

16. [9.16](#) `str = String.valueOf(number);`

17. [9.17](#) Once you create a `String` object, you cannot change the object's value.

18. [9.18](#) It would be more efficient to use `StringBuilder` objects because they are not immutable. Making a change to a `StringBuilder` object does not cause a new object to be created.

19. [9.19](#) `StringBuilder city = new StringBuilder("Asheville");`
20. [9.20](#) The `append` method.
21. [9.21](#) The `insert` method.
22. [9.22](#) The `deleteCharAt` method.
23. [9.23](#) The `setCharAt` method.
24. [9.24](#) While the `String` class's `replace` method replaces the occurrences of one character with another character, the `StringBuilder` class's `replace` method replaces a specified substring with a string.
25. [9.25](#) The tokens are “apples”, “pears”, and “bananas”. The delimiter is the space character.
26. [9.26](#) `4` will be stored in `x`.  
  
`"one"` will be stored in `stuff`.
27. [9.27](#)  
  

```
String str = "/home/rjones/mydata.txt"
String[] tokens = str.split("/");
```
28. [9.28](#)  
  


```
String str = "dog$cat@bird%squirrel"
String[] tokens = str.split("[$@%]");
```
29. [9.29](#) `str = Integer.toString(i);`
30. [9.30](#) The `toBinaryString`, `toHexString`, and `toOctalString` methods are static members of the `Integer` and `Long` wrapper classes.
31. [9.31](#) These variables hold the minimum and maximum values for a particular data type.

# Chapter 10

1. [10.1](#) `Vehicle` is the superclass and `Truck` is the subclass.

2. [10.2](#)

1. `Shape` is the superclass, `Circle` is the subclass.



2.

## [L.10-2 Full Alternative Text](#)

3. `setArea`, `getArea`, `setRadius`, `getradius`

4. `area`

5.

`c.setRadius(10.0);` legal

`s.setRadius(10.0);` illegal

`System.out.println(c.getArea());` legal

`System.out.println(s.getArea());` legal



3. [10.3](#) The superclass constructor (class A) will execute first, then the subclass constructor (class B) will execute.

4. [10.4](#)

```
You are on the ground.
You are in the sky.
```

5. [10.5](#)

```
The ground is green
The sky is blue
```

6. [10.6](#) When the superclass method is inadequate for the subclass's purpose.

7. [10.7](#) A subclass may call an overridden superclass method by prefixing its name with the `super` key word and a dot (`.`).

8. [10.8](#) It overrides the superclass method.

9. [10.9](#) It overloads the superclass method.

10. [10.10](#) You declare the method with the `final` modifier.

11. [10.11](#) Protected members of class may be accessed by methods in a subclass, and by methods in the same package as the class.

12. [10.12](#) Private members may be accessed only by methods in the same class. A protected member of a class may be directly accessed by methods of the same class or methods of a subclass. In addition, protected members may be accessed by methods of any class that are in the same package as the protected member's class.

13. [10.13](#) Because any other class that inherits from the class, or is in the same package, has unrestricted access to the protected member.

14. [10.14](#) Private members may be accessed only by methods in the same class. Any method in the same package as the class may directly access the class's members that have package access.

15. [10.15](#) When you accidentally leave out the access specifier, the member will have package access.

16. [10.16](#) `ClassD` still inherits from `Object`, because all classes ultimately inherit from `Object`.

17. [10.17](#) Because those methods are members of the `Object` class.

18. [10.18](#)

1. Legal, because a `Cube` is a `Rectangle`.

2.

`System.out.println(r.getLength());`      Legal

`System.out.println(r.getWidth());`      Legal

`System.out.println(r.getHeight());`      Illegal

`System.out.println(r.getSurfaceArea());`      Illegal

3. Illegal, because a `Rectangle` is not a `Cube`. The `Cube` class has capabilities beyond those of the `Rectangle` class, so a `Cube` variable cannot reference a `Rectangle` object.

19. [10.19](#) Abstract methods are used to ensure that a subclass implements the method.
20. [10.20](#) Override the abstract method.
21. [10.21](#) An abstract class serves as a superclass for other classes. It represents the generic or abstract form of all the classes that inherit from it.
22. [10.22](#) An abstract class cannot be instantiated. It must serve as a superclass.
23. [10.23](#) To specify behavior for other classes.
24. [10.24](#) It cannot be instantiated.
25. [10.25](#) An interface only specifies methods; it does not define them. In addition, all members of an interface are public.
26. [10.26](#) As `final` and `static`.
27. [10.27](#) `public class Customer implements Relatable`
28. [10.28](#) `public class Employee implements Payable, Listable`

# Chapter 11

1. [11.1](#) An exception is an object that is generated as the result of an error or an unexpected event.
2. [11.2](#) To throw an exception means to generate an exception object.
3. [11.3](#) Unless an exception is detected by the application and dealt with, it causes the application to halt.
4. [11.4](#) The `Throwable` class.
5. [11.5](#) Classes that extend the `Error` class are for exceptions that are thrown when a critical error occurs, such as an internal error in the Java Virtual Machine or running out of memory. An application should not try to handle these exceptions. Exception classes that extend the `Exception` class are general exceptions that an application can handle.
6. [11.6](#) A try block is one or more statements that are executed and can potentially throw an exception. A catch block is code that appears immediately after a `catch` clause, and is executed in response to a particular exception.
7. [11.7](#) The program will resume with the code that appears after the entire `try/catch` construct.
8. [11.8](#) Each exception object has a method named `getMessage` that can be used to retrieve the error message for the exception.
9. [11.9](#) When an exception is thrown by code in the try block, the JVM begins searching the `try` statement for a `catch` clause that can handle it. It searches the `catch` clauses from top to bottom and passes control of the program to the first `catch` clause with a parameter that is compatible with the exception.
10. [11.10](#) Statements in the finally block are always executed after the try block has executed and after any catch blocks have executed if an exception was thrown. The finally block executes whether an exception is thrown or not.
11. [11.11](#) A call stack is an internal list of all the methods that are currently executing.  
  
A stack trace is a list of all the methods in the call stack.
12. [11.12](#) Because method B does not handle the exception, control is passed back to method A. Method A doesn't handle the exception either, so control is passed back to method `main`. Because `main` doesn't handle the exception, the program is halted and the default exception handler handles the exception.
13. [11.13](#) Unchecked exceptions are those that are extended from the `Error` class or the

`RuntimeException` class. A program should not attempt to handle unchecked exceptions. All the remaining exceptions are checked exceptions. Either a checked exception must be handled by the program, or applicable `throws` clauses must be used with all of the methods that can potentially throw the checked exception.

14. [11.14](#) When code in the method can potentially throw a checked exception, but does not handle the exception.
15. [11.15](#) The `throw` statement causes an exception object to be created and thrown.
16. [11.16](#) The argument contains a custom error message that can be retrieved from the exception object's `getMessage` method. If you do not pass a message to the constructor, the exception will have a null message.
17. [11.17](#) The `throw` statement causes an exception to be thrown. The `throws` clause informs the compiler that a method throws one or more exceptions.
18. [11.18](#) No. If the method throws an unchecked exception, it does not have to have a `throws` clause.
19. [11.19](#) You can create a checked exception by extending its class from `Exception`. You can create an unchecked exception by extending its class from `Error`.
20. [11.20](#) All of the data stored in a text file is formatted as text. Numeric data is converted to text when it is stored in a text file. You can view the data stored in a text file by opening it with a text editor, such as Notepad. In a binary file, data is not formatted as text. Subsequently, you cannot view a binary file's contents with a text editor.
21. [11.21](#) To write data to a binary file you use the `FileOutputStream` and `DataOutputStream` classes. To read data from a file you use the `FileInputStream` and `DataInputStream` classes.
22. [11.22](#) With sequential access, when a file is opened for input, its read position is at the very beginning of the file. This means that the first time data is read from the file, the data will be read from its beginning. As the reading continues, the file's read position advances sequentially through the file's contents. In random file access, a program may immediately jump to any location in the file without first reading the preceding bytes.
23. [11.23](#) `RandomAccessFile`
24. [11.24](#) The two modes are "r" for reading, and "rw" for reading and writing. When a file is opened with "r" as the mode, the program can only read from the file. When a file is opened with "rw" as the mode, the program can read from the file and write to it.
25. [11.25](#) The class must implement the `Serializable` interface. You can then use the `ObjectOutputStream` class's `writeObject` method to serialize objects of the class.

# Chapter 12

1. [12.1](#) A computer's user interface is the part of the computer with which the user interacts.
2. [12.2](#) A command line interface, which is also known as a console interface, requires the user to type commands. If a command is typed correctly, it is executed and the results are displayed. If a command is not typed correctly, an error message is displayed.
3. [12.3](#) In a text-based environment, such as a command line interface, programs determine the order in which things happen.
4. [12.4](#) A program that responds to the user's actions, such as the clicking of a button.
5. [12.5](#) A standard Java library for developing rich applications that employ graphics.
6. [12.6](#) A scene
7. [12.7](#) A tree-like hierarchical data structure that holds the objects in a GUI.
8. [12.8](#) Root, branch, and leaf.
9. [12.9](#) Branch nodes can have children. Leaf nodes cannot have children.
10. [12.10](#) `Application`
11. [12.11](#) The `launch` method creates a `Stage` object that will be the application's window. It then calls the `start` method, passing a reference to the `Stage` object as an argument.
12. [12.12](#) `start` is the entry point for the application.
13. [12.13](#) Set the text that is to be displayed in the window's title var.
14. [12.14](#) It displays the window.
15. [12.15](#) An `HBox` arranges controls in a single horizontal row. A `VBox` arranges controls in a single vertical row.
16. [12.16](#) `GridPane`
17. [12.17](#) `javafx.scene`
18. [12.18](#) With the `setAlignment` method
19. [12.19](#) `javafx.geometry`

20. [12.20](#) `javafx.scene.image`
21. [12.21](#) An `ImageView` object
22. [12.22](#) You call the `ImageView` class's `setFitWidth` and `setFitHeight` methods.
23. [12.23](#) The `ImageView` class's `setPreserveRatio` method.
24. [12.24](#) You call the `ImageView` class's `setImage` method.
25. [12.25](#) Spacing is space that appears between controls that are inside the container, and padding is space that appears around the inside edge of the container.
26. [12.26](#) `javafx.geometry`
27. [12.27](#) With the container's `setHgap` and `setVgap` methods
28. [12.28](#) You call the container's `add` method. The arguments are: the control being added, the column, and the row.
29. [12.29](#) An event is an action that takes place while a program is running. For example, each time the user clicks a `Button` control, an event takes place.
30. [12.30](#) An object that responds to events. If an event source is connected to an event handler, a specific method in the event handler is called and the event object is passed as an argument to the method.
31. [12.31](#) `javafx.event`
32. [12.32](#) You call the `Button` class's `setOnAction` method.
33. [12.33](#) `javafx.scene.control`
34. [12.34](#) You call the control's `getText` method. The method returns the value that has been entered into the `TextField` as a `String`.
35. [12.35](#)
  - Writing the definition of an inner class that implements the `EventHandler` interface. Then, instantiating that class, and registering it with a control.
  - Instantiating an anonymous inner class that implements the `EventHandler` interface, and registering the object with a control.
  - Using a lambda expression to instantiate an anonymous inner class that implements the `EventHandler` interface, and registering the object with a control.

36. [12.36](#) Top, bottom, left, center, and right

37. [12.37](#) Constructor #1: The no-arg constructor creates an empty `BorderPane` container.

Constructor #2: This constructor accepts one argument. The node that passed as the argument is placed in the `BorderPane`'s center region.

Constructor #3: This constructor accepts five nodes as arguments: one to place in each region.

# Chapter 13

1. [13.1](#)

```
.button
.text-field
.abel
.image-view
```

2. [13.2](#) `.menu-button`

3. [13.3](#) `-fx-font-family`

4. [13.4](#) `-fx-font-weight`

5. [13.5](#) `-fx-border-style`

6. [13.6](#) `scene.getStylesheets().add("stylesheet.css");`

7. [13.7](#) `.root`

8. [13.8](#) The `.root` definition

9. [13.9](#) red, green, blue

10. [13.10](#) blue

11. [13.11](#) An ID selector

12. [13.12](#) You add the `RadioButtons` to a `ToggleGroup`. Only one of the `RadioButton` controls in a `ToggleGroup` may be selected at any time.

13. [13.13](#) The `RadioButton` has a method, `isSelected`, that returns `true` if it is selected, or `false` otherwise.

14. [13.14](#) Call the `RadioButton`'s `setSelected` method, passing `true` as the argument.

15. [13.15](#) `ActionEvent`

16. [13.16](#) The `CheckBox` control has a method, `isSelected`, that returns `true` if it is selected, or `false` otherwise.

17. [13.17](#) Call the `CheckBox`'s `setSelected` method, passing `true` as the argument.

18. [13.18](#) `ActionEvent`



19. [13.19](#) With the control's `setPreferredSize` method.
20. [13.20](#) There are two ways: (1) You can use the `ListView` class's `getSelectionModel().getSelectedItem()` method get the item that is currently selected, or (2) You can use the `ListView` class's `getSelectionModel().getSelectedIndex()` method get the index of the item that is currently selected.
21. [13.21](#) In single selection mode, only one item can be selected at a time in a `ListView` control.
22. [13.22](#) In multiple interval selection mode, multiple items may be selected in a `ListView` control.
23. [13.23](#) You can set the orientation to either vertical or horizontal with the `ListView` class's `setOrientation` method. When you call the method, you pass it one of the following enum constants: `Orientation.VERTICAL` or `Orientation.HORIZONTAL`.
24. [13.24](#) `ObservableList`
25. [13.25](#) You can use the `ComboBox` class's `getValue()` method get the item that is currently selected.
26. [13.26](#) `ActionEvent`
27. [13.27](#) An editable `ComboBox` allows the user to select an item from a drop-down list, or type input into a field that is similar to a `TextField`. By default, `ComboBox` controls are uneditable, which means that the user cannot type input into the control; he or she can only select items from a drop-down list.
28. [13.28](#)
  1. `setMajorTickUnit(value)`
  2. `setMinorTickCount(value)`
  3. `setShowTickMarks(value)`
  4. `setShowTickLabels(value)`
29. [13.29](#) Change event
30. [13.30](#) You call the `Slider`'s `getValue()` method.
31. [13.31](#) A `TextField` can accept multiple lines of input.
32. [13.32](#) 40 columns. You can change the number of columns with the `setPrefColumnCount()` method.
33. [13.33](#) 10 rows. You can change the number of rows with the `setPrefRowCount()` method.

34. [13.34](#) When text is entered into the control and the end of a line is reached, the text wraps around to the next line. You can enable text wrapping with the `TextArea` class's `setWrapText` method.
35. [13.35](#) You call the control's `getText()` method.
36. [13.36](#)
- Menu Bar. At the top of the window, is a menu bar. The *menu bar* lists the names of one or more menus.
  - Menu. A *menu* is a drop-down list of menu items. The user may activate a menu by clicking on its name on the menu bar. In the figure, the *Edit* menu has been activated.
  - Menu Item. A *menu item* can be selected by the user. When a menu item is selected, some type of action is usually performed.
  - Check menu item. A *check menu item* appears may be selected or deselected. When it is selected, a check mark appears next to the item. When it is deselected, the check does not appear. Check menu items are normally used to turn an option on or off. The user toggles the state of a check menu item each time he or she selects it.
  - Radio menu item. A *radio menu item* may be selected or deselected. A radio menu item looks like a check menu item. When it is selected, a check mark appears next to the item. When it is deselected, the check does not appear. When a set of radio menu items are grouped with a `ToggleGroup` object, only one of them can be selected at a time. When the user selects a radio menu item, the one that was previously selected is deselected.
  - Submenu. A menu within a menu is called a *submenu*. Some of the commands on a menu are actually the names of submenus. You can tell when a command is the name of a submenu because a small right arrow appears to its right. Activating the name of a submenu causes the submenu to appear.
  - Separator bar. A separator bar is a horizontal bar used to separate groups of items on a menu. Separator bars are only used as a visual aid and cannot be selected by the user.
37. [13.37](#) `MenuBar`
38. [13.38](#) `Menu`
39. [13.39](#) `MenuItem`
40. [13.40](#) `RadioMenuItem`. Call the object's `setSelected` method to select it.
41. [13.41](#) Add them to a `ToggleGroup`.
42. [13.42](#) `CheckMenuItem`. Call the object's `setSelected` method to select it.

43. [13.43](#) `ActionEvent`
44. [13.44](#) `javafx.stage`
45. [13.45](#) Call the `FileChooser`'s `showOpenDialog` method.
46. [13.46](#) Call the `FileChooser`'s `showSaveDialog` method.
47. [13.47](#) The `showOpenDialog` and `showSaveDialog` methods return a `File` object containing information about the selected file.

# Chapter 14

1. [14.1](#) (0,0)
2. [14.2](#) (639, 479)
3. [14.3](#) In the Cartesian coordinate system, the *Y* coordinates decrease as you move downward. In the screen coordinate system, the *Y* coordinates increase as you move downward, toward the bottom of the screen.
4. [14.4](#) `javafx.scene.shape`
5. [14.5](#) `javafx.scene.paint`
6. [14.6](#) `javafx.scene.layout`
7. [14.7](#) `setStroke()`
8. [14.8](#) `setFill()`
9. [14.9](#) `setRotate()`
10. [14.10](#) `setScaleX()`
11. [14.11](#) `setScaleY()`
12. [14.12](#) `Triangle`
13. [14.13](#) `TranslateTransition`
14. [14.14](#) `RotateTransition`
15. [14.15](#) `ScaleTransition`
16. [14.16](#) `StrokeTransition`
17. [14.17](#) `FillTransition`
18. [14.18](#) `FadeTransition`
19. [14.19](#) With the `Duration` class
20. [14.20](#) `javafx.util`
21. [14.21](#) You call the transition class's `setInterpolator` method.

22. [14.22](#) `javafx.scene.effect`
23. [14.23](#) `DropShadow`
24. [14.24](#) `InnerShadow`
25. [14.25](#) `ColorAdjust`
26. [14.26](#) `BoxBlur`, `GaussianBlur`, and `MotionBlur`
27. [14.27](#) `SepiaTone`
28. [14.28](#) `Glow`
29. [14.29](#) `Reflection`
30. [14.30](#) `setInput`
31. [14.31](#) `Media` and `MediaPlayer`
32. [14.32](#) `javafx.scene.media`
33. [14.33](#) If you pass `true` to the method, the sound file will be immediately played once the `MediaPlayer` object is created and ready.
34. [14.34](#) It specifies the number of times that the sound file should be played.
35. [14.35](#) `Media`, `MediaPlayer`, and `MediaPlayer`
36. [14.36](#) `javafx.scene.media`
37. [14.37](#) With the `MediaView` class's `setFitWidth` and `setFitHeight` methods.
38. [14.38](#) `KEY_PRESSED`
39. [14.39](#) `KEY_RELEASED`
40. [14.40](#) `KEY_TYPED`
41. [14.41](#) `getCode()`
42. [14.42](#) The `getCode()` method returns a `KeyCode` value, which is an `enum` in the `javafx.scene.input` package. The `KeyCode` `enum` contains an extensive set of constants representing the keys on a computer keyboard.

# Chapter 15

1. [15.1](#) A recursive algorithm requires multiple method calls. Each method call requires several actions to be performed by the JVM. These actions include allocating memory for parameters and local variables, and storing the address of the program location where control returns after the method terminates. All of these actions are known as overhead. In an iterative algorithm, which uses a loop, such overhead is unnecessary.
2. [15.2](#) A case in which the problem can be solved without recursion.
3. [15.3](#) Cases in which the problem is solved using recursion.
4. [15.4](#) When it reaches the base case.
5. [15.5](#) In direct recursion, a recursive method calls itself. In indirect recursion, method A calls method B, which in turn calls method A.

# Chapter 16

1. [16.1](#) Bubble sort
2. [16.2](#) Insertion sort
3. [16.3](#) Selection sort
4. [16.4](#) Quicksort
5. [16.5](#) The sequential search algorithm simply uses a loop to step through each element of an array, comparing each element's value with the value being searched for. The binary search algorithm, which requires the values in the array to be sorted in order, starts searching at the element in the middle of the array. If the middle element's value is greater than the value being searched for, the algorithm next tests the element in the middle of the first half of the array. If the middle element's value is less than the value being searched for, the algorithm next tests the element in the middle of the last half of the array. Each time the array tests an array element and does not find the value being searched for, it eliminates half of the remaining portion of the array. This method continues until the value is found, or there are no more elements to test. The binary search is more efficient than the sequential search.
6. [16.6](#) 10,000
7. [16.7](#) 15
8. [16.8](#) The items frequently searched for can be stored near the beginning of the array.
9. [16.9](#) A basic operation is one that requires constant time, regardless of the size of the problem that is being solved.
10. [16.10](#) The worst case complexity function  $f(n)$  of an algorithm is a measure of the time required by the algorithm to solve a problem instance of size  $n$  that requires the most time.
11. [16.11](#) Because  $100n$  and  $25n$  differ by a constant factor and constant factors are not significant, the two algorithms are equivalent in efficiency.
12. [16.12](#) To say that  $f(n)$  is not in  $O(g(n))$  means that  $f(n)$  cannot be bounded by any constant multiple of  $g(n)$ .
13. [16.13](#) To say that  $f(n)$  is in  $O(g(n))$  means that there is a positive constant  $K$  such that  $f(n) \leq Kg(n)$  for all  $n \geq 1$ . This means that for large problem sizes, an algorithm with complexity function  $f(n)$  is no worse than one with complexity function  $g(n)$ .
14. [16.14](#) To show that  $100n^3 + 50n^2 + 75$  is in  $O(20n^3)$ , observe that

$$100n^3 + 50n^2 + 75 \cdot 20n^3 = 5 + 5 \cdot 2n + 75 \cdot 20n^3 \leq 5 + 5 + 75 \leq 85$$

for all  $n \geq 1$ . Take  $K=85$  in Equation 17.1.

15. [16.15](#) The problem definition and algorithm are as follows:

### MAIN ALGORITHM

INPUT: two integer arrays  $a[ ]$  and  $b[ ]$  of size  $n$

SIZE OF INPUT: The number  $n$  of array entries

OUTPUT: true if each element of  $a[ ]$  is also an element of  $b[ ]$ , false otherwise

```
boolean contained = true
int i = 0
While contained && i < n do
 contained = LINSEARCH(a[i], b)
 i++;
End While
print contained
```

### SEARCH ALGORITHM

INPUT: an integer  $X$  and an array  $b[ ]$  of size  $n$

SIZE OF INPUT: The number  $n$  of array entries

RETURN: true if  $X$  is contained in  $b[ ]$ , false otherwise

```
boolean LINSEARCH(int X, int [] b)
Begin
 int k = 0;
 boolean found = false;
 While !found && k < n do
 If (X == b[k])
 found = true;
 else
 k++;
 END If
 End While
 return found
End LINSEARCH
```

16. [16.16](#) See solution of 16.15 for the algorithm. In the worst case, the main algorithm makes  $n$  calls to the `LINSEARCH` procedure. Each call to `LINSEARCH` makes at most  $n^2$  comparisons. Therefore, the algorithm makes a total of comparisons.
17. [16.17](#) Let  $f(n)$  be in  $O(20)$ . Then there is a positive  $K$  such that  $f(n) \leq 20K$  for all  $n \geq 1$ . Let  $K_1 = 20K$ . Then  $f(n) \leq 1K_1$  for all  $n \geq 1$ , so  $f(n)$  is in  $O(1)$ .



18. [16.18](#) Assuming that  $g(n) \geq 1$  for all  $n \geq 1$ , we have  $100 \leq 100g(n)$  for all  $n \geq 1$ . This implies that  $g(n) + 100 \leq g(n) + 100g(n) = 101g(n)$  for all  $n \geq 1$ . Now if  $f(n)$  is in  $O(g(n) + 100)$ , there exists a positive  $K$  such that  $f(n) \leq K(g(n) + 100) \leq 101Kg(n)$  for all  $n \geq 1$ . Taking  $K_1 = 101K$ , we see that  $f(n) \leq K_1g(n)$  for all  $n \geq 1$ .

# Chapter 17

1. [17.1](#) Because it opens an opportunity for the application to throw an exception at run time. If we retrieve an object from the `ArrayList`, but cast it to the wrong type, an exception occurs.
2. [17.2](#) In non-generic mode, the `get` method returns an `Object` reference. When the `get` method is called, the compiler does not know the exact type of the object being returned, only that it is an `Object`. To assign the reference to any other type of variable we have to cast it to the appropriate type.
3. [17.3](#) `Rectangle` objects.
4. [17.4](#) Compile time.
5. [17.5](#) A type parameter is an identifier in a generic class declaration. It represents an actual type that will be used when a generic class is instantiated.
6. [17.6](#) Only reference types may be passed to a type parameter.
7. [17.7](#) Yes, it is possible. Doing so instantiates the class as a raw type. This is not recommended, however, because you give up the benefits of type-safety that you would have otherwise.
8. [17.8](#) No, it is not required, but it is the standard convention.
9. [17.9](#) That any type argument can be used in its place.
10. [17.10](#) It means that any type that extends `Number`, or is `Number`, can be passed as a type argument.
11. [17.11](#) It means that any type that is `Integer` or a superclass of `Integer` can be passed as a type argument.
12. [17.12](#) An upper boundary
13. [17.13](#) A lower boundary
14. [17.14](#) No, you cannot use both together.
15. [17.15](#) When you call a generic method, the compiler determines which type to use from the context in which you are using the method.
16. [17.16](#) To prevent the class from being instantiated with a type argument that does not support some operation that the class is performing.
17. [17.17](#) A generic class's type can be constrained with either an upper bound or a lower bound.

18. [17.18](#) A generic class can be a superclass and/or a subclass.
19. [17.19](#) Yes. The “is-a” relationship is in effect between the classes.
20. [17.20](#) Yes.
21. [17.21](#) Yes.
22. [17.22](#) You would write the type parameter definition as:

```
<T extends Comparable>
```

23. [17.23](#) When the compiler encounters a generic class, interface, or method with an unbound type parameter, such as `<T>` or `<E>`, it replaces all occurrences of the type parameter with `Object`.
24. [17.24](#) When the compiler encounters a class, interface, or method with a bound type parameter, such as `<T extends Number>` or `<E extends Comparable>`, it replaces all occurrences of the type parameter with the bound that is applied to the parameter.
25. [17.25](#) No.

# Chapter 18

1. [18.1](#) The `java.util` package.
2. [18.2](#) Lists, sets, and maps.
3. [18.3](#) Elements are stored in a list according to position, and a list can have duplicate values. A set has no concept of position and does not allow duplicates.
4. [18.4](#) Each element of a map consists of a key associated with a value.
5. [18.5](#) `List` and `Set` extend the `Collection` interface.
6. [18.6](#) `forEach()` and `iterator()`.
7. [18.7](#) An iterator is an object that can move through a collection and provide access to each element of that collection.
8. [18.8](#) Contiguous allocation.
9. [18.9](#) Linked allocation.
10. [18.10](#) So you can change the class that implements the interface without having to change the rest of the code.
11. [18.11](#) The `Iterator` interface.
12. [18.12](#) A `ListIterator` object can move forward and backward through a collection, while an `Iterator` object can only move forward.
13. [18.13](#) Call the `listIterator()` method.
14. [18.14](#) By calling the `hasNext()` method.
15. [18.15](#) The `hasPrevious()` method.
16. [18.16](#) The element returned by the last call to the `next()` method.
17. [18.17](#) The element returned by the last call to `next()` or `previous()`.
18. [18.18](#) The capacity is the length of the array that underlies the `ArrayList`, while the size is the number of elements currently stored.
19. [18.19](#) The `LinkedList` class.

20. [18.20](#) The `LinkedList` class.
21. [18.21](#) A hash code of an object is an integer value that is characteristic of that object.
22. [18.22](#) The `hashCode()` method.
23. [18.23](#) A collision occurs when two different objects hash to the same value.
24. [18.24](#) The duplicate item is not added to the set.
25. [18.25](#) The `hashCode()` method.
26. [18.26](#) The `LinkedHashSet` class.
27. [18.27](#) The `TreeSet` class.
28. [18.28](#) A comparator for a class is an object that can compare two objects of that class and determine whether they are equal, and if not, which comes before the other. A comparator implements the `Comparator` interface.
29. [18.29](#) Each element of map consists of a key and a value.
30. [18.30](#) The key.
31. [18.31](#) Insertion order and access order.
32. [18.32](#) By key order.
33. [18.33](#) A stream is an object that permits a pipeline of operations to be performed on a sequence of elements drawn from a source.
34. [18.34](#) An intermediate operation transforms a stream into another stream, while a terminal operation transforms a stream into a non-stream object or result.
35. [18.35](#) The `Stream` interface, the `Collection` interface, and the `Arrays` class all have methods for creating streams.
36. [18.36](#) A pipeline is a sequence of operations where the output of one operation becomes the input to the next operation in the sequence.
37. [18.37](#) A terminal operation and a reduction are the same thing.
38. [18.38](#) A collector is an object that performs a reduction on a stream.
39. [18.39](#) A stream pipeline can have at most one reduction, which must be the last operation in the pipeline.

# Chapter 19

## 1. [19.1](#)

```
Node greenNode = new Node("green");
greenNode.next = new Node("blue");
Node redNode = new Node("red", greenNode);
```

## 2. [19.2](#) BrownYellowPurple

## 3. [19.3](#)

```
// Let p point to the new Node
Node p = new Node(e, succ, succ.prev);
succ.prev = p;
if (p.prev == null)
 first = p;
else
 p.prev.succ = p;
```

## 4. [19.4](#)

```
last.next = new Node(e, null, last);
last = last.next;
```

# Chapter 20

1. [20.1](#) push
2. [20.2](#) pop
3. [20.3](#) Throw an exception
4. [20.4](#) Last in first out
5. [20.5](#) Some examples are a stack of plates or trays in a cafeteria, a line of cars parked single file in a narrow driveway, and a person wearing an undershirt, a shirt, a jacket, and an overcoat.
6. [20.6](#) enqueue
7. [20.7](#) First in first out
8. [20.8](#) Some examples are a grocery store checkout line, a line of cars going through a toll booth.
9. [20.9](#) By the time the add operation reaches the end of the array, the remove operation may have emptied the spaces of the beginning of the array, allowing the add operation to wrap around and start storing elements at the beginning. Similarly, when the remove empties spaces at the end of the array, it should wrap around to the beginning.

# Chapter 21

## 1. [21.1](#)

1. Ancestors of  $A$ :  $\{A\}$ ; ancestors of  $B$ :  $\{B, A\}$ ; ancestors of  $C$ :  $\{C, A\}$ ; ancestors of  $D$ :  $\{D, B, A\}$ ; ancestors of  $E$ :  $\{E, C, A\}$ ; ancestors of  $F$ :  $\{F, C, A\}$ ; ancestors of  $G$ :  $\{G, F, C, A\}$ .
2. Descendants of  $A$ :  $\{A, B, D\}$ ; descendants of  $B$ :  $\{B, D\}$ ; descendants of  $D$ :  $\{D\}$ ; descendants of  $C$ :  $\{C, E, F, G\}$ ; descendants of  $E$ :  $\{E\}$ ; descendants of  $F$ :  $\{F, G\}$ ; descendants of  $G$ :  $\{G\}$ .
3. Leaves of  $T$ :  $\{D, E, G\}$ .
4. Non-leaf nodes of  $T$ :  $\{A, B, C, F\}$ .
5. Preorder:  $\{A, B, D, C, E, F, G\}$ .
6. Postorder:  $\{D, B, E, G, F, C, A\}$ .
7. Inorder:  $\{B, D, A, E, C, G, F\}$ .

## 2. [21.2](#) The tree on the left : $x \ 12 +$


The tree on the right:  $x \ 12 + y * 2 +$

## 3. [21.3](#) The tree on the left: $+ \ x \ 12$

The tree on the right:  $+ \ * + x \ 12 \ y \ 2$

## 4. [21.4](#) Nodes in preorder: 160 10 20 50 60 180

Drawing of the tree:







# Chapter 22 (Available on the Book's Companion Website)

1. [22.1](#) Traditional text and binary files are not practical when a large amount of data must be stored and manipulated. Many businesses keep hundreds of thousands, or even millions, of data items in files. When a text or binary file contains this much data, simple operations such as searching, inserting, and deleting become cumbersome and inefficient.
2. [22.2](#) The DBMS handles the actual reading, writing, and searching of data. The Java application programmer needs only to know how to communicate with the DBMS.
3. [22.3](#) JDBC, or Java Database Connectivity
4. [22.4](#) SQL, or Structured Query Language
5. [22.5](#) A string listing the protocol that should be used to access a database, the name of the database, and potentially other items
6. [22.6](#) `jdbc:derby:InventoryDB`
7. [22.7](#) `DriverManager.getConnection`
8. [22.8](#) The data that is stored in a table is organized into rows and columns. A row is a complete set of information about a single item. The data that is stored in a row is divided into columns. Each column is an individual piece of information about the item.
9. [22.9](#) A primary key is a column that holds a unique value for each row, and can be used to identify specific rows.
10. [22.10](#)
  - `int`
  - `int`
  - `float`
  - `String`
  - `String`
  - `String`

- double

11. [22.11](#) A `ResultSet` is an object that contains the results of an SQL statement.

12. [22.12](#)

1. The table is `Account`.

2. The column is `Id`.

13. [22.13](#)

1. `SELECT * FROM Inventory`

2. `SELECT ProductID FROM Inventory`

3. `SELECT ProductID, QtyOnHand FROM Inventory`

4. `SELECT ProductID FROM Inventory`

`WHERE Cost < 17.00`

5. `SELECT * FROM Inventory`

`WHERE ProductID LIKE '%ZZ'`

14. [22.14](#) The `LIKE` operator can be used in a search criterion to search for a string that contains a pattern of characters.

15. [22.15](#) The `%` symbol represents any sequence of 0 or more characters. The underscore character (`_`) represents only a single character.

16. [22.16](#) You use the `ORDER BY` clause.

17. [23.17](#)

```
Connection conn = DriverManager.getConnection(DB_URL);
Statement stmt = conn.createStatement();
ResultSet result = stmt.executeQuery(sql);
stmt.close();
conn.close();
```

18. [22.18](#) You create a `Statement` object (`Statement` in an interface in `java.sql`), and you pass a string containing the `SELECT` statement to the `Statement` object's `executeQuery` method.

19. [22.19](#) The cursor initially points before the first row. To move it forward you call the `ResultSet` object's `next` method.

20. [22.20](#) You call the `ResultSet` object's `getString` method, passing 3 as an argument.

21. [22.21](#)

```
INSERT INTO Coffee
 (Description, ProdNum, Price)
VALUES
 ('Eastern Blend', '30-001', 18.95)
```

22. [22.22](#)

```
INSERT INTO Coffee
 (Description, ProdNum, Price)
VALUES
 ('Honduran Dark', '22-001', 8.65)
```

23. [22.23](#)

```
UPDATE Coffee
SET Price = 4.95
WHERE Description Like '%Decaf%'
```

24. [22.24](#) DELETE FROM Coffee WHERE Description LIKE '%Decaf%'

25. [22.25](#)

```
CREATE TABLE Book
(Publisher CHAR(25),
 Author, CHAR(25),
 Pages INTEGER,
 Isbn CHAR(10))
```

26. [22.26](#) DROP TABLE Book

# Chapter 23 (Available on the Book's Companion Website)

1. [23.1](#) A frame is a basic window that has a border around it, a title bar, and a set of buttons for minimizing, maximizing, and closing the window. In a Swing application, you create a frame from the `JFrame` class.
2. [23.2](#) With the `setSize` method.
3. [23.3](#) With the `setVisible` method.
4. [23.4](#) A content pane is a container that is part of every `JFrame` object. You cannot see the content pane and it does not have a border, but any component that is added to a `JFrame` must be added to its content pane.
5. [23.5](#) Panels cannot be displayed by themselves.
6. [23.6](#) It is an object that responds to events.
7. [23.7](#) The class must implement the `ActionListener` interface, and it must have a method named `actionPerformed`. The method is executed when the user clicks the button.
8. [23.8](#) With the `addActionListener` method.
9. [23.9](#) You use the `setBackground` method to set the color of a component, and the `setForeground` method to set the color of text displayed on a component.
10. [23.10](#) By calling the container's `setLayout` method and passing a reference to a layout manager object as the argument.
11. [23.11](#) `BorderLayout`
12. [23.12](#) `FlowLayout`
13. [23.13](#) `GridLayout`
14. [23.14](#) Only one, in both cases.
15. [23.15](#) By placing the component in a panel, then adding the panel to the region. The layout manager resizes the panel, not the component inside the panel.
16. [23.16](#) By calling the `pack` method.

17. [23.17](#) `BorderLayout` is the default layout manager for a `JFrame` object's content pane.  
  
`FlowLayout` is the default layout manager for a `JPanel` object.
18. [23.18](#) Radio buttons
19. [23.19](#) Check boxes
20. [23.20](#) A `ButtonGroup` object can contain radio buttons. It creates a mutually exclusive relationship between the radio buttons that it contains.
21. [23.21](#) Radio buttons.
22. [23.22](#) An action event.
23. [23.23](#) An item event.
24. [23.24](#) You use the `isSelected` method to determine whether a `JRadioButton` component is selected. The method returns a `boolean` value. If the check box is selected, the method returns `true`. Otherwise, it returns `false`.
25. [23.25](#) You use the `isSelected` method to determine whether a `JCheckBox` component is selected. The method returns a `boolean` value. If the check box is selected, the method returns `true`. Otherwise, it returns `false`.
26. [23.26](#) `setBorder`
27. [23.27](#) You should use the `BorderFactory` class to create them for you. The `BorderFactory` class has static methods that return various types of borders.

# Chapter 24 (Available on the Book's Companion Website)

1. [24.1](#) You call its `setEditable` method and pass `false` as the parameter. To store text in the text field, use the `setText` method.
2. [24.2](#) The index of the first item is 0. The index of the twelfth item would be 11.
3. [24.3](#) You use the `getSelectedValue` method to retrieve the selected item, and the `getSelectedIndex` method to get the index of the selected item.
4. [24.4](#) First, set the number of visible rows for the list component. Next, create a scroll pane object and add the list component to it.
5. [24.5](#) You use the `getSelectedItem` method to retrieve the selected item, and the `getSelectedIndex` method to get the index of the selected item.
6. [24.6](#) An uneditable combo box combines a button with a list and allows the user to select only items from its list. An editable combo box combines a text field and a list. In addition to selecting items from the list, the user may also type input into the text field. The default type of combo box is uneditable.
7. [24.7](#) To display an image, you first create an instance of the `ImageIcon` class, which can read the contents of an image file. Next, you display the image in a label by passing the `ImageIcon` object as an argument to the `JLabel` constructor. You can create a `JLabel` with both an image and text by creating the `JLabel` in the usual way, with the text passed to the constructor. Then, you can use the `setIcon` method to display an image.
8. [24.8](#) To display an image, you first create an instance of the `ImageIcon` class, which can read the contents of an image file. Next, you display the image in a button by passing the `ImageIcon` object as an argument to the `JButton` constructor. To create a button with an image and text, pass a `String` and an `ImageIcon` object to the constructor.
9. [24.9](#) `setIcon`
10. [24.10](#) A mnemonic is a key on the keyboard that you press in combination with the Alt key to quickly access a component such as a button. These are sometimes referred to as short-cut keys, or hot keys. You assign an access key to a component through the component's `setMnemonic` method.
11. [24.11](#) A tool tip is text that is displayed in a small box when the user holds the mouse cursor over a component. The box usually gives a short description of what the component does. You assign a tool tip to a component with the `setToolTipText` method.

## 12. [24.12](#)

1. **Menu bar.** At the top of the window, just below the title bar, is a menu bar. The menu bar lists the names of one or more menus.
2. **Menu item.** A menu item can be selected by the user. When a menu item is selected, some type of action is usually performed.
3. **Check box menu item.** A check box menu item appears with a small box beside it. The item may be selected or deselected. When it is selected, a check mark appears in the box. When it is deselected, the box appears empty. Check box menu items are normally used to turn an option on or off. The user toggles the state of a check box menu item each time he or she selects it.
4. **Radio button menu item.** A radio button menu item may be selected or deselected. A small circle appears beside it that is filled in when the item is selected and appears empty when the item is deselected. Like a check box menu item, a radio button menu item can be used to turn an option on or off. When a set of radio button menu items are grouped with a `ButtonGroup` object, only one of them can be selected at a time. When the user selects a radio button menu item, the one that was previously selected is deselected.
5. **Submenu.** A menu within a menu is called a submenu. Some of the commands on a menu are actually the names of submenus. You can tell when a command is the name of a submenu because a small right arrow appears to its right. Activating the name of a submenu causes the submenu to appear.
6. **Separator bar.** A separator bar is a horizontal bar used to separate groups of items on a menu. Separator bars are used only as a visual aid and cannot be selected by the user.

13. [24.13](#) The `JMenuItem` class. You pass a string, which is displayed on the menu item.

14. [24.14](#) The `JRadioButtonMenuItem` class. You pass a string to the constructor, which is displayed on the menu item. To cause it to be initially selected, you pass `true` as an optional second argument to the constructor.

15. [24.15](#) You add them to a `ButtonGroup`.

16. [24.16](#) The `JCheckBoxMenuItem` class. You pass a string to the constructor, which is displayed on the menu item. To cause it to be initially selected, you pass `true` as an optional second argument to the constructor.

17. [24.17](#) The `JMenu` class. You pass a string to the constructor, which is displayed on the menu.

18. [24.18](#) The `JMenuBar` class.

19. [24.19](#) With the `JFrame` object's `setJMenuBar` method.



20. [24.20](#) Action events.
21. [24.21](#) With the `setPreferredSize` method.
22. [24.22](#) The first argument to the `Dimension` class constructor is the component's width, and the second argument is the component's height.
23. [24.23](#) One constructor accepts the size of the text area, in rows and columns. Another constructor accepts the string that is to be displayed initially in the text area, as well as the size of the text area in rows and columns.
24. [24.24](#) With the `getText` method.
25. [24.25](#) `JTextArea` components do not automatically display scroll bars. To display scroll bars on a `JTextArea` component, you must add it to the scroll pane.
26. [24.26](#) Line wrapping causes text to automatically wrap to the next line when a line is filled. There are two different styles of line wrapping: word wrapping and character wrapping. When *word wrapping* is performed, the line breaks always occur between words, never in the middle of a word. When character wrapping is performed, lines are broken between characters. You use the `JTextArea` class's `setLineWrap` method to turn line wrapping on. The method accepts a `boolean` argument. If you pass `true`, then line wrapping is turned on. If you pass `false`, line wrapping is turned off. You specify the style of line wrapping that you prefer with the `JTextArea` class's `setWrapStyleWord` method. This method accepts a `boolean` argument. If you pass `true`, then the text area will perform word wrapping. If you pass `false`, the text area will perform character wrapping.
27. [24.27](#) A `Font` object.
28. [24.28](#) The first argument is the name of a font. The second argument is an `int` that represents the style of the text. The third argument is the size of the text in points.
29. [24.29](#) It generates a change event.
30. [24.30](#)
  1. `setMajorTickSpacing`
  2. `setMinorTickSpacing`
  3. `setPaintTicks`
  4. `setPaintLabels`

# Chapter 25 (Available on the Book's Companion Website)

1. [25.1](#) When a user accesses a Web page with his or her browser, any applet associated with the Web page is transmitted to the user's system. The applet is then executed on the user's system.
2. [25.2](#) To prevent malicious code from attacking or spying on unsuspecting users.
3. [25.3](#) `<html>` marks the beginning and `</html>` marks the end.
4. [25.4](#) `<head></head>`
5. [25.5](#) You would use: `<title>My Web Page</title>`  
This would go in the document head section.
6. [25.6](#) `<body></body>`
7. [25.7](#) `<h1>Student Roster</h1>`
8. [25.8](#) `<center><b>My Resume</b></center>`
9. [25.9](#) `<b><i>Hello World</i></b>`
10. [25.10](#) The `<br />` tag causes a line break. The `<p />` tag causes a paragraph break. The `<hr />` tag displays a horizontal rule.
11. [25.11](#) `<a href= "http://java.sun.com">Click Here</a>`
12. [25.12](#) `JApplet`
13. [25.13](#) `init`
14. [25.14](#) Because the browser creates an instance of the class automatically.
15. [25.15](#) `<applet code="MyApplet.class" width=400 height=200>`
16. [25.16](#) `Applet`
17. [25.17](#) You simply add them to the `Frame` or `Applet` object. These classes have an `add` method.
18. [25.18](#) You should override the `paint` method.
19. [25.19](#) You override the `paintComponent` method.

20. [25.20](#) These methods are automatically called when the component is first displayed and are called again any time the component needs to be redisplayed.
21. [25.21](#) The superclass's `paint` or `paintComponent` method should be called.
22. [25.22](#) Call the `repaint` method.
23. [25.23](#) A rectangle.
24. [25.24](#) The first array contains the  $X$  coordinates for each vertex, and the second array contains the  $Y$  coordinates for each vertex.
25. [25.25](#)
  1. `drawLine`
  2. `fillRect`
  3. `fillOval`
  4. `fillArc`
  5. `setColor`
  6. `drawRect`
  7. `drawOval`
  8. `drawArc`
  9. `drawString`
  10. `setFont`
26. [25.26](#) A mouse press event indicates that the mouse button was pressed. A mouse click event indicates that the mouse button was pressed, and then released.
27. [25.27](#)
  - To handle a mouse click event: `MouseListener`
  - To handle a mouse press event: `MouseListener`
  - To handle a mouse dragged event: `MouseMotionListener`
  - To handle a mouse release event: `MouseListener`
  - To handle a mouse move event: `MouseMotionListener`

28. [25.28](#) They accept `MouseEvent` objects. The `getX` and `getY` methods return the mouse's *X* and *Y* coordinates.
29. [25.29](#) No, they cannot be left out. Empty definitions can be written, or an adapter class can be used.
30. [25.30](#) An adapter class implements an interface and provides empty definitions for all of the required methods. When you extend a class from an adapter class, it inherits the empty methods. In your subclass, you can override the methods you want and forget about the others.
31. [25.31](#) Action events
32. [25.32](#) In milliseconds.
33. [25.33](#) By calling its `start` method.
34. [25.34](#) By calling its `stop` method.
35. [25.35](#) The `play` method.
36. [25.36](#) The `Applet` class's `play` method loads a sound file, plays it one time, and then releases it for garbage collection. If you need to load a sound file to be played multiple times, you should use an `AudioClip` object. An `AudioClip` object loads a sound file and retains it in memory. It provides methods for playing, looping, and stopping the audio play.
37. [25.37](#) The `play` method plays the sound file. The `loop` method plays the sound file repeatedly. The `stop` method stops the sound file from playing.
38. [25.38](#) The `getDocumentBase` method returns a `URL` object containing the location of the HTML file that invoked the applet. The `getCodeBase` method returns a `URL` object containing the location of the applet's `.class` file.

# APPENDIX M Answers to Odd-Numbered Review Questions\*

\*From *Starting Out with Java: From Control Structures through Data Structures*, Fourth Edition.  
Tony Gaddis. Copyright © 2019 by Pearson Education, Inc. All rights reserved.

# Chapter 1

## Multiple Choice

1. [1.](#) b
2. [3.](#) a
3. [5.](#) b
4. [7.](#) c
5. [9.](#) a
6. [11.](#) a
7. [13.](#) b

## Find the Error

1. [1.](#) The algorithm performs the math operation at the wrong time. It multiplies width by length before getting values for those variables.

## Algorithm Workbench

1. [1.](#) *Display “What is the customer’s maximum amount of credit?”*

*Input maxCredit.*

*Display “What is the amount of credit used by the customer?”*

*Input creditUsed.*

$availableCredit = maxCredit - creditUsed.$

*Display availableCredit.*

2. [3.](#) *Display “What is the account’s starting balance?”*

*Input startingBalance.*

*Display “What is the total amount of the deposits made?”*

*Input deposits.*

*Display “What is the total amount of the withdrawals made?”*

*Input withdrawals.*

*Display “What is the monthly interest rate?”*

*Input interestRate.*

$\text{balance} = \text{startingbalance} + \text{deposits} - \text{withdrawals}.$

$\text{interest} = \text{balance} * \text{interestRate}.$

$\text{balance} = \text{balance} + \text{interest}.$

*Display balance.*

## Predict the Result

1. [1.7](#)

## Short Answer

1. [1.](#) Main memory, or RAM, holds the sequences of instructions in the programs that are running and the data those programs are using. Main memory, or RAM, is usually volatile. Secondary storage is a type of memory that can hold data for long periods of time—even when there is no power to the computer.
2. [3.](#) An operating system is a set of programs that manages the computer’s hardware devices and controls their processes. Windows and UNIX are examples of operating systems. Application software refers to programs that make the computer useful to the user. These programs solve specific problems or perform general operations that satisfy the needs of the user. Word processing, spreadsheet, and database packages are all examples of application software.
3. [5.](#) Because machine language programs are streams of binary numbers, and high-level language programs are made up of words.
4. [7.](#) *Syntax errors* are mistakes that the programmer has made that violate the rules of the programming language. *Logical errors* are mistakes that cause the program to produce erroneous results.

5. [9.](#) A program that translates source code into executable code.
6. [11.](#) Machine language code is executed directly by the CPU. Byte code is executed by the JVM.
7. [13.](#) Object-oriented programming
8. [15.](#) The object's methods.
9. [17.](#) No
10. [19.](#) `javac LabAssignment.java`



# Chapter 2

## Multiple Choice and True/False

1. [1.](#) c
2. [3.](#) a
3. [5.](#) a, c, and d
4. [7.](#) c
5. [9.](#) a
6. [11.](#) b
7. [13.](#) a
8. [15.](#) a
9. [17.](#) True
10. [19.](#) True
11. [21.](#) False

## Predict the Output

1. [1.](#)

```
0
100
```

2. [3.](#)

```
I am the incrediblecomputing
machine
and I will
amaze
you.
```

3. [5.](#)

```
23
```

# Find the Error

- The comment symbols in the first line are reversed. They should be `/*` and `*/`.
- The word `class` is missing in the second line. It should read `public class MyProgram`.
- The `main` header should not be terminated with a semicolon.
- The fifth line should have a left brace, not a right brace.
- The first four lines inside the `main` method are missing their semicolons.
- The comment in the first line inside the `main` method should begin with forward slashes (`//`), not backward slashes.
- The last line inside the `main` method, a call to `println`, uses a string literal, but the literal is enclosed in single quotes. It should be enclosed in double quotes, like this: `"The value of c is"`.
- The last line inside the `main` method passes `C` to `println`, but it should pass `c` (lowercase).
- The class is missing its closing brace.

# Algorithm Workbench

1. [1.](#) `double temp, weight, age;`

2. [3.](#)

1. `a. b = a + 2;`

2. `b. a = b * 4;`

3. `c. b = a / 3.14;`

4. `d. a = b - 8;`

5. `e. c = 'K';`

6. `f. c = 66;`

3. [5.](#)

1. a. 3.287E6
2. b. -9.7865E12
3. c. 7.65491E-3

4. [7.](#)<sub>10 20 1</sub>

5. [9.](#)<sub>a</sub>

6. [11.](#)

```
int speed, time, distance;
speed = 20;
time = 10;
distance = speed * time;
System.out.println(distance);
```

7. [13.](#)

```
double income;
// Create a Scanner object for keyboard input.
Scanner keyboard = new Scanner(System.in);
// Ask the user to enter his or her desired income
System.out.print("Enter your desired annual income: ");
income = keyboard.nextDouble();
```

8. [15.](#)<sub>total = (float)number;</sub>

## Short Answer

1. [1.](#) Multi-line style
2. [3.](#) A self-documenting program is written in such a way that you get an understanding of what the program is doing just by reading its code.
3. [5.](#) The `print` and `println` methods are members of the `out` object. The `out` object is a member of the `System` class. The `System` class is part of the Java API.
4. [7.](#) You should always choose names for your variables that give an indication of what they are used for. The rather nondescript name, `x`, gives no clue as to what the variable's purpose is.
5. [9.](#) In both cases you are storing a value in a variable. An assignment statement can appear anywhere in a program. An initialization, however, is part of a variable declaration.
6. [11.](#) Programming style refers to the way a programmer uses spaces, indentations, blank lines, and punctuation characters to visually arrange a program's source code. An inconsistent programming style can create confusion for a person reading the code.

7. [13.](#) javadoc SalesAverage.java

# Chapter 3

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) a
3. [5.](#) c
4. [7.](#) a
5. [9.](#) a
6. [11.](#) a
7. [13.](#) c
8. [15.](#) c
9. [17.](#) True
10. [19.](#) True
11. [21.](#) True

## Find the Error

1. [1.](#) Each `if` clause is prematurely terminated by a semicolon.
2. [3.](#) The conditionally executed blocks of code should be enclosed in braces.
3. [5.](#) The `!` operator is only applied to the variable `x`, not the expression. The code should read:  

```
if (!(x > 20))
```
4. [7.](#) The statement should use the `||` operator instead of the `&&` operator.
5. [9.](#) The `equalsIgnoreCase` method should be used instead of the `equals` method.

## Algorithm Workbench

1. [1.](#)

```
if (y == 0)
x = 100;
```

2. [3.](#)

```
if (sales < 10000)
 commission = .10;
else if (sales <= 15000)
 commission = .15;
else
 commission = .20;
```

3. [5.](#)

```
if (amount1 > 10)
{
 if (amount2 < 100)
 {
 if (amount1 > amount2)
 {
 System.out.println(amount1);
 }
 else
 {
 System.out.println(amount2);
 }
 }
}
```

4. [7.](#)

```
if (temperature >= -50 && temperature <= 150)
 System.out.println("The number is valid.");
```

5. [9.](#)

```
if (title1.compareTo(title2) < 0)
 System.out.println(title1 + " " + title2);
else
 System.out.println(title2 + " " + title1);
```

6. [11.](#) C, A, B

7. [13.](#) `System.out.printf("%,.2f", number);`

## Short Answer

1. [1.](#) Conditionally executed code is executed only under a condition, such as an expression being true.

2. [3.](#) By indenting the conditionally executed statements, you are causing them to stand out visually. This is so you can tell at a glance what part of the program the `if` statement executes.
3. [5.](#) A flag is a `boolean` variable that signals when some condition exists in the program. When the flag variable is set to `false`, it indicates the condition does not yet exist. When the flag variable is set to `true`, it means the condition does exist.
4. [7.](#) It takes two `boolean` expressions as operands and creates a boolean expression that is true only when both subexpressions are true.
5. [9.](#) They determine whether a specific relationship exists between two values. The relationships are greater-than, less-than, equal-to, not equal-to, greater-than or equal-to, and less-than or equal-to.

# Chapter 4

## Multiple Choice and True/False

1. [1.](#) a
2. [3.](#) c
3. [5.](#) a
4. [7.](#) b
5. [9.](#) c
6. [11.](#) a
7. [13.](#) a
8. [15.](#) d
9. [17.](#) d
10. [19.](#) True
11. [21.](#) False
12. [23.](#) False
13. [25.](#) True

## Find the Error

1. [1.](#) The conditionally executed statements should be enclosed in a set of braces. Also, the `again` variable should be initialized with either 'y' or 'Y'.
2. [3.](#) The expression being tested by the `do-while` loop should be `choice == 1`. Also, the `do-while` loop must be terminated by a semicolon.

## Algorithm Workbench



1. [1.](#)

```
Scanner keyboard = new Scanner(System.in);
int product = 0, num;
while (product < 100)
{
 num = keyboard.nextInt();
 product = num * 10;
}
```

2. [3.](#) The following code simply prints the numbers, separated by spaces.

```
for (int x = 0; x <= 1000; x += 10)
 System.out.print(x + " ");
```

The following code prints the numbers separated by commas.

```
for (int x = 0; x <= 1000; x += 10)
{
 if (x < 1000)
 System.out.print(x + ", ");
 else
 System.out.print(x);
}
```

3. [5.](#)

```
double total = 0;
for (int num = 1, denom = 30; num <= 30; num++, denom--)
 total += num / denom;
```

4. [7.](#)

```
Scanner keyboard = new Scanner(System.in);
int x;
do
{
 System.out.print("Enter a number: ");
 x = keyboard.nextInt();
} while (x > 0);
```

5. [9.](#)

```
for (int count = 0; count > 50; count++)
 System.out.println("count is " + count);
```

6. [11.](#)

```
int number;
Scanner keyboard = new Scanner(System.in);
System.out.print("Enter a number in the range " +
 "of 1 through 4: ");
number = keyboard.nextInt();
while (number < 1 || number > 4)
```

```

{
 System.out.print("Invalid number. Enter a " +
 "number in the range " +
 "of 1 through 4: ");
 number = keyboard.nextInt();
}

```

7. [13.](#)

```

for (int r = 7; r > 0; r--)
{
 for (int c = 0; c < r; c++)
 {
 System.out.print("*");
 }
 System.out.println();
}

```

8. [15.](#)

```

import java.util.Random;
public class ReviewQuestion15
{
 public static void main(String[] args)
 {
 Random rand = new Random();
 System.out.println(rand.nextInt(10) + 1);
 }
}

```

9. [17.](#)

```

PrintWriter outFile = new PrintWriter("NumberList.txt");
for (int i = 1; i <= 100; i++)
 outFile.println(i);
outFile.close();

```

10. [19.](#)

```

File file = new File("NumberList.txt");
Scanner inFile = new Scanner(file);
String input;
int number, total = 0;
while (inFile.hasNext())
{
 input = inFile.nextLine();
 number = Integer.parseInt(input);
 total += number;
}
System.out.println("The total is " + total);
inFile.close();

```

# Short Answer

1. [1.](#) In postfix mode the operator is placed after the operand. In prefix mode the operator is placed before the variable operand. Postfix mode causes the increment or decrement operation to happen after the value of the variable is used in the expression. Prefix mode causes the increment or decrement to happen first.
2. [3.](#) A pretest loop tests its test expression before each iteration. A posttest loop tests its test expression after each iteration.
3. [5.](#) The `while` loop is a pretest loop and the `do-while` loop is a posttest loop.
4. [7.](#) The `do-while` loop.
5. [9.](#) An accumulator is used to keep a running total of numbers. In a loop, a value is usually added to the current value of the accumulator. If it is not properly initialized, it will not contain the correct total.
6. [11.](#) There are many possible examples. A program that asks the user to enter a business's daily sales for a number of days, and then displays the total sales is one example.
7. [13.](#) Sometimes the user has a list of input values that is very long, and doesn't know the number of items there are. When the sentinel value is entered, it signals the end of the list, and the user doesn't have to count the number of items in the list.
8. [15.](#) There are many possible examples. One example is a program that asks for the average temperature for each month, for a period of five years. The outer loop would iterate once for each year and the inner loop would iterate once for each month.
9. [17.](#) Closing a file writes any unsaved data remaining in the file buffer.
10. [19.](#) After the `println` method writes its data, it writes a newline character. The `print` method does not write the newline character.
11. [21.](#) The file does not exist.
12. [23.](#) You create an instance of the `FileWriter` class to open the file. You pass the name of the file (a string) as the constructor's first argument, and the boolean value `true` as the second argument. Then, when you create an instance of the `PrintWriter` class, you pass a reference to the `FileWriter` object as an argument to the `PrintWriter` constructor. The file will not be erased if it already exists and new data will be written to the end of the file.

# Chapter 5

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) a
3. [5.](#) b
4. [7.](#) b
5. [9.](#) c
6. [11.](#) True
7. [13.](#) False
8. [15.](#) True
9. [17.](#) False
10. [19.](#) False

## Find the Error

1. [1.](#) The header should not be terminated with a semicolon.
2. [3.](#) The method should have a `return` statement that returns a `double` value.

## Algorithm Workbench

1. [1.](#) `doSomething(25);`
2. [3.](#) The value 3 will be stored in `a`, 2 will be stored in `b`, and 1 will be stored in `c`.
3. [5.](#) `result = cube(4);`
4. [7.](#)

```
public static double timesTen(double num)
{
 return num * 10.0;
}
```

## 5. [9.](#)

```
// Assume java.util.Scanner has been imported.
public static String getName()
{
 String name;
 Scanner keyboard = new Scanner(System.in);
 System.out.print("Enter your first name: ");
 name = keyboard.nextLine();
 return name;
}
```

# Short Answer

1. [1.](#) A large complex problem is broken down into smaller manageable pieces. Each smaller piece of the problem is then solved.
2. [3.](#) An argument is a value that is passed into a method when the method is called. A parameter variable is a variable that is declared in the method header, and receives the value of an argument when the method is called.
3. [5.](#) When an argument is passed to a method, only a copy of the argument is passed. The method cannot access the actual argument.

# Chapter 6

## Multiple Choice and True/False

1. [1.](#) a
2. [3.](#) d
3. [5.](#) b
4. [7.](#) d
5. [9.](#) b
6. [11.](#) c
7. [13.](#) b
8. [15.](#) True
9. [17.](#) False
10. [19.](#) False

## Find the Error

1. [1.](#) The constructor cannot have a return type, not even `void`.
2. [3.](#) The parentheses are missing. The statement should read:  

```
Rectangle box = new Rectangle();
```
3. [5.](#) The `square` methods must have different parameter lists. Both accept an `int`.

## Algorithm Workbench

1. [1.](#)
  1. a. UML diagram:

| Pet                                                                                                                                                              |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| - name : String<br>- animal : String<br>- age : int                                                                                                              |
| + setName(n : String) : void<br>+ setAnimal(a : String) : void<br>+ setAge(a : int) : void<br>+ getName() : String<br>+ getAnimal() : String<br>+ getAge() : int |

### [M.6-1 Full Alternative Text](#)

#### 2. b. Class code:

```

public class Pet
{
 private String name; // The pet's name
 private String animal; // The type of animal
 private int age; // The pet's age
 /**
 * setName method
 * @param n The pet's name.
 */
 public void setName(String n)
 {
 name = n;
 }
 /**
 * setAnimal method
 * @param a The type of animal.
 */
 public void setAnimal(String a)
 {
 animal = a;
 }
 /**
 * setAge method
 * @param a The pet's age.
 */
 public void setAge(int a)
 {
 age = a;
 }
 /**
 * getName method
 * @return The pet's name.
 */
 public String getName()

```

```

 {
 return name;
 }
 /**
 * getAnimal method
 * @return The type of animal.
 */
 public String getAnimal()
 {
 return animal;
 }
 /**
 * getAge method
 * @return The pet's age.
 */
 public int getAge()
 {
 return age;
 }
}

```

2. [3.](#)

1. a.

```

public Square()
{
 sideLength = 0.0;
}

```

2. b.

```

public Square(double s)
{
 sideLength = s;
}

```

## Short Answer

1. [1.](#) A class is a collection of programming statements that specify the attributes and methods that a particular type of object may have. You should think of a class as a “blueprint” that describes an object. An instance of a class is an actual object that exists in memory.
2. [3.](#) An accessor method is a method that gets a value from a class’s field but does not change it. A mutator method is a method that stores a value in a field or in some other way changes the value of a field.
3. [5.](#) Methods that are members of the same class.
4. [7.](#) It looks in the current folder or directory for the file `Customer.class`. If that file does not exist, the compiler searches for the file `Customer.java` and compiles it. This creates the file



`Customer.class`, which makes the `Customer` class available. The same procedure is followed when the compiler searches for the `Account` class.

5. [9.](#) If you do not write a constructor for a class, Java automatically provides one.
6. [11.](#) By their signatures, which include the method name and the data types of the method parameters, in the order that they appear.

# Chapter 7

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) b
3. [5.](#) c
4. [7.](#) b
5. [9.](#) d
6. [11.](#) c
7. [13.](#) a
8. [15.](#) False
9. [17.](#) True
10. [19.](#) True
11. [21.](#) False
12. [23.](#) True

## Find the Error

1. [1.](#) The size declarator cannot be negative.
2. [3.](#) The loop uses the values 1 through 10 as subscripts. It should use 0 through 9.
3. [5.](#) A subscript should be used with words, such as `words[0].toUpperCase()`.

## Algorithm Workbench

1. [1.](#)

```
for (int i = 0; i < 20; i++)
 System.out.println(names[i]);
```

## 2. [3.](#)

### 1. a.

```
String[] scientists = {"Einstein", "Newton",
 "Copernicus", "Kepler"};
```

### 2. b.

```
for (int i = 0; i < scientists.length; i++)
 System.out.println(scientists[i]);
```

### 3. c.

```
int total = 0;
for (int i = 0; i < scientists.length; i++)
 total += scientists[i].length();
System.out.println("The total length is " + total);
```

## 3. [5.](#)

### 1. a.

```
// Define the arrays.
int[] id = new int[10];
double[] weeklyPay = new double[10];
```

### 2. b.

```
// Display each employee's gross weekly pay.
for (int i = 0; i < 10; i++)
{
 System.out.println("The pay for employee " +
 id[i] + " is $" + weeklyPay[i]);
}
```

## 4. [7.](#)

```
final int NUM_ROWS = 30;
final int NUM_COLS = 10;
int total = 0;
double average;
for (int row = 0; row < grades.length; row++)
{
 for (int col = 0; col < grades[row].length; col++)
 {
 total += grades[row][col];
 }
}
average = (double) total / (NUM_ROWS * NUM_COLS);
```

5. [9.](#)

```
double total = 0.0; // Accumulator
// Sum the values in the array.
for (int row = 0; row < 10; row++)
{
 for (int col = 0; col < 20; col++)
 total += values[row][col];
}
```

6. [11.](#)

```
// Create an ArrayList.
ArrayList<String> cars = new ArrayList<String>();
// Add three car names to the ArrayList.
cars.add("Porsche");
cars.add("BMW");
cars.add("Jaguar");
// Display the contents of cars.
for (String str : cars)
 System.out.println(str);
```

## Short Answer

1. [1.](#) The size declarator is used in a definition of an array to indicate the number of elements the array will have. A subscript is used to access a specific element in an array.
2. [3.](#)
  1. a. 2
  2. b. 14
  3. c. 8
3. [5.](#) Because this statement merely makes `array1` reference the same array that `array2` references. Both variables will reference the same array. To copy the contents of `array2` to `array1`, the contents of `array2`'s individual elements will have to be assigned to the elements of `array1`.
4. [7.](#) It will have to read all 10,000 elements to find the value stored in the last element.

# Chapter 8

## Multiple Choice and True/False

1. [1.](#) c
2. [3.](#) a
3. [5.](#) b
4. [7.](#) d
5. [9.](#) c
6. [11.](#) b
7. [13.](#) False
8. [15.](#) False
9. [17.](#) True

## Find the Error

1. [1.](#) The static method `setValues` cannot refer to the non-static fields `x` and `y`.

## Algorithm Workbench

1. [1.](#)

1. a.

```
public String toString()
{
 String str;
 str = "Radius: " + radius + " Area: " + getArea();
 return str;
}
```

2. b.

```
public boolean equals(Circle c)
{
 boolean status;
 if (c.getRadius() == radius)
 status = true;
 else
 status = false;
 return status;
}
```

3. c.

```
public boolean greaterThan(Circle c)
{
 boolean status;
 if (c.getArea() > getArea())
 status = true;
 else
 status = false;
 return status;
}
```

2. [3.](#) enum Pet { DOG, CAT, BIRD, HAMSTER }

## Short Answer

1. [1.](#) Access a non-static member.
2. [3.](#) When a variable is passed as an argument, a copy of the variable's contents is passed. The receiving method does not have access to the variable itself. When an object is passed as an argument, a reference to the object (which is the object's address) is passed. This allows the receiving method to have access to the object.
3. [5.](#) It means that an aggregate relationship exists. When an object of class B is a member of class A, it can be said that class A "has a" class B object.
4. [7.](#) It is not advisable because it will allow access to the private fields. The exception to this is when the field is a `String` object. This is because `String` objects are immutable, meaning that they cannot be changed.
5. [9.](#)

1. a) `Color`

2. b) `Color.RED, Color.ORANGE, Color.GREEN, Color.BLUE`

3. c) `Color myColor = Color.BLUE;`

6. [11.](#) When there are no references to it.



# Chapter 9

## Multiple Choice and True/False

1. [1.](#) c
2. [3.](#) a
3. [5.](#) a
4. [7.](#) b
5. [9.](#) d
6. [11.](#) a
7. [13.](#) d
8. [15.](#) False
9. [17.](#) False
10. [19.](#) True
11. [21.](#) True
12. [23.](#) False

## Find the Error

1. [1.](#) The `valueOf` method is static. It should be called like this:

```
str = String.valueOf(number);
```

2. [3.](#) The very first character is at position 0, so the statement should read:

```
str.setCharAt(0, 'Z');
```

## Algorithm Workbench

1. [1.](#) `if (Character.toUpperCase(choice) == 'Y')`



or

```
if (Character.toLowerCase(choice) == 'y')
```

2. [3.](#)

```
int total = 0;
for (int i = 0; i < str.length(); i++)
{
 if (Character.isDigit(str.charAt(i)))
 total++;
}
```

3. [5.](#)

```
public static boolean dotCom(String str)
{
 boolean status;
 if (str.endsWith(".com"))
 status = true;
 else
 status = false;
 return status;
}
```

4. [7.](#)

```
public static void upperT(StringBuilder str)
{
 for (int i = 0; i < str.length(); i++)
 {
 if (str.charAt(i) == 't')
 str.setCharAt(i, 'T');
 }
}
```

5. [9.](#)

```
if (d <= Integer.MAX_VALUE)
 i = (int) d;
```

## Short Answer

1. [1.](#) This will improve the program's efficiency by reducing the number of `String` objects that must be created and then removed by the garbage collector.
2. [3.](#) Converts a number to a string.

# Chapter 10

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) d
3. [5.](#) a
4. [7.](#) c
5. [9.](#) a
6. [11.](#) a
7. [13.](#) c
8. [15.](#) c
9. [17.](#) c
10. [19.](#) d
11. [21.](#) c
12. [23.](#) True
13. [25.](#) False
14. [27.](#) False
15. [29.](#) True
16. [31.](#) True
17. [33.](#) True

## Find the Error

1. [1.](#) The `Car` class header should use the word `extends` instead of `expands`.

2. [3.](#) Because the `Vehicle` class does not have a default constructor or a no-arg constructor, the `Car` class constructor must call the `Vehicle` class constructor.

# Algorithm Workbench

1. [1.](#) `public class Poodle extends Dog`
2. [3.](#)

```
public abstract class B
{
 private int m;
 protected int n;

 public void setM(int value)
 {
 m = value;
 }
 public void setN(int value)
 {
 n = value;
 }
 public int getM()
 {
 return m;
 }
 public int getN()
 {
 return n;
 }
 public abstract double calc();
}

public class D extends B
{
 private double q;
 protected double r;
 public void setQ(double value)
 {
 q = value;
 }
 public void setR(double value)
 {
 r = value;
 }
 public double getQ()
 {
 return q;
 }
 public double getR()
 {
 return r;
 }
 public double calc()
```

```
{
 return q * r;
}
```

3. [5.](#) `setValue(10);`

or

```
super.setValue(10);
```

4. [7.](#)

```
public class Stereo extends SoundSystem
 implements CDPlayable,
 TunerPlayable,
 CassettePlayable
```

5. [9.](#) `Computable half = x -> x / 2;`

## Short Answer

1. [1.](#) When an “is a” relationship exists between objects, it means that the specialized object has all of the characteristics of the general object, plus additional characteristics that make it special.
2. [3.](#) `Dog` is the superclass and `Pet` is the subclass.
3. [5.](#) No.
4. [7.](#) Overloading is when a method has the same name as one or more other methods, but a different parameter list. Although overloaded methods have the same name, they have different signatures. When a method overrides another method, however, they both have the same signature.
5. [9.](#) At runtime.
6. [11.](#) An abstract class is not instantiated itself, but serves as a superclass for other classes. The abstract class represents the generic or abstract form of all the classes that inherit from it.
7. [13.](#) The class must implement an interface, or extend a superclass.
8. [15.](#) An expression that creates an object that implements a functional interface.

# Chapter 11

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) a
3. [5.](#) b
4. [7.](#) d
5. [9.](#) c
6. [11.](#) c
7. [13.](#) d
8. [15.](#) c
9. [17.](#) True
10. [19.](#) False
11. [21.](#) True
12. [23.](#) False

## Find the Error

1. [1.](#) The `try` block must appear first.
2. [3.](#) The `catch (Exception e)` statement and its block should appear after the other catch blocks, because this is a more general exception than the others.

## Algorithm Workbench

1. [1.](#)

B  
D

## 2. [3.](#)

```
public static int arraySearch(int[] array, int value)
 throws Exception
{
 int i; // Loop control variable
 int element; // Element the value is found at
 boolean found; // Flag indicating search results

 // Element 0 is the starting point of the search.
 i = 0;

 // Store the default values element and found.
 element = -1;
 found = false;

 // Search the array.
 while (!found && i < array.length)
 {
 if (array[i] == value)
 {
 found = true;
 element = i;
 }
 i++;
 }
 if (element == -1)
 throw new Exception("Element not found.");
 else
 return element;
}
```

## 3. [5.](#)

```
public class NegativeNumber extends Exception
{
 /**
 * No-arg constructor
 */

 public NegativeNumber()
 {
 super("Error: Negative number");
 }

 /**
 * The following constructor accepts the number that
 * caused the exception.
 * @param n The number.
 */
 public NegativeNumber(int n)
 {
 super("Error: Negative number: " + n);
 }
}
```

4. [7.](#) `public int getValueFromFile() throws IOException, FileNotFoundException`

5. [9.](#)

```
FileOutputStream fstream =
 new FileOutputStream("Configuration.dat");
```

6. [11.](#)

```
FileOutputStream outputStream =
 new FileOutputStream("ObjectData.dat");
ObjectOutputStream objectOutputStream =
 new ObjectOutputStream(outputStream);
objectOutputStream.writeObject(r);
```

## Short Answer

1. [1.](#) An exception object has been created in response to an error that has occurred.
2. [3.](#) Control of the program is passed to the previous method in the call stack (that is, the method that called the offending method). If that method cannot handle the exception, then control is passed again, up the call stack, to the previous method. This continues until control reaches the `main` method. If the `main` method does not handle the exception, then the program is halted and the default exception handler handles the exception.
3. [5.](#) The first statement after that `try/catch` construct.
4. [7.](#) Any object that inherits from the `Throwable` class.
5. [9.](#) Unchecked exceptions are those that inherit from the `Error` class or the `RuntimeException` class. You should not handle these exceptions because the conditions that cause them can rarely be dealt with within the program. All of the remaining exceptions (that is, those that do not inherit from `Error` or `RuntimeException`) are *checked* exceptions. These are the exceptions that you should handle in your program.
6. [11.](#) All of the data stored in a text file is formatted as text. Even numeric data is converted to text when it is stored in a text file. You can view the data stored in a text file by opening it with a text editor, such as Notepad. In a binary file, data is not formatted as text. Subsequently, you cannot view a binary file's contents with a text editor.
7. [13.](#) When an object is serialized it is converted to a series of bytes, which are written to a file. When the object is deserialized, the series of bytes are converted back into an object.

# Chapter 12

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) a
3. [5.](#) c
4. [7.](#) a
5. [9.](#) b
6. [11.](#) a
7. [13.](#) b
8. [15.](#) False
9. [17.](#) True

## Find the Error

1. [1.](#) The parameter for the overridden `start` method must be a `Stage` object, not a `Scene` object.
2. [3.](#) The argument that is passed to the `setAlignment` method should be `Pos.CENTER`.
3. [5.](#) The correct code is:

```
myImageView.setFitWidth(100);
myImageView.setFitHeight(100);
```

## Algorithm Workbench

1. [1.](#) `HBox hbox = new HBox(label1, label2, label3);`
2. [3.](#) `primaryStage.setScene(scene);`
3. [5.](#)



```
Image image = new Image("file:Cat.png");
ImageView imageView = new ImageView(image);
```

4. [7.](#) `HBox hbox = new HBox(10, label1, label2, label3);`

5. [9.](#) `gridPane.add(button, 2, 5);`

6. [11.](#)

```
class ButtonClickHandler implements EventHandler<ActionEvent>
{
 @Override
 void handle(ActionEvent event)
 {
 outputLabel.setText("Hello World");
 }
}
```

7. [13.](#)

```
myButton.setOnAction(event ->
{
 outputLabel.setText("Hello World");
});
```

## Short Answer

1. [1.](#) A program that responds to the user's actions.
2. [3.](#) What is the purpose of the `Application` class's abstract `start` method? It is the entry point for a JavaFX application.
3. [5.](#) `javafx.scene.control`
4. [7.](#) The aspect ratio is the ratio of the image's width to the image's height. If you change an image's aspect ratio, the image will appear stretched (either horizontally or vertically). The `ImageView` class has a method named `setPreserveRatio` that you can call to make sure that an image's aspect ratio is preserved.
5. [9.](#) You call the control's `getText` method. The method returns the value that has been entered into the `TextField` as a `String`.
6. [11.](#) `BorderPane`

# Chapter 13

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) d
3. [5.](#) a
4. [7.](#) a
5. [9.](#) c
6. [11.](#) c
7. [13.](#) a
8. [15.](#) d
9. [17.](#) True
10. [19.](#) True
11. [21.](#) False
12. [23.](#) True

## Find the Error

1. [1.](#) The=operator should be a colon.
2. [3.](#) The last two statements should be:

```
radio1.setToggleGroup(radioGroup);
radio2.setToggleGroup(radioGroup);
```

## Algorithm Workbench

1. [1.](#)

```
.label {
 -fx-font-size: 24pt;
 -fx-font-style: italic;
 -fx-font-weight: bold;
}
```

2. [3.](#)

```
.label, .button, .text-field {
 -fx-font-size: 24pt;
 -fx-font-style: italic;
 -fx-font-weight: bold;
}
```

3. [5.](#) `rgb(100, 20, 255);`

4. [7.](#) `calcButton.setId("calc-button");`

5. [9.](#)

```
radio1.setToggleGroup(radioGroup);
radio2.setToggleGroup(radioGroup);
radio3.setToggleGroup(radioGroup);
```

6. [11.](#)

```
ListView<String> snackListView = new ListView<>();
snackListView.getItems().addAll("cheese", "olives", "crackers");
```

7. [13.](#)

```
Slider slider = new Slider(0.0, 1000.0, 500.0);
slider.setOrientation(Orientation.HORIZONTAL);
slider.setShowTickMarks(true);
slider.setMajorTickUnit(100);
slider.setMinorTickCount(25);
slider.setShowTickLabels(true);
```

## Short Answer

1. [1.](#) A type selector is used to apply a style definition to all nodes of a specific type. An ID selector can be used to apply a style definition to a node with a specific ID.
2. [3.](#) `RadioButtons`
3. [5.](#) You group `RadioButtons` together with a `ToggleGroup` so only one of the `RadioButtons` can be selected at a time.
4. [7.](#) An editable `ComboBox` allows the user to select an item from a drop-down list, or type input into a field that is similar to a `TextField`. By default, `ComboBox` controls are uneditable, which

means that the user cannot type input into the control; he or she can only select items from a dropdown list.

## 5. [9.](#) ComboBox

# Chapter 14

## Multiple Choice and True/False

1. [1.](#) d
2. [3.](#) a
3. [5.](#) c
4. [7.](#) a
5. [9.](#) d
6. [11.](#) True
7. [13.](#) True
8. [15.](#) False

## Algorithm Workbench

1. [1.](#)

```
Line myLine = new Line(50, 25, 150, 125);
```
2. [3.](#)

```
Rectangle myRectangle = new Rectangle(100, 200, 175, 150);
```
3. [5.](#)  

```
Polygon poly = new Polygon(100.0, 100.0,
 200.0, 100.0,
 200.0, 200.0,
 100.0, 200.0);
```
4. [7.](#)

```
myBox.setRotate(45.0);
```
5. [9.](#)  

```
Circle myCircle = new Circle(0, 100, 30);
TranslateTransition ttrans =
 new TranslateTransition(new Duration(5000), myCircle);
ttrans.setFromX(0);
ttrans.setFromY(100);
ttrans.setToX(200);
ttrans.setToY(150);
```

```
ttrans.play();
```

# Short Answer

1. [1.](#)(639, 0)
2. [3.](#)setFromX and setFromY
3. [5.](#)setFromAngle
4. [7.](#)setByAngle
5. [9.](#)setToX and setToY
6. [11.](#)BoxBlur, GuassianBlur, and MotionBlur

# Chapter 15

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) d
3. [5.](#) d
4. [7.](#) True
5. [9.](#) False

## Find the Error

1. [1.](#) The recursive method, `myMethod`, has no base case. So, it has no way of stopping.

## Algorithm Workbench

1. [1.](#)

```
public static void main(String [] args)
{
 String str = "test string";
 display(str, 0);
}
public static void display(String str, int pos)
{
 if (pos < str.length())
 {
 System.out.print(str.charAt(pos));
 display(str, pos + 1);
 }
}
```

2. [3.](#) 10
3. [5.](#) 55
4. [7.](#)

```
public static int factorial(int num)
```

```
{
 int fac = 1;
 for (int i = 1; i <= num; i++)
 {
 fac = fac * i;
 }
 return fac;
}
```

## Short Answer

1. [1.](#) An iterative algorithm uses a loop to solve the problem, while a recursive algorithm uses a method that calls itself.
2. [3.](#) For Question 3 the base case is reached when `arg` is equal to 10. For Question 4 the base case is also reached when `arg` is equal to 10. For Question 5 the base case is reached when `num` is less than or equal to 0.
3. [5.](#) Recursive algorithms are usually less efficient than iterative algorithms. This is because a method call requires several actions to be performed by the JVM, such as allocating memory for parameters and local variables, and storing the address of the program location where control returns after the method terminates.
4. [7.](#) The value of an argument is usually reduced.



# Chapter 16

## Multiple Choice and True/False

- 1. [1.](#) a
- 2. [3.](#) a
- 3. [5.](#) c
- 4. [7.](#) a
- 5. [9.](#) b
- 6. [11.](#) True
- 7. [13.](#) True

## Find the Error

- 1. [1.](#) The numbers in the array are not sorted.

## Algorithm Workbench

- 1. [1.](#)

| Array Size                                 | 50       | 500      | 10,000   | 100,000  | 10,000,000 |
|--------------------------------------------|----------|----------|----------|----------|------------|
|                                            | Elements | Elements | Elements | Elements | Elements   |
| Sequential Search<br>(Average Comparisons) | 25       | 250      | 5,000    | 50,000   | 5,000,000  |
| Sequential Search<br>(Maximum Comparisons) | 50       | 500      | 10,000   | 100,000  | 10,000,000 |

Binary Search

|                       |   |   |    |    |    |
|-----------------------|---|---|----|----|----|
|                       | 6 | 9 | 14 | 17 | 24 |
| (Maximum Comparisons) |   |   |    |    |    |

## Short Answer

1. [1.](#) It will have to read all 10,000 elements of the array.
2. [3.](#) One time.
3. [5.](#) Because it moves the items in the array only by one element at a time.
4. [7.](#)
  - Move 1 to element 0. The array now looks like this:
    - 1, 4, 3, 2
  - Move 2 to element 1. The array now looks like this:
    - 1, 2, 3, 4
  - The values are now in order. No further swaps are made.

# Chapter 17

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) a
3. [5.](#) a, b, and c
4. [7.](#) a
5. [9.](#) d
6. [11.](#) False
7. [13.](#) False
8. [15.](#) True
9. [17.](#) False

## Find the Error

1. [1.](#) The type argument must also be provided in the variable declaration. The statement should read:  

```
ArrayList<String> myList = new ArrayList<String>();
```
2. [3.](#) The static method cannot refer to the class's type parameter `T`. It must define its own type parameter.

## Algorithm Workbench

1. [1.](#)

```
ArrayList<Customer> customers = new ArrayList<Customer>();
```
2. [3.](#)

```
public class MyType<T extends String>
```
3. [5.](#)

```
public static <T extends Comparable> T max(T a, T b)
```

```
{
 if (a.compareTo(b) > 0)
 return a;
 else
 return b;
}
```

4. [7.](#) `public class MyType<T extends Number, S extends String>`

## Short Answer

1. [1.](#) At compile time.
2. [3.](#) When the Java compiler encounters generic code, it uses a process known as erasure to compile the code. The process is called erasure because the compiler erases the generic notation and substitutes an actual type for each type parameter.
3. [5.](#) `Object`.

# Chapter 18

## Multiple Choice and True/False

1. [1.](#) a
2. [3.](#) b
3. [5.](#) b
4. [7.](#) a
5. [9.](#) c
6. [11.](#) a
7. [13.](#) d
8. [15.](#) True
9. [17.](#) False
10. [19.](#) True
11. [21.](#) False

## Find the Error

1. [1.](#) The `LinkedHashSet` class does not implement the `List` interface. Replace with `LinkedList` or `ArrayList`.
2. [3.](#) You cannot use a `ListIterator` from an object that implements the `Set` interface.

## Algorithm Workbench

1. [1.](#)

```
List<String> sList = new ArrayList<>();
```
2. [3.](#)  

```
while (it.hasNext())
{
```

```
System.out.println(it.next());
}
```

# Short Answer

1. [1.](#) A source
2. [3.](#) You can change the class that implements the interface without changing the rest of the code.
3. [5.](#) By calling the `iterator()` method. Or if you want to, you can call the `listIterator()` method.
4. [7.](#) By internally generating code that uses an iterator.
5. [9.](#) Because `IntStream` allows streams of the primitive type `int`. You cannot get a stream of primitive types from `Stream<Integer>`.
6. [11.](#) Yes.
7. [13.](#) A map associates keys with values. A `HashMap` object will use the hash codes of the keys to find the keys and then from there, to find the values.

# Chapter 19

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) c
3. [5.](#) a
4. [7.](#) d
5. [9.](#) c
6. [11.](#) True
7. [13.](#) False
8. [15.](#) True
9. [17.](#) False
10. [19.](#) False

## Find the Error

1. [1.](#) Replace `ref++` with `ref = ref.next;`
2. [3.](#) You need to set the `prev` reference of the second node. Add the code:

```
if (myList.next != null)
 myList.next.prev = myList;
```

3. [5.](#) You need to check to see if the list is now empty, and you need to set the `next` reference of the last node to null.

```
if (last == null)
 first = null;
else
 last.next = null;
```

## Algorithm Workbench

1. [1.](#)

```
static void print(Node ref)
{
 if (ref != null)
 {
 // Print the first element
 System.out.print(ref.value + " ");
 // Print elements in the tail
 print(ref.next);
 }
}
```

2. [3.](#)

```
Node reverse(Node list)
{
 // Check base case
 if (list == null || list.next == null)
 return null;
 // reverse the tail
 Node revTail = reverse(list.next);
 // Put head at the end of the reversed tail
 list.next = null;
 Node last = revTail;
 while (last.next != null)
 last = last.next;
 last.next = list;

 // revTail is now head of the reversed list
 return revTail;
}
```

## Short Answer

1. [1.](#) They need to be passed a `Node` reference as parameter. `Node` is an implementation detail that should not be made public.
2. [3.](#) It is the zero-based position of the element in the list.
3. [5.](#) Successive elements of the list are not necessarily stored in consecutive memory locations. Instead, the list stores with each element a reference to its successor.



# Chapter 20

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) b
3. [5.](#) c
4. [7.](#) False
5. [9.](#) False

## Find the Error

1. [1.](#) The flow of execution never reaches the statement `top --;`
2. [3.](#) `pop` forgets to remove the value from the stack.
3. [5.](#) `dequeue` should remove from the front, not the rear. Also, the index should wrap to the beginning of the array when it gets to the end.

## Algorithm Workbench

1. [1.](#)

```
static boolean palindrome(String s)
{
 Stack<Character> stack = new Stack<Character>();
 // Push a copy of the string onto the stack
 for(int k = 0; k < s.length(); k++)
 stack.push(s.charAt(k));
 // Compare the two copies using the stack
 for (int k = 0; k < s.length(); k++)
 {
 if (s.charAt(k) != stack.pop())
 return false;
 }
 // No discrepancy found, so s is a palindrome
 return true;
}
```

2. [3.](#) To add  $x$  to the queue, push  $x$  onto the first stack. To remove an item from the queue, pop everything from the first stack and store it into the second stack, then pop the item at the top of the second stack and save it to be returned. Pop everything from the second stack back onto the first stack. Return the saved item.

## Short Answer

1. [1.](#) A stack
2. [3.](#) A stack
3. [5.](#) To facilitate garbage collection.

# Chapter 21

## Multiple Choice and True/False

1. [1.](#) d
2. [3.](#) c
3. [5.](#) c
4. [7.](#) a
5. [9.](#) b
6. [11.](#) False
7. [13.](#) True
8. [15.](#) True
9. [17.](#) True
10. [19.](#) True
11. [21.](#) False

## Find the Error

1. [1.](#) The function does not use the values returned from the recursive calls. Should be:

```
if (tree!= null)
 return NodeCount(tree.left)+NodeCount(tree.right)+1;
else
 return 0;
```

2. [3.](#) The code will fail if `tree` is null.

## Short Answer

1. [1.](#) It keeps the height of the tree shallow, making algorithms that traverse a path of the tree from the root to a leaf more efficient.

2. [3.](#) The length of the longest path from the root to a leaf.
3. [5.](#) A collection that stores elements from an ordered set, and from which elements are removed according to that order, with the least element in the order being removed first.
4. [7.](#) The AVL tree must satisfy a balance condition: the difference in the heights of the two subtrees of any node must be at most 1.

# Algorithm Workbench

1. [1.](#)

```
Node preorder(Node bTree)
{
 if (tree != null)
 {
 System.out.print(bTree.value + " ");
 preorder(bTree.left);
 preorder(bTree.right);
 }
}
```

2. [3.](#)

```
int treeSize(Node bTree)
{
 if (bTree == null)
 return 0;
 else
 return treeSize(bTree.left) + treeSize(bTree.right) + 1;
}
```

3. [5.](#)

```
boolean contains(Node btree, int x)
{
 if (btree == null)
 return false;
 if (btree.value == x)
 return true;
 if (contains(btree.left, x))
 return true;
 if (contains(btree.right, x))
 return true;
 else
 return false;
}
```

4. [7.](#)

```
static class DecOrder implements Comparator<String>
```

```
{
 public int compare(String x, String y)
 {
 if (x.length() == y.length())
 return x.compareTo(y);
 else
 return x.length()-y.length();
 }
}
```

# Chapter 22 (Available on the Book's Companion Website)

## Multiple Choice and True/False

1. [1.](#) b
2. [3.](#) a
3. [5.](#) d
4. [7.](#) b
5. [9.](#) a
6. [11.](#) c
7. [13.](#) b
8. [15.](#) d
9. [17.](#) c
10. [19.](#) True
11. [21.](#) False
12. [23.](#) False

## Find the Error

1. [1.](#) The string “French Roast Dark” should be enclosed in single quotes instead of double quotes.
2. [3.](#) The last statement should call the `executeQuery` method instead of the `execute` method.

## Algorithm Workbench

1. [1.](#)

- INTEGER **or** INT
- REAL
- CHARACTER(*n*) **or** CHAR(*n*) **or** VARCHAR(*n*)
- DOUBLE

2. [3.](#) SELECT \* FROM Stock

3. [5.](#) SELECT TradingSymbol, NumShares FROM Stock

4. [7.](#)

- SELECT \* FROM Stock
- WHERE TradingSymbol LIKE 'SU%'

5. [9.](#)

- SELECT TradingSymbol, NumShares FROM Stock
- WHERE SellingPrice > PurchasePrice AND
- NumShares > 100
- ORDER BY NumShares

6. [11.](#)

- UPDATE Stock
- SET TradingSymbol = 'ABC'
- WHERE TradingSymbol = 'XYZ'

7. [13.](#) Connection conn = DriverManager.getConnection(DB\_URL);

8. [15.](#) ResultSet result = stmt.executeQuery(sqlStatement);

9. [17.](#)

- CREATE TABLE Car
- ( Manufacturer CHAR(25),
- Year INTEGER,

# Short Answer

1. [1.](#) Traditional text and binary files are not practical when a large amount of data must be stored and manipulated. Many businesses keep hundreds of thousands, or even millions, of data items in files. When a text or binary file contains this much data, simple operations such as searching, inserting, and deleting become cumbersome and inefficient.
2. [3.](#) The data that is stored in a database is organized into one or more tables. Each table holds a collection of related data. The data that is stored in a table is then organized into rows and columns. A row is a complete set of information about a single item. The data that is stored in a row is divided into columns. Each column is an individual piece of information about the item.
3. [5.](#) A result set is an object that is somewhat similar to a collection, and contains the results of an SQL statement.
4. [7.](#) The first row is row 1. The first column is column 1.
5. [9.](#) A foreign key is a column in one table that references a primary key in another table.



# Chapter 23 (Available on the Book's Companion Website)

## Multiple Choice and True/False

1. [1.](#) c
2. [3.](#) c
3. [5.](#) d
4. [7.](#) b
5. [9.](#) b
6. [11.](#) a
7. [13.](#) False
8. [15.](#) True
9. [17.](#) False
10. [19.](#) False

## Find the Error

1. [1.](#) The `x` is missing in the package name. The statement should read:

```
import javax.swing.*;
```

2. [3.](#) The arguments passed to the `GridLayout` constructor are reversed. This statement creates 10 rows and 5 columns.
3. [5.](#) You do not create an instance of `BorderFactory`. Instead you call one of its static methods to create a `Border` object. The statement should read:

```
panel.setBorder(BorderFactory.createTitledBorder("Choices"));
```

## Algorithm Workbench

1. [1.](#) `myWindow.setSize(500, 250);`
2. [3.](#) `myWindow.setVisible(true);`
3. [5.](#) `setLayout(new FlowLayout(FlowLayout.LEFT));`
4. [7.](#) `panel.add(button, BorderLayout.WEST);`
5. [9.](#) `panel.setBorder(BorderFactory.createLineBorder(Color.blue, 2));`

## Short Answer

1. [1.](#) The window is hidden from view, but the application does not end.
2. [3.](#) Radio buttons are normally used to select one of several possible items. Because a mutually exclusive relationship usually exists between radio buttons, only one of the items may be selected. Check boxes, which may appear alone or in groups, allow the user to make yes/no or on/off selections. Because there is not usually a mutually exclusive relationship between check boxes, the user can select any number of them when they appear in a group.

# Chapter 24 (Available on the Book's Companion Website)

## Multiple Choice and True/False

1. [1.](#) d
2. [3.](#) b
3. [5.](#) a
4. [7.](#) c
5. [9.](#) b
6. [11.](#) c
7. [13.](#) a
8. [15.](#) b
9. [17.](#) a
10. [19.](#) b
11. [21.](#) a
12. [23.](#) False
13. [25.](#) False
14. [27.](#) True
15. [29.](#) True
16. [31.](#) True
17. [33.](#) False
18. [35.](#) True

# Find the Error

1. [1.](#) The argument `false` should have been passed to the `setEditable` method.
2. [3.](#) You should pass `list` as an argument to the `JScrollPane` constructor:

```
JScrollPane scrollPane = new JScrollPane(list);
```

3. [5.](#) The second statement should read:

```
label.setIcon(image);
```

4. [7.](#) The statement should read:

```
JTextArea textArea = new JTextArea (5, 20);
```

# Algorithm Workbench

1. [1.](#) `JTextField textField = new JTextField(20);`

```
textField.setEditable(false);
```

2. [3.](#) `dayList.setVisibleRowCount(4);`

```
JScrollPane scrollPane = new JScrollPane(dayList);
```

3. [5.](#) `selectionIndex = myComboBox.getSelectedIndex();`

4. [7.](#) `ImageIcon image = new ImageIcon("picture.gif");`

```
label.setIcon(image);
```

5. [9.](#) `JFileChooser fileChooser = new JFileChooser();`

```
int status = fileChooser.showOpenDialog(null);
if (status == JFileChooser.APPROVE_OPTION)
{
 File selectedFile = fileChooser.getSelectedFile();
 String filename = selectedFile.getPath();
}
```

6. [11.](#)

```
// Create an Open menu item.
JMenuItem openItem = new JMenuItem("Open");
openItem.setMnemonic(KeyEvent.VK_O);
openItem.addActionListener(new OpenListener());
// Create a Print menu item.
```

```
JMenuItem printItem = new JMenuItem("Print");
printItem.setMnemonic(KeyEvent.VK_P);
printItem.addActionListener(new PrintListener());
// Create an Exit menu item.
JMenuItem exitItem = new JMenuItem("Exit");
exitItem.setMnemonic(KeyEvent.VK_X);
exitItem.addActionListener(new ExitListener());
// Create a JMenu object for the File menu.
JMenu fileMenu = new JMenu("File");
fileMenu.setMnemonic(KeyEvent.VK_F);
// Add the menu items to the File menu.
fileMenu.add(openItem);
fileMenu.add(printItem);
fileMenu.add(exitItem);
// Create the menu bar.
JMenuBar menuBar = new JMenuBar();
// Add the file menu to the menu bar.
menuBar.add(fileMenu);
// Set the window's menu bar.
setJMenuBar(menuBar);
```

## Short Answer

1. [1.](#) Single selection mode
2. [3.](#) An uneditable combo box combines a button with a list, and allows the user to select only items from its list. An editable combo box combines a text field and a list. In addition to selecting items from the list, the user may also type input into the text field. The default type of combo box is uneditable.
3. [5.](#) A mnemonic is a key on the keyboard that you press in combination with the Alt key to quickly access a component such as a button. When you assign a mnemonic to a button, the user can click the button by holding down the a key and pressing the mnemonic key.
4. [7.](#) A tool tip is text that is displayed in a small box when the user holds the mouse cursor over a component. The box usually gives a short description of what the component does.
5. [9.](#) The item is deselected, which causes the check mark to disappear. The checked menu item component also generates an action event.
6. [11.](#) Because, as the user moves the `JSlider` component's knob, it will only take on values within its established range.

# Chapter 25 (Available on the Book's Companion Website)

## Multiple Choice and True/False

1. [1.](#) c
2. [3.](#) b
3. [5.](#) b
4. [7.](#) d
5. [9.](#) c
6. [11.](#) a
7. [13.](#) d
8. [15.](#) c
9. [17.](#) c
10. [19.](#) c
11. [21.](#) b
12. [23.](#) True
13. [25.](#) True
14. [27.](#) False
15. [29.](#) True
16. [31.](#) True
17. [33.](#) False
18. [35.](#) False

# Find the Error

1. [1.](#) The tag should specify the file `MyApplet.class` instead of `MyApplet.java`.
2. [3.](#) Call `repaint` instead of `paint`.
3. [5.](#) The class must provide all of the methods specified by the `MouseListener` interface.

## Algorithm Workbench

1. [1.](#) `<center><h1>My Home Page</h1></center>`
2. [3.](#)

- Line 1: Change `JFrame` to `JApplet`
- Line 3: Change to `public void init()`
- Line 5: Delete
- Line 6: Delete
- Line 8: Delete
- Line 9: Delete
- Line 15: Delete
- Line 16: Delete
- Line 17: Delete

3. [5.](#)

```
private class MyMouseMotionListener extends MouseAdapter
{
 public void mouseMoved(MouseEvent e)
 {
 mouseMovements += 1;
 }
}
```

## Short Answer

1. [1.](#) It is executed by the user's system.
2. [3.](#) Applets are important because they can be used to extend the capabilities of a Web page. Web pages are normally written in Hypertext Markup Language (HTML). HTML is limited, however, because it merely describes the content and layout of a Web page, and creates links to other files and Web pages. HTML does

not have sophisticated abilities such as performing math calculations and interacting with the user. A programmer can write a Java applet to perform these types of operations and associate it with a Web page.

3. [5.](#) Some browsers, such as older versions of Microsoft Internet Explorer, do not directly support the Swing classes in applets. These browsers require a plug-in in order to run applets that use Swing components. If you are writing an applet for other people to run on their computers, there is no guarantee that they will have the required plug-in. If this is the case, you should use the AWT classes instead of the Swing classes for the components in your applet.
4. [7.](#) When the component is first displayed and is called again any time the component needs to be redisplayed.
5. [9.](#) If you want to load the sound file and keep it in memory so it can be played more than once, or if you want to play the sound file repeatedly.



# CASE STUDY 1 Calculating Sales Commission\*

*\*From Starting Out with Java: From Control Structures through Objects, Sixth Edition. Tony Gaddis. Copyright © 2016 by Pearson Education, Inc. All rights reserved.*

Hal’s Home Computer Emporium is a retail seller of home computers. Hal’s sales staff works strictly on commission. At the end of the month, each salesperson’s commission is calculated according to [Table CS1-1](#).

## Table CS1-1 Sales commission rates

### Sales This Month Commission Rate

|                    |     |
|--------------------|-----|
| less than \$10,000 | 5%  |
| \$10,000–14,999    | 10% |
| \$15,000–17,999    | 12% |
| \$18,000–21,999    | 15% |
| \$22,000 or more   | 16% |

For example, a salesperson with \$16,000 in monthly sales will earn a 12% commission (\$1,920). Another salesperson with \$20,000 in monthly sales will earn a 15% commission (\$3,000).

Because the staff gets paid once per month, Hal allows each employee to take up to \$1,500 per month in advance. When sales commissions are calculated, the amount of each employee’s advanced pay is subtracted from the commission. If any salesperson’s commission is less than the amount of this advance, he or she must reimburse Hal for the difference.

Here are two examples: Beverly and John have \$21,400 and \$12,600 in sales, respectively. Beverly’s commission is \$3,210 and John’s commission is \$1,260. Both Beverly and John took

\$1,500 in advance pay. At the end of the month, Beverly gets a check for \$1,710, but John must pay \$240 back to Hal.

Now we will examine a program that eases the task of calculating the end-of-month commission.

# Program Design

The program must perform the following general steps:

1. Ask the user to enter the salesperson's monthly sales.
2. Ask the user to enter the amount of advanced pay.
3. Use the amount of monthly sales to determine the commission rate.
4. Calculate the commission.
5. Calculate the salesperson's pay by subtracting the amount of advanced pay from the commission. If the amount is negative, the salesperson must reimburse the company.

[Table CS1-2](#) lists and describes the program's variables.

## Table CS1-2 Variables

| Variable                | Description                                                                                  |
|-------------------------|----------------------------------------------------------------------------------------------|
| <code>sales</code>      | A <code>double</code> variable to hold a salesperson's total monthly sales.                  |
| <code>rate</code>       | A <code>double</code> variable to hold the salesperson's commission rate.                    |
| <code>commission</code> | A <code>double</code> variable to hold the commission.                                       |
| <code>advance</code>    | A <code>double</code> variable to hold the amount of advanced pay the salesperson has drawn. |
| <code>Pay</code>        | A <code>double</code> variable to hold the salesperson's amount of gross pay.                |

Programmers often write parts of a program in pseudocode before writing the actual Java code. Recall from Chapter 1 that *pseudocode* is a cross between human language and a programming language. It is an informal way of writing down each operation that a program must perform, and is

especially helpful when designing an algorithm. Although the computer can't understand pseudocode, most programmers find pseudocode helpful in creating a model of a program. When the pseudocode model is complete, it can be converted to actual code. Here is a pseudocode model for the sales commission program:

```
Display "Enter the amount of monthly sales".
Read sales.
Display "Enter the amount of advanced pay".
Read advance.
If sales is less than 10,000
 Store 0.05 in rate.
else if sales is less than 15,000
 Store 0.1 in rate.
else if sales is less than 18,000
 Store 0.12 in rate.
else if sales is less than 22,000
 Store 0.15 in rate.
else
 Store 0.16 in rate.
commission = sales * rate.
pay = commission - advance.
Display pay.
```

[Code Listing CS1-1](#) shows the actual program. [Figure CS1-1](#) shows an example of interaction with the program.




**Figure CS1-1 Interaction with the  
HalsCommission program**

[Figure CS1-1 Full Alternative Text](#)

# Code Listing CS1-1 (HalsCommission.java)

```
1 import javax.swing.JOptionPane; // Needed for JOptionPane
2
3 /**
4 * This program calculates a salesperson's gross
5 * pay at Hal's Computer Emporium.
6 */
7
8 public class HalsCommission
9 {
10 public static void main(String[] args)
11 {
12 String input; // To hold the user's input
13 double sales; // Monthly sales
14 double rate; // Commission rate
15 double commission; // Amount of commission
16 double advance; // Advanced pay
17 double pay; // Salesperson's pay
18
19 // Get the monthly sales.
20 input = JOptionPane.showInputDialog("Enter the amount " +
21 "of monthly sales.");
22 sales = Double.parseDouble(input);
23
24 // Get the advanced pay.
25 input = JOptionPane.showInputDialog("Enter the amount " +
26 "of advanced pay.");
27 advance = Double.parseDouble(input);
28
29 // Determine the rate of commission.
30 if (sales < 10000)
31 rate = 0.05; // 5% commission rate
32 else if (sales < 15000)
33 rate = 0.1; // 10% commission rate
34 else if (sales < 18000)
35 rate = 0.12; // 12% commission rate
36 else if (sales < 22000)
37 rate = 0.15; // 15% commission rate
38 else
39 rate = 0.16; // 16% commission rate
40
41 // Calculate the amount of commission.
42 commission = rate * sales;
43
44 // Calculate the salesperson's pay.
45 pay = commission - advance;
46
47 // Display the salesperson's commission rate and pay.
48 JOptionPane.showMessageDialog(null,
49 String.format("Commission rate is %.2f. The amount of pay is $%,.2f",
50 rate, pay));
51 System.exit(0);
52 }
53 }
```



# CASE STUDY 2 The Amortization Class\*

\*From *Starting Out with Java: From Control Structures through Objects*, Sixth Edition. Tony Gaddis. Copyright © 2016 by Pearson Education, Inc. All rights reserved.

The loan officer at one of the Central Mountain Credit Union's branch offices has asked you to write a loan amortization application to run on her desktop PC. The application should allow the user to enter the amount of a loan, the number of years of the loan, and the annual interest rate. An amortization report should then be saved to a text file.

# Calculations

The credit union uses the following formula to calculate the monthly payment of a loan:

$$\text{Payment} = \text{Loan} \times \text{Rate} \times \frac{1 - (1 + \text{Rate}/12)^{-\text{Term}}}{1 - (1 + \text{Rate}/12)^{-12}}$$

where:

- Loan=the amount of the loan,
- Rate=the annual interest rate, and
- Term= $(1 + \text{Rate}/12)^{\text{Years}} \times 12$



# Report Requirements

The report produced by the program should show the monthly payment amount and the following information for each month in the loan period: amount applied to interest, amount applied to principal, and the balance. The following report may be used as a model. It shows all the required information on a one-year \$5,000 loan at 5.9 percent annual interest.

## Monthly Payment: \$430.10

| Month | Interest | Principal | Balance |
|-------|----------|-----------|---------|
|-------|----------|-----------|---------|

|    |       |        |          |
|----|-------|--------|----------|
| 1  | 24.58 | 405.52 | 4,594.48 |
| 2  | 22.59 | 407.51 | 4,186.97 |
| 3  | 20.59 | 409.52 | 3,777.45 |
| 4  | 18.57 | 411.53 | 3,365.92 |
| 5  | 16.55 | 413.55 | 2,952.37 |
| 6  | 14.52 | 415.59 | 2,536.78 |
| 7  | 12.47 | 417.63 | 2,119.15 |
| 8  | 10.42 | 419.68 | 1,699.47 |
| 9  | 8.36  | 421.75 | 1,277.72 |
| 10 | 6.28  | 423.82 | 853.90   |
| 11 | 4.20  | 425.90 | 428.00   |

The core of the program will be a class, `Amortization`, that holds the primary data, performs the mathematical calculations, and displays the report. [Figure CS2-1](#) shows a UML diagram for the class.




**Figure CS2-1 UML diagram for the `Amortization` class**

[Figure CS2-1 Full Alternative Text](#)

[Table CS2-1](#) lists and describes the class's fields.

**Table CS2-1 `Amortization` class fields**

| Field                     | Description                                                      |
|---------------------------|------------------------------------------------------------------|
| <code>loanAmount</code>   | A <code>double</code> variable to hold the amount of the loan.   |
| <code>interestRate</code> | A <code>double</code> variable to hold the annual interest rate. |

`loanBalance`    A `double` variable to hold the loan balance.

`term`            A `double` variable used in the calculation of the monthly payment.

`payment`        A `double` variable to hold the amount of the monthly payment.

`loanYears`      An `int` variable to hold the number of years of the loan.

[Table CS2-2](#) lists and describes the class's methods.

## Table CS2-2 Amortization class methods

| Method                           | Description                                                                                                                                                                                                                                 |
|----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Constructor                      | The constructor accepts three arguments: the loan amount, the annual interest rate, and the number of years of the loan. These values are stored in their corresponding fields. The private method <code>calcPayment</code> is then called. |
| <code>calcPayment</code>         | A private method that is used to calculate the monthly payment amount. The result is stored in the <code>payment</code> field.                                                                                                              |
| <code>getNumberOfPayments</code> | Returns as an <code>int</code> the number of loan payments.                                                                                                                                                                                 |
| <code>saveReport</code>          | Saves the amortization report to a text file.                                                                                                                                                                                               |
| <code>getLoanAmount</code>       | Returns as a <code>double</code> the amount of the loan.                                                                                                                                                                                    |
| <code>getInterestRate</code>     | Returns as a <code>double</code> the annual interest rate.                                                                                                                                                                                  |
| <code>getLoanYears</code>        | Returns as an <code>int</code> the number of years of the loan.                                                                                                                                                                             |

[Code Listing CS2-1](#) shows the code for the class.

## Code Listing CS2-1 (Amortization.java)

```
1 import java.io.*; // For file-related classes
2
3 /**
4 * This class stores loan information and creates a
5 * text file containing an amortization report.
6 */
7
8 public class Amortization
9 {
10 private double loanAmount; // Loan Amount
11 private double interestRate; // Annual Interest Rate
12 private double loanBalance; // Monthly Balance
13 private double term; // Payment Term
14 private double payment; // Monthly Payment
15 private int loanYears; // Years of Loan
16
17 /**
18 * The constructor accepts the loan amount, the annual
19 * interest rate, and the number of years of the loan
20 * as arguments. The private method CalcPayment is then
21 * called.
22 * @param loan The loan amount.
23 * @param rate The annual interest rate.
24 * @param years The number of years of the loan.
25 */
26
27 public Amortization(double loan, double rate, int years)
28 {
29 loanAmount = loan;
30 loanBalance = loan;
31 interestRate = rate;
32 loanYears = years;
33 calcPayment();
34 }
35
36 /**
37 * The calcPayment method calculates the monthly payment
38 * amount. The result is stored in the payment field.
39 */
40
41 private void calcPayment()
42 {
43 // Calculate value of Term
44 term =
45 Math.pow((1+interestRate/12.0), 12.0 * loanYears);
46
47 // Calculate monthly payment
48 payment =
```

```

49 (loanAmount * interestRate/12.0 * term) / (term - 1);
50 }
51
52 /**
53 * The getNumberOfPayments method returns the total number of
54 * payments to be made for the loan.
55 * @return The number of loan payments.
56 */
57
58 public int getNumberOfPayments()
59 {
60 return 12 * loanYears;
61 }
62
63 /**
64 * The saveReport method saves the amortization report to
65 * the file named by the argument.
66 * @param filename The name of the file to create.
67 */
68
69 public void saveReport(String filename) throws IOException
70 {
71 double monthlyInterest; // The monthly interest rate
72 double principal; // The amount of principal
73 //DecimalFormat dollar = new DecimalFormat("#,##0.00");
74 FileWriter fwriter = new FileWriter(filename);
75 PrintWriter outputFile = new PrintWriter(fwriter);
76
77 // Print monthly payment amount.
78 outputFile.println(String.format(
79 "Monthly Payment: $%.2f", payment));
80
81 // Print the report header.
82 outputFile.println("Month\tInterest\tPrincipal\tBalance");
83 outputFile.println("-----" +
84 "-----");
85
86 // Display the amortization table.
87 for (int month = 1; month <= getNumberOfPayments(); month++)
88 {
89 // Calculate monthly interest.
90 monthlyInterest = interestRate / 12.0 * loanBalance;
91
92 if (month != getNumberOfPayments())
93 {
94 // Calculate payment applied to principal
95 principal = payment - monthlyInterest;
96 }
97 else // This is the last month.
98 {
99 principal = loanBalance;
100 payment = loanBalance + monthlyInterest;
101 }
102
103 // Calculate the new loan balance.
104 loanBalance -= principal;
105
106 // Display a line of data.

```

```
107 outputFile.println(String.format("%d\t%.2f\t\t%.2f\t\t%.2f",
108 month, monthlyInterest, principal,
109 loanBalance));
110 }
111
112 // Close the file.
113 outputFile.close();
114 }
115
116 /**
117 The getLoanAmount method returns the loan amount.
118 @return The value in the loanAmount field.
119 */
120
121 public double getLoanAmount()
122 {
123 return loanAmount;
124 }
125
126 /**
127 The getInterestRate method returns the interest rate.
128 @return The value in the interestRate field.
129 */
130
131 public double getInterestRate()
132 {
133 return interestRate;
134 }
135
136 /**
137 The getLoanYears method returns the years of the loan.
138 @return The value in the loanYears field.
139 */
140
141 public int getLoanYears()
142 {
143 return loanYears;
144 }
145 }
```

# The Main Program

The main program code is shown in [Code Listing CS2-2](#). First, it gets the amount of the loan, the annual interest rate, and the years of the loan as input from the user. It then creates an instance of the `Amortization` class and passes this data to the class's constructor. The program then saves the amortization report in the file `LoanAmortization.txt`. It asks the user whether he or she wants to run another report. If so, the program repeats these steps. [Figure CS2-2](#) shows an example of interaction with the program.




**Figure CS2-2 Interaction with the `LoanReport` program**

# Code Listing CS2-2 (LoanReport.java)

```
1 import javax.swing.JOptionPane; // For the JOptionPane class
2
3 // The following import statement is required because the main
4 // method has a throws IOException clause. Although this program
5 // doesn't have code that directly performs file I/O, the import
6 // statement is still required because of IOException.
7 import java.io.*;
8
9 /**
10 This program displays a loan amortization report.
11 */
12
13 public class LoanReport
14 {
15 public static void main(String[] args) throws IOException
16 {
17 String input; // To hold user input
18 double loan; // Loan amount
19 double interestRate; // Annual interest rate
20 int years; // Years of the loan
21 char again; // To indicate if loop should repeat
22
23 do
24 {
25 // Get the loan amount.
26 input = JOptionPane.showInputDialog("Enter the " +
27 "loan amount.");
28 loan = Double.parseDouble(input);
29
30 // Validate the loan amount.
31 // (No negative amounts.)
32 while (loan < 0)
33 {
34 input = JOptionPane.showInputDialog("Invalid amount. " +
35 "Enter the loan amount.");
36 loan = Double.parseDouble(input);
37 }
38
39 // Get the annual interest rate.
40 input = JOptionPane.showInputDialog("Enter the " +
41 "annual interest rate.");
42 interestRate = Double.parseDouble(input);
43 // Validate the interest rate
44 // (No negative amounts.)
45 while (interestRate < 0)
46 {
47 input = JOptionPane.showInputDialog("Invalid amount. " +
48 "Enter the annual interest rate.");
49 interestRate = Double.parseDouble(input);
50 }
51
52 // Get the years of the loan.
53 input = JOptionPane.showInputDialog("Enter the " +
```



```

54 "years of the loan.");
55 years = Integer.parseInt(input);
56
57 // Validate the number of years.
58 // (No negative amounts.)
59 while (years < 0)
60 {
61 input = JOptionPane.showInputDialog("Invalid amount. " +
62 "Enter the years of the loan.");
63 years = Integer.parseInt(input);
64 }
65
66 // Create and initialize an Amortization object.
67 Amortization am =
68 new Amortization(loan, interestRate, years);
69
70 // Save the report.
71 am.saveReport("LoanAmortization.txt");
72 JOptionPane.showMessageDialog(null, "Report saved to " +
73 "the file LoanAmortization.txt.");
74
75 // Do another report?
76 input = JOptionPane.showInputDialog("Would you like " +
77 "to run another report? Enter Y for " +
78 "yes or N for no: ");
79 again = input.charAt(0);
80
81 } while (again == 'Y' || again == 'y');
82
83 System.exit(0);
84 }
85 }

```

Contents of the file *LoanAmortization.txt*:

# Monthly Payment:

## \$430.10

|  | Month | Interest | Principal | Balance |
|--|-------|----------|-----------|---------|
|--|-------|----------|-----------|---------|

|   |       |        |          |
|---|-------|--------|----------|
| 1 | 24.58 | 405.52 | 4,594.48 |
|---|-------|--------|----------|

|   |       |        |          |
|---|-------|--------|----------|
| 2 | 22.59 | 407.51 | 4,186.97 |
|---|-------|--------|----------|

|   |       |        |          |
|---|-------|--------|----------|
| 3 | 20.59 | 409.52 | 3,777.45 |
|---|-------|--------|----------|

|   |       |        |          |
|---|-------|--------|----------|
| 4 | 18.57 | 411.53 | 3,365.92 |
|---|-------|--------|----------|

|    |       |        |          |
|----|-------|--------|----------|
| 5  | 16.55 | 413.55 | 2,952.37 |
| 6  | 14.52 | 415.59 | 2,536.78 |
| 7  | 12.47 | 417.63 | 2,119.15 |
| 8  | 10.42 | 419.68 | 1,699.47 |
| 9  | 8.36  | 421.75 | 1,277.72 |
| 10 | 6.28  | 423.82 | 853.90   |
| 11 | 4.20  | 425.90 | 428.00   |
| 12 | 2.10  | 428.00 | 0.00     |

First, notice that a `do-while` loop (in lines 23 through 81) controls everything done in this program. Here is a condensed version of the loop:

```

23 do
24 {

(Code to get the loan data and display the report.)

75 // Do another report?
76 input = JOptionPane.showInputDialog("Would you like " +
77 "to run another report? Enter Y for " +
78 "yes or N for no: ");
79 again = input.charAt(0);
80
81 } while (again == 'Y' || again == 'y');
```

During each iteration, this loop first gathers the necessary loan data as input and then writes the amortization report. At the end of the iteration, the user is asked whether he or she wants to run another report. If the user enters `Y` or `y` for yes, the loop repeats. Otherwise, the loop terminates and the program ends.

Inside the loop, the code that gathers input from the user also validates the input. For example, here is the code in lines 25 through 37 that gets and validates the loan amount:

```

25 // Get the loan amount.
```

```

26 input = JOptionPane.showInputDialog("Enter the " +
27 "loan amount.");
28 loan = Double.parseDouble(input);
29
30 // Validate the loan amount.
31 // (No negative amounts.)
32 while (loan < 0)
33 {
34 input = JOptionPane.showInputDialog
35 ("Invalid amount. Enter the loan amount.");
36 loan = Double.parseDouble(input);
37 }

```

If the user enters a negative number, the `while` loop displays an error message and asks the user to enter the loan amount again. The code that gets the annual interest rate and the years of the loan, in lines 39 through 64, also performs similar input validation. Once these values have been correctly entered, an instance of the `Amortization` class is created in lines 67 and 68:

```

Amortization am =
 new Amortization(loan, interestRate, years);

```

Then, the `saveReport` method is called in line 71 to save the amortization report to the file `LoanAmortization.txt`:

```

am.printReport("LoanAmortization.txt");

```

## **NOTE:**

Notice that [Code Listing CS2-1](#) has the following import statement in line 1:

```

import java.io.*;

```

Although this program does not have any code that directly performs file I/O, we still have to have this import statement because of the `throws IOException` clause in the `main` method header. The `main` method has to have the `throws IOException` clause because it calls the `saveReport` method in the `Amortization` class.

# CASE STUDY 3 The PinTester Class\*

\*From *Starting Out with Java: From Control Structures through Objects*, Sixth Edition. Tony Gaddis. Copyright © 2016 by Pearson Education, Inc. All rights reserved.

The National Commerce Bank has hired you as a contract programmer. Your first assignment is to write a class that will be used by the bank’s automated teller machines (ATMs) to validate a customer’s personal identification number (PIN). Your class, which will be named `PinTester`, will be incorporated into a larger program that asks the customer to input his or her PIN on the ATM’s numeric keypad. (PINs are multi-digit numbers. The program stores each digit in an element of an `int` array.) The program also retrieves a copy of the customer’s actual PIN from a database, as an `int` array. If these arrays match, then the customer’s identity is validated.

[Table CS3-1](#) lists the fields that the class should have.

## Table CS3-1 PinTester class fields

| Field Name            | Description                                                                                                                                                           |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>userPIN</code>  | This field should be an <code>int</code> array reference variable. It will be used to reference an array containing the PIN digits entered by the user.               |
| <code>validPIN</code> | This field should be an <code>int</code> array reference variable. It will be used to reference an array containing the valid PIN digits retrieved from the database. |

[Table CS3-2](#) lists the methods that the class should have.

## Table CS3-2 PinTester class methods

| Method | Description                                                                                                                                                             |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|        | The constructor is to accept references to two <code>int</code> arrays as arguments. The array referenced by the first argument will contain the numbers entered by the |

customer. The array referenced by the second argument will contain the Constructor numbers retrieved from the bank's database. The constructor should create two arrays and reference them by the `userPIN` and `validPIN` fields. Then, it should copy the contents of the arrays passed as arguments to the `userPIN` and `validPIN` arrays.

`isValid` This boolean function should return `true` if the `userPIN` and `validPIN` arrays are identical. Otherwise, it should return `false`.

[Code Listing CS3-1](#) shows the `PinTester` class.

## Code Listing CS3-1 (`PinTester.java`)

```
1 /**
2 The PinTester class is used to compare the PIN that
3 was entered by a user to the valid PIN.
4 */
5
6 public class PinTester
7 {
8 private int[] userPIN; // The PIN entered
9 private int[] validPIN; // The valid PIN
10
11 /**
12 Constructor
13 @param user The PIN digits entered by the user.
14 @param valid The valid PIN digits.
15 */
16
17 PinTester(int[] user, int[] valid)
18 {
19 // Create an array to hold a copy of the
20 // PIN digits entered by the user.
21 userPIN = new int[user.length];
22
23 // Copy the user array to the userPIN array.
24 for (int i = 0; i < user.length; i++)
25 userPIN[i] = user[i];
26
27 // Create an array to hold a copy of the
28 // valid PIN digits.
29 validPIN = new int[valid.length];
30
31 // Copy the valid array to the validPIN array.
32 for (int i = 0; i < valid.length; i++)
33 validPIN[i] = valid[i];
34 }
35
36 /**
```

```

37 The isValid method compares the PIN digits entered
38 by the user and determines whether they are valid.
39 @return true if valid, false otherwise.
40 */
41
42 public boolean isValid()
43 {
44 int i; // Loop control variable
45 boolean isEqual; // Flag
46
47 // Compare the lengths of the two arrays. If
48 // the arrays are not the same size, then they
49 // are not equal.
50
51 if (userPIN.length != validPIN.length)
52 isEqual = false;
53 else // Compare the contents of the two arrays
54 {
55 i = 0;
56 isEqual = true;
57
58 while (isEqual && i < userPIN.length)
59 {
60 if (userPIN[i] != validPIN[i])
61 isEqual = false;
62 i++;
63 }
64 }
65
66 return isEqual;
67 }
68 }

```

Because you have only been asked to write a method that performs the comparison between the customer's input and the PIN that was retrieved from the database, you will also need to write a program that tests the method. [Code Listing CS3-2](#) shows the complete program.

## Code Listing CS3-2 (PinTesterDemo.java)

```

1 /**
2 * This program tests the PinTester class with
3 * various PIN values.
4 */
5
6 public class PinTesterDemo
7 {
8 public static void main(String [] args)
9 {
10 // Create two sample arrays that are identical.
11 int [] pin1 = {2, 4, 1, 8, 7, 9, 0};
12 int [] pin2 = {2, 4, 1, 8, 7, 9, 0};
13
14 // Create two sample arrays that are different.
15 int [] pin3 = {2, 4, 6, 8, 7, 9, 0};

```

```

16 int [] pin4 = {1, 2, 3, 4, 5, 6, 7};
17
18 // Create two PinTester objects.
19 PinTester pt1 = new PinTester(pin1, pin2);
20 PinTester pt2 = new PinTester(pin3, pin4);
21
22 // Compare pin1 and pin2. These arrays are identical.
23 if (pt1.isValid())
24 {
25 System.out.println("SUCCESS: pin1 and pin2 " +
26 "are the same.");
27 }
28 else
29 {
30 System.out.println("ERROR: pin1 and pin2 " +
31 "are different.");
32 }
33
34 // Compare pin3 and pin4. All of the elements in these
35 // arrays are different.
36
37 if (pt2.isValid())
38 {
39 System.out.println("ERROR: pin3 and pin4 " +
40 "are the same.");
41 }
42 else
43 {
44 System.out.println("SUCCESS: pin3 and pin4 " +
45 "are different.");
46 }
47 }
48 }

```

# Program Output

```

SUCCESS: pin1 and pin2 are the same.
SUCCESS: pin3 and pin4 are different.

```

# CASE STUDY 4 Parallel Arrays\*

\*From *Starting Out with Java: From Control Structures through Objects*, Sixth Edition. Tony Gaddis. Copyright © 2016 by Pearson Education, Inc. All rights reserved.

Sometimes it's useful to store related data in two or more arrays. For example, assume a program uses the following arrays:

```
String[] names = new String[5];
String[] addresses = new String[5];
```

The `names` array is used to store the names of five persons, and the `addresses` array is used to store the addresses of the same five persons. The data for one person is stored in the same relative location in each array. For example, the first person's name is stored at `names[0]` and that same person's address is stored at `addresses[0]`. This relationship between the arrays is illustrated in [Figure CS4-1](#).




**Figure CS4-1 The relationship between the names and addresses arrays**

[Figure CS4-1 Full Alternative Text](#)

To access the data, use the same subscript with both arrays. For example, the following loop displays each person's name and address.



```

for (int i = 0; i < names.length; i++)
{
 System.out.println("Name: " + names[i]);
 System.out.println("Address: " + addresses[i]);
}

```

The names and addresses arrays are examples of parallel arrays. *Parallel arrays* are two or more arrays that hold related data, and the related elements of each array are accessed with a common subscript.

Parallel arrays are especially useful when the related data is of unlike types. For example, the program in Code Listing [Code Listing CS4-1](#) is a payroll program. It uses two arrays: one to store the hours worked by each employee (as `int` values), and another to store each employee's hourly pay rate (as `double` values).

## Code Listing CS4-1 (ParallelArrays.java)

```

1 import java.util.Scanner;
2
3 /**
4 * This program demonstrates parallel arrays.
5 */
6
7 public class ParallelArrays
8 {
9 public static void main(String [] args)
10 {
11 final int NUM_EMPLOYEES = 3;
12 int[] hours = new int[NUM_EMPLOYEES];
13 double[] payRates = new double[NUM_EMPLOYEES];
14
15 // Get the hours worked by each employee.
16 getPayData(hours, payRates);
17
18 // Display each employee's gross pay.
19 displayGrossPay(hours, payRates);
20 }
21
22 /**
23 * The getPayData method accepts as arguments arrays
24 * for employees' hours and pay rates. The user enters
25 * values for these arrays.
26 */
27
28 private static void getPayData(int[] hours, double[] payRates)
29 {
30 // Create a Scanner object for keyboard input.
31 Scanner keyboard = new Scanner(System.in);
32
33 // Get each employee's hours worked and pay rate.
34 for (int i = 0; i < hours.length; i++)
35 {
36 // Get the hours worked for an employee.

```

```

37 System.out.printf(
38 "Enter the hours worked by employee #d: ",
39 (i + 1));
40 hours[i] = keyboard.nextInt();
41
42 // Get the hourly pay rate for this employee.
43 System.out.printf(
44 "Enter the hourly pay rate for employee #d: ",
45 (i + 1));
46 payRates[i] = keyboard.nextDouble();
47 }
48 }
49
50 /**
51 The displayGrossPay method accepts as arguments
52 arrays for employees' hours and pay rates. The
53 method uses these arrays to calculate and display
54 each employee's gross pay.
55 */
56
57 private static void displayGrossPay(int [] hours,
58 double [] payRates)
59 {
60 double grossPay; // To hold gross pay
61
62 // Display each employee's gross pay.
63 for (int i = 0; i < hours.length; i++)
64 {
65 // Calculate the gross pay.
66 grossPay = hours[i] * payRates[i];
67
68 // Display the gross pay.
69 System.out.printf(
70 "The gross pay for employee #d is $%,.2f.\n",
71 (i + 1), grossPay);
72 }
73 }
74 }

```

# Program Output with Example Input Shown in Bold

```

Enter the hours worked by employee #1: 15 [Enter]
Enter the hourly pay rate for employee #1: 12.00 [Enter]
Enter the hours worked by employee #2: 20 [Enter]
Enter the hourly pay rate for employee #2: 18.00 [Enter]
Enter the hours worked by employee #3: 40 [Enter]
Enter the hourly pay rate for employee #3: 25.50 [Enter]
The gross pay for employee #1 is $180.00.
The gross pay for employee #2 is $360.00.
The gross pay for employee #3 is $1,020.00.

```

# CASE STUDY 5 The FeetInches Class\*

\*From *Starting Out with Java: From Control Structures through Objects*, Sixth Edition. Tony Gaddis. Copyright © 2016 by Pearson Education, Inc. All rights reserved.

Recall from Chapter 3 that the `String` class has a method, `equals`, that accepts a `String` object as an argument. As a review, here is an example of how the code can be used:

```
String name1 = "Joe";
String name2 = "Jo";
if (name1.equals(name2))
 System.out.println("The names are the same.");
else
 System.out.println("The names are not the same.");
```

This is an example of a *same class operation*, where an object performs an operation involving another object of the same class. In this section we use a class named `FeetInches` to study same class operations. The `FeetInches` class is designed to hold distances or measurements expressed in feet and inches. [Figure CS5-1](#) shows a UML diagram for the class.




**Figure CS5-1 UML diagram for the FeetInches class**

The class has two `int` fields: `feet` and `inches`. Together these fields hold a distance measured in feet and inches, such as 12 feet 7 inches. The `feet` field holds the feet part and the `inches` field holds the inches part. [Table CS5-1](#) describes the class’s methods.

**Table CS5-1 The FeetInches class methods**

| Method                                 | Description                                                                                                                                                                                                                                    |
|----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Constructor<br>#1 (no-arg constructor) | This constructor assigns 0 to both the <code>feet</code> and <code>inches</code> fields.                                                                                                                                                       |
| Constructor<br>#2                      | This constructor accepts two <code>int</code> arguments that are assigned to the <code>feet</code> and <code>inches</code> fields. The <code>simplify</code> method is also called.                                                            |
| Constructor<br>#3                      | This constructor accepts a <code>FeetInches</code> object as its argument. The object that is being created will become a copy of the object passed as an argument. This type of constructor is sometimes referred to as a “copy constructor.” |
| <code>simplify</code>                  | This method adjusts any set of values where the <code>inches</code> field is greater than 11. For example, 3 feet 14 inches would be adjusted to read 4 feet 2 inches.                                                                         |
| <code>setFeet</code>                   | This method accepts an <code>int</code> argument that is assigned to the <code>feet</code> field.                                                                                                                                              |
| <code>setInches</code>                 | This method accepts an <code>int</code> argument that is assigned to the <code>inches</code> field. The <code>simplify</code> method is then called.                                                                                           |
| <code>getFeet</code>                   | This method returns the value in the <code>feet</code> field.                                                                                                                                                                                  |
| <code>getInches</code>                 | This method returns the value in the <code>inches</code> field.                                                                                                                                                                                |

This method returns a string representing the distance held by the object. For

|                       |                                                                                                                                                                                                                                  |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>toString</code> | example, if an object's <code>feet</code> field holds 3 and its <code>inches</code> field holds 7, the <code>toString</code> method would return the string "3 feet 7 inches".                                                   |
| <code>add</code>      | This method accepts a <code>FeetInches</code> object as its argument. It returns a reference to a <code>FeetInches</code> object that is the sum of the calling object and the object that was passed as the argument.           |
| <code>equals</code>   | This method accepts a <code>FeetInches</code> object as its argument. It returns the boolean value <code>true</code> if the calling object and the argument object hold the same data. Otherwise it returns <code>false</code> . |
| <code>copy</code>     | This method returns a reference to a new <code>FeetInches</code> object that is a copy of the calling object.                                                                                                                    |

Before looking at the code for the `FeetInches` class, it might help to see a program that uses the class. The program in [Code Listing CS5-1](#) demonstrates the class.

## Code Listing CS5-1 (`DistanceDemo.java`)

```

1 import java.util.Scanner;
2
3 /**
4 * This program demonstrates the FeetInches class.
5 */
6
7 public class DistanceDemo
8 {
9 public static void main(String[] args)
10 {
11 int feet, inches; // To hold feet and inches
12 FeetInches distance; // To reference a FeetInches object
13
14 // Create a Scanner object for keyboard input.
15 Scanner keyboard = new Scanner(System.in);
16
17 // Get a distance in feet and inches.
18 System.out.println("Enter a distance in feet "
19 + "and inches.");
20 System.out.print("First enter the feet: ");
21 feet = keyboard.nextInt();
22 System.out.print("Now enter the inches: ");
23 inches = keyboard.nextInt();
24
25 // Create a FeetInches object, initialized with
26 // the values entered by the user.

```

```

27 distance = new FeetInches(feet, inches);
28
29 // Display the object's contents.
30 System.out.println("The distance you entered is "
31 + distance.toString());
32 }
33 }

```

## Program Output with Example Input Shown in Bold

```

Enter a distance in feet and inches.
First enter the feet: 5 [Enter]
Now enter the inches: 7 [Enter]
The distance you entered is 5 feet 7 inches

```

## Program Output with Example Input Shown in Bold

```

Enter a distance in feet and inches.
First enter the feet: 5 [Enter]
Now enter the inches: 14 [Enter]
The distance you entered is 6 feet 2 inches

```

Notice that the `FeetInches` class automatically adjusts measurements so that the `inches` field is never greater than 11. For example, the distance 5 feet 14 inches was adjusted to the more proper measurement 6 feet 2 inches. This adjustment is performed by a private method named `simplify`, which is called from the constructor and the `setInches` method.

Also notice that the class has a `toString` method which returns a string representing the distance held by an object. This method is useful for displaying a distance. [Code Listing CS5-2](#) shows the `FeetInches` class.

## Code Listing CS5-2 (`FeetInches.java`)

```

1 /**
2 * The FeetInches class holds distances measured in
3 * feet and inches.
4 */
5
6 public class FeetInches
7 {
8 private int feet; // The number of feet
9 private int inches; // The number of inches

```

```

10
11 /**
12 This constructor assigns 0 to the feet
13 and inches fields.
14 */
15
16 public FeetInches()
17 {
18 feet = 0;
19 inches = 0;
20 }
21
22 /**
23 This constructor accepts two arguments which
24 are assigned to the feet and inches fields.
25 The simplify method is then called.
26 @param f The value to assign to feet.
27 @param i The value to assign to inches.
28 */
29
30 public FeetInches(int f, int i)
31 {
32 feet = f;
33 inches = i;
34 simplify();
35 }
36
37 /**
38 The following is a copy constructor. It accepts a
39 reference to another FeetInches object. The feet
40 and inches fields are set to the same values as
41 those in the argument object.
42 @param object2 The object to copy.
43 */
44
45 public FeetInches(FeetInches object2)
46 {
47 feet = object2.feet;
48 inches = object2.inches;
49 }
50
51 /**
52 The simplify method adjusts the values
53 in feet and inches to conform to a
54 standard measurement.
55 */
56
57 private void simplify()
58 {
59 if (inches > 11)
60 {
61 feet = feet + (inches / 12);
62 inches = inches % 12;
63 }
64 }
65
66 /**
67 The setFeet method assigns a value to

```

```

68 the feet field.
69 @param f The value to assign to feet.
70 */
71
72 public void setFeet(int f)
73 {
74 feet = f;
75 }
76
77 /**
78 The setInches method assigns a value to
79 the inches field.
80 @param i The value to assign to inches.
81 */
82
83 public void setInches(int i)
84 {
85 inches = i;
86 simplify();
87 }
88
89 /**
90 getFeet method
91 @return The value in the feet field.
92 */
93
94 public int getFeet()
95 {
96 return feet;
97 }
98
99 /**
100 getInches method
101 @return The value in the inches field.
102 */
103
104 public int getInches()
105 {
106 return inches;
107 }
108
109 /**
110 toString method
111 @return a reference to a String stating
112 the feet and inches.
113 */
114
115 public String toString()
116 {
117 String str = feet + " feet " +
118 inches + " inches";
119 return str;
120 }
121
122 /**
123 The add method returns a FeetInches object
124 that holds the sum of this object and another

```



```

125 FeetInches object.
126 @param object2 The other FeetInches object.
127 @return A reference to a FeetInches object.
128 */
129
130 public FeetInches add(FeetInches object2)
131 {
132 int totalFeet, // To hold the sum of feet
133 totalInches; // To hold the sum of inches
134
135 totalFeet = feet + object2.feet;
136 totalInches = inches + object2.inches;
137 return new FeetInches(totalFeet, totalInches);
138 }
139
140 /**
141 * The equals method compares this object to the
142 * argument object. If both have the same values,
143 * the method returns true.
144 * @return true if the objects are equal, false
145 * otherwise.
146 */
147
148 public boolean equals(FeetInches object2)
149 {
150 boolean status;
151
152 if (object2 == null)
153 status = false;
154 else if (feet == object2.feet &&
155 inches == object2.inches)
156 status = true;
157 else
158 status = false;
159 return status;
160 }
161
162 /**
163 * The copy method makes a copy of the
164 * calling object.
165 * @return A reference to the copy.
166 */
167
168 public FeetInches copy()
169 {
170 // Make a new FeetInches object and
171 // initialize it with the same data
172 // as the calling object.
173 FeetInches newObject =
174 new FeetInches(feet, inches);
175
176 // Return a reference to the new object.
177 return newObject;
178 }
179 }

```

# More about the toString Method

The program in [Code Listing CS5-1](#) explicitly called the `FeetInches` class's `toString` method to display the contents of the `distance` object. The following statement appears in lines 30 and 31:

```
System.out.println("The distance you entered is "
 + distance.toString());
```

In actuality, it is unnecessary to explicitly call the `toString` method in this statement. If you write a `toString` method for a class, Java will automatically call the method any time you concatenate an object of the class with a string. For example, the following statement would implicitly call the `distance` object's `toString` method:

```
System.out.println("The distance you entered is "
 + distance);
```

Java also implicitly calls an object's `toString` method when the object is passed to the `print` or `println` method. Assuming that `distance` is a `FeetInches` object, here is an example:

```
System.out.println(distance);
```

This statement is equivalent to:

```
System.out.println(distance.toString());
```

## NOTE:

Every class automatically has a `toString` method that returns a string containing the object's class name, followed by the `@` symbol, followed by an integer that is usually based on the object's memory address. This method is called when necessary if you have not provided your own `toString` method. This is discussed further in [Chapter 8](#)

# The add Method

The `add` method is the first `FeetInches` method we will study that performs a same class operation. It allows us to add one `FeetInches` object to another. Let's take a closer look at the method, which appears in lines 130 through 138. For your convenience, the method is shown here:

```
130 public FeetInches add(FeetInches object2)
131 {
132 int totalFeet, // To hold the sum of feet
133 totalInches; // To hold the sum of inches
134
135 totalFeet = feet + object2.feet;
136 totalInches = inches + object2.inches;
137 return new FeetInches(totalFeet, totalInches);
138 }
```

This method accepts a `FeetInches` object as its argument. The parameter variable `object2` will reference the object that was passed as an argument. The statement in line 135 adds the `feet` field of the calling object to the `feet` field of the argument object, and stores the result in the local variable `totalFeet`. Notice that the method can access `object2`'s `feet` field directly. Because the `add` method is a member of the `FeetInches` class, it is allowed to access `object2`'s private fields.

Next, in line 136, the method adds the `inches` field of the calling object to the `inches` field of `object2`, and stores the result in the local variable `totalInches`.

The last statement, in line 137, creates a new `FeetInches` object, passes `totalFeet` and `totalInches` to the constructor, and returns a reference to the object. The program in [Code Listing CS5-3](#) demonstrates the method.

## Code Listing CS5-3 (DistanceAdd.java)

```
1 /**
2 * This program uses the FeetInches class's add method
3 * to add two distances.
4 */
5
6 public class DistanceAdd
7 {
8 public static void main(String[] args)
9 {
10 // Create a FeetInches object holding 5 feet, 9 inches.
11 FeetInches distance1 = new FeetInches(5, 9);
12
13 // Create a FeetInches object holding 2 feet, 5 inches.
14 FeetInches distance2 = new FeetInches(2, 5);
15
16 // Declare a FeetInches variable that does not
```

```
17 // reference an object.
18 FeetInches distance3;
19
20 // Add distance 1 and distance 2, store the
21 // result in distance3.
22 distance3 = distance1.add(distance2);
23
24 // Display the results.
25 System.out.println("The first distance is " + distance1);
26 System.out.println("The second distance is " + distance2);
27 System.out.println("The sum of these distances is " + distance3);
28 }
29 }
```

## Program Output

```
The first distance is 5 feet 9 inches
The second distance is 2 feet 5 inches
The sum of these distances is 8 feet 2 inches
```

# The equals Method

Now let's take a closer look at the `equals` method, which also performs a same class operation: It compares the contents of two `FeetInches` objects and determines whether the two objects are equal. This method is useful because you cannot use the `==` operator to compare the contents of two objects. For example, the following code appears to compare two `FeetInches` objects, but in reality does not.

```
FeetInches distance1 = new FeetInches(6, 5);
FeetInches distance2 = new FeetInches(6, 5);
if (distance1 == distance2) // This is a mistake.
 System.out.println("The objects are the same.");
else
 System.out.println("The objects are not the same.");
```

When you use the `==` operator with reference variables, the operator compares the memory addresses that the variables contain, not the contents of the objects referenced by the variables. This is illustrated in [Figure CS5-2](#).




**Figure CS5-2 The `if` statement tests the contents of the reference variables, not the contents of the objects the variables reference**

Because the two variables reference different objects in memory, they will contain different addresses. Therefore, the result of the boolean expression `distance1 == distance2` is false and the code reports that the objects are not the same. Instead of using the `==` operator to compare two `FeetInches` objects, we should use the `equals` method.

Here is the code for the method, which appears in lines 148 through 160:

```
148 public boolean equals(FeetInches object2)
149 {
150 boolean status;
151
152 if (object2 == null)
153 status = false;
154 else if (feet == object2.feet &&
155 inches == object2.inches)
156 status = true;
157 else
158 status = false;
159 return status;
160 }
```

This method accepts a `FeetInches` object as its argument. The parameter variable `object2` will reference the object that was passed as an argument. Lines 154 and 155 perform the following comparison: If the `feet` field of the calling object is equal to the `feet` field of `object2`, and the `inches` field of the calling object is equal to the `inches` field of `object2`, then the two objects contain the same distance. In this case, the local variable `status` (a boolean) is set to `true`. Otherwise, `status` is set to `false`. The method returns the value of the `status` variable in line 159. The program in [Code Listing CS5-4](#) demonstrates the `equals` method.

## Code Listing CS5-4 (`DistanceCompare.java`)

```
1 /**
2 * This program uses the FeetInches class's equals
3 * method to compare two distances.
4 */
5
6 public class DistanceCompare
7 {
8 public static void main(String[] args)
9 {
10 // Create a FeetInches object holding 5 feet, 9 inches.
11 FeetInches distance1 = new FeetInches(5, 9);
12
13 // Create a FeetInches object holding 5 feet, 9 inches.
14 FeetInches distance2 = new FeetInches(5, 9);
15
16 // Create a FeetInches object holding 7 feet, 4 inches.
17 FeetInches distance3 = new FeetInches(7, 4);
18
19 // Display the distances.
20 System.out.println("The first distance is " + distance1);
```

```

21 System.out.println("The second distance is " + distance2);
22 System.out.println("The third distance is " + distance3);
23
24 // Compare distance1 and distance2.
25 if (distance1.equals(distance2))
26 System.out.println("distances 1 and 2 are the same.");
27 else
28 System.out.println("distances 1 and 2 are not the same.");
29
30 // Compare distance1 and distance3.
31 if (distance1.equals(distance3))
32 System.out.println("distances 1 and 3 are the same.");
33 else
34 System.out.println("distances 1 and 3 are not the same.");
35 }
36 }

```

## Program Output

```

The first distance is 5 feet 9 inches
The second distance is 5 feet 9 inches
The third distance is 7 feet 4 inches
distances 1 and 2 are the same.
distances 1 and 3 are not the same.

```

If you want to be able to compare the objects of a given class, you should always write an `equals` method for the class.

## NOTE:

Every class automatically has an `equals` method, which works the same as the `==` operator. This method is called when necessary if you have not provided your own `equals` method. This is discussed further in [Chapter 10](#).

# The copy Method

The `copy` method also performs a same class operation: It returns a `FeetInches` object that is a copy of the calling object. This method is necessary because you cannot copy objects with a simple assignment statement, as you would with primitive variables. For example, look at the following code:

```
FeetInches distance1 = new FeetInches(4, 9);
FeetInches distance2 = distance1;
```

The first statement creates a `FeetInches` object and assigns its address to the `distance1` variable. The second statement assigns `distance1` to `distance2`. This does not make a copy of the object referenced by `distance1`. Rather, it makes a copy of the address that is stored in `distance1` and stores it in `distance2`. After this statement executes, both the `distance1` and `distance2` variables will reference the same object. This is illustrated in [Figure CS5-3](#).




**Figure CS5-3 Both variables reference the same object**

The `copy` method, which appears in lines 168 through 178, creates a new `FeetInches` object and passes the calling object's `feet` and `inches` fields as arguments to the constructor. This makes the new object a copy of the calling object. The method then returns a reference to the new object. Here is the method's code:

```
168 public FeetInches copy()
169 {
170 // Make a new FeetInches object and
171 // initialize it with the same data
172 // as the calling object.
173 FeetInches newObject =
174 new FeetInches(feet, inches);
```



```
175
176 // Return a reference to the new object.
177 return newObject;
178 }
```

# The Copy Constructor

The `FeetInches` class has three constructors: a no-arg constructor, a constructor that accepts arguments for the `feet` and `inches` fields, and a third constructor that accepts a `FeetInches` object as its argument. Here is the code for the third constructor, which appears in lines 45 through 49:

```
45 public FeetInches(FeetInches object2)
46 {
47 feet = object2.feet;
48 inches = object2.inches;
49 }
```

This type of constructor is sometimes referred to as a *copy constructor*. It makes the object that is being created a copy of the object that was passed as an argument. For example, the following code segment creates two `FeetInches` objects. The second one is a copy of the first one.

```
// Create a FeetInches object holding 5 feet, 9 inches.
FeetInches distance1 = new FeetInches(5, 9);
// Create another one that is a copy.
FeetInches distance2 = new FeetInches(distance1);
```

# CASE STUDY 6 The `SerialNumber` Class\*

\*From *Starting Out with Java: From Control Structures through Objects*, Sixth Edition. Tony Gaddis. Copyright © 2016 by Pearson Education, Inc. All rights reserved.

In this case study you examine the `SerialNumber` class, which is used by the Home Software Company to validate software serial numbers. A valid software serial number is in the form `LLLLL-DDDD-LLLL`, where `L` indicates an alphabetic letter and `D` indicates a numeric digit. For example, `WRXTQ-7786-PGVZ` is a valid serial number. Notice that a serial number consists of three groups of characters, delimited by hyphens. [Figure CS6-1](#) shows a UML diagram for the `SerialNumber` class.




**Figure CS6-1 UML diagram for the `SerialNumber` class**

[Figure CS6-1 Full Alternative Text](#)

The fields `first`, `second`, and `third` are used to hold the first, second, and third groups of characters in a serial number. The `valid` field is set to `true` by the constructor to indicate a valid serial number, or `false` to indicate an invalid serial number. [Table CS6-1](#) describes the class’s methods.

**Table CS6-1 `SerialNumber` class’s methods**

| Method | Description                                                      |
|--------|------------------------------------------------------------------|
|        | The constructor accepts a string argument that contains a serial |

|                                 |                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Constructor                     | number. The string is tokenized and its tokens are stored in the <code>first</code> , <code>second</code> , and <code>third</code> fields. The <code>validate</code> method is called.                                        |
| <code>isValid</code>            | This method returns the value in the <code>valid</code> field.                                                                                                                                                                |
| <code>validate</code>           | This method calls the <code>isFirstGroupValid</code> , <code>isSecondGroupValid</code> , and <code>isThirdGroupValid</code> methods to validate the <code>first</code> , <code>second</code> , and <code>third</code> fields. |
| <code>isFirstGroupValid</code>  | This method returns <code>true</code> if the value stored in the <code>first</code> field is valid. Otherwise, it returns <code>false</code> .                                                                                |
| <code>isSecondGroupValid</code> | This method returns <code>true</code> if the value stored in the <code>second</code> field is valid. Otherwise, it returns <code>false</code> .                                                                               |
| <code>isThirdGroupValid</code>  | This method returns <code>true</code> if the value stored in the <code>third</code> field is valid. Otherwise, it returns <code>false</code> .                                                                                |

The code for the `SerialNumber` class is shown in [Code Listing CS6-1](#).

## Code Listing CS6-1 (`SerialNumber.java`)

```

1 /**
2 The SerialNumber class takes a software serial number in
3 the form of LLLLL-DDDD-LLLL where each L is a letter
4 and each D is a digit. The serial number has three groups
5 of characters, separated by hyphens. The class extracts
6 the three groups of characters and validates them.
7 */
8
9 public class SerialNumber
10 {
11 private String first; // First group of characters
12 private String second; // Second group of characters
13 private String third; // Third group of characters
14 private boolean valid; // Flag indicating validity
15
16 /**
17 The constructor breaks a serial number into
18 three groups and each group is validated.
19 @param sn A serial number.
20 */

```

```

21
22 public SerialNumber(String sn)
23 {
24 // Trim the argument.
25 sn = sn.trim();
26
27 // Tokenize the argument.
28 String tokens[] = sn.split("-");
29
30 // Validate the tokens.
31 if (tokens.length != 3)
32 valid = false;
33 else
34 {
35 first = tokens[0];
36 second = tokens[1];
37 third = tokens[2];
38 validate();
39 }
40
41 }
42
43 /**
44 * The isValid method returns a boolean value
45 * indicating whether the serial number is valid.
46 * @return true if the serial number is valid,
47 * otherwise false.
48 */
49
50 public boolean isValid()
51 {
52 return valid;
53 }
54
55 /**
56 * The validate method sets the valid field to true
57 * if the serial number is valid. Otherwise it sets
58 * valid to false.
59 */
60
61 private void validate()
62 {
63 if (isFirstGroupValid() && isSecondGroupValid() &&
64 isThirdGroupValid())
65 valid = true;
66 else
67 valid = false;
68 }
69
70 /**
71 * The isFirstGroupValid method validates the first
72 * group of characters.
73 * @return true if the group is valid, otherwise false.
74 */
75
76 private boolean isFirstGroupValid()
77 {

```

```

78 boolean groupGood = true; // Flag
79 int i = 0; // Loop control variable
80
81 // Check the length of the group.
82 if (first.length() != 5)
83 groupGood = false;
84
85 // See if each character is a letter.
86 while (groupGood && i < first.length())
87 {
88 if (!Character.isLetter(first.charAt(i)))
89 groupGood = false;
90 i++;
91 }
92
93 return groupGood;
94 }
95
96 /**
97 The isSecondGroupValid method validates the second
98 group of characters.
99 @return true if the group is valid, otherwise false.
100 */
101
102 private boolean isSecondGroupValid()
103 {
104 boolean groupGood = true; // Flag
105 int i = 0; // Loop control variable
106
107 // Check the length of the group.
108 if (second.length() != 4)
109 groupGood = false;
110
111 // See if each character is a digit.
112 while (groupGood && i < second.length())
113 {
114 if (!Character.isDigit(second.charAt(i)))
115 groupGood = false;
116 i++;
117 }
118
119 return groupGood;
120 }
121
122 /**
123 The isThirdGroupValid method validates the third
124 group of characters.
125 @return true if the group is valid, otherwise false.
126 */
127
128 private boolean isThirdGroupValid()
129 {
130 boolean groupGood = true; // Flag
131 int i = 0; // Loop control variable
132
133 // Check the length of the group.
134 if (third.length() != 4)
135 groupGood = false;
136
137 // See if each character is a digit.
138 while (groupGood && i < third.length())
139 {
140 if (!Character.isDigit(third.charAt(i)))
141 groupGood = false;
142 i++;
143 }
144
145 return groupGood;
146 }

```

```

135 // Check the length of the group.
136 if (third.length() != 4)
137 groupGood = false;
138
139 // See if each character is a letter.
140 while (groupGood && i < third.length())
141 {
142 if (!Character.isLetter(third.charAt(i)))
143 groupGood = false;
144 i++;
145 }
146
147 return groupGood;
148 }
149 }

```

Let's take a closer look at the constructor.

Notice that in line 25 we call the argument's `trim` method to remove any leading and/or trailing whitespace characters. This is important because whitespace characters are not used as delimiters in this code. If the argument contains leading whitespace characters, they will be included as part of the first token. Trailing whitespace characters will be included as part of the last token. Then, in line 27 we call the argument's `split` method to tokenize the string, using the “-” character as a delimiter. The tokens are returned to an array named `tokens`.

Next, the `if` statement in lines 31 through 39 executes:

```

31 if (tokens.length != 3)
32 valid = false;
33 else
34 {
35 first = tokens[0];
36 second = tokens[1];
37 third = tokens[2];
38 validate();
39 }

```

A valid serial number must have three groups of characters, so the `if` statement determines whether the `tokens` array has three elements. If not, the `valid` field is set to `false`. Otherwise, the three tokens are assigned to the `first`, `second`, and `third` fields. Last, the `validate` method is called. The `validate` method calls the `isFirstGroupValid`, `isSecondGroupValid`, and `isThirdGroupValid` methods to validate the three groups of characters. In the end, the `valid` field will be set to `true` if the serial number is valid, or `false` otherwise. The program in [Code Listing CS6-2](#) demonstrates the class.

## Code Listing CS6-2

### (SerialNumberTester.java)

```

2 This program demonstrates the SerialNumber class.
3 */
4
5 public class SerialNumberTester
6 {
7 public static void main(String[] args)
8 {
9 String serial1 = "GHTRJ-8975-AQWR"; // Valid
10 String serial2 = "GHT7J-8975-AQWR"; // Invalid
11 String serial3 = "GHTRJ-8J75-AQWR"; // Invalid
12 String serial4 = "GHTRJ-8975-AQ2R"; // Invalid
13
14 // Validate serial1.
15
16 SerialNumber sn = new SerialNumber(serial1);
17 if (sn.isValid())
18 System.out.println(serial1 + " is valid.");
19 else
20 System.out.println(serial1 + " is invalid.");
21
22 // Validate serial2.
23
24 sn = new SerialNumber(serial2);
25 if (sn.isValid())
26 System.out.println(serial2 + " is valid.");
27 else
28 System.out.println(serial2 + " is invalid.");
29
30 // Validate serial3.
31
32 sn = new SerialNumber(serial3);
33 if (sn.isValid())
34 System.out.println(serial3 + " is valid.");
35 else
36 System.out.println(serial3 + " is invalid.");
37
38 // Validate serial4.
39
40 sn = new SerialNumber(serial4);
41 if (sn.isValid())
42 System.out.println(serial4 + " is valid.");
43 else
44 System.out.println(serial4 + " is invalid.");
45 }
46 }

```

# Program Output

```

GHTRJ-8975-AQWR is valid.
GHT7J-8975-AQWR is invalid.
GHTRJ-8J75-AQWR is invalid.
GHTRJ-8975-AQ2R is invalid.

```



# CASE STUDY 7 A Simple Text Editor

## Application\*

\*From *Starting Out with Java: From Control Structures through Objects*, Sixth Edition. Tony Gaddis. Copyright © 2016 by Pearson Education, Inc. All rights reserved.

To demonstrate the `JTextArea` component, fonts, menus, and file choosers, we present a simple text editor application. This application allows you to create new text files and open existing text files. The file contents are displayed in a text area. You can also change the font and style of the text that is displayed in the text area. The application's window is shown in [Figure CS7-1](#). Line wrapping is turned on, using word wrap style, and the text area is in a scroll pane.





**Figure CS7-1 The text editor window**

The application uses a menu system to perform these operations, which is shown in [Figure CS7-2](#).



## Figure CS7-2 Menu system for the text editor application

[Figure CS7-2 Full Alternative Text](#)

Each menu item generates an action event that is handled by an action listener. The following section presents a summary of the actions performed by each menu item.

# File Menu

**New** This menu item clears any text that is stored in the text area. In addition, the class's `filename` field is set to `null`. The `filename` field contains the path and name of the file that is currently displayed in the text area.

**Open** This menu item displays a file chooser that allows the user to select a file to open. If the user selects a file, it is opened and its contents are read into the text area. The path and name of the file are stored in the `filename` field.

**Save** This menu item saves the contents of the text area. The contents are saved to the file with the name and path stored in the `filename` field. If the `filename` field is set to `null`, which would indicate that the file has not been saved yet, then this menu item performs the same action as the Save As menu item.

**Save As** This menu item displays a file chooser that allows the user to select a location and file name. The contents of the text area are written to the selected file, and the path and file name are stored in the `filename` field. (Be careful when using this menu item. As it is currently written, this application does not warn you when you are about to overwrite an existing file!)

**Exit** This menu item ends the application.

# Font Menu

**Monospaced** This radio button menu item changes the text area's font to Monospaced.

**Serif** This radio button menu item changes the text area's font to Serif.


**SansSerif** This radio button menu item changes the text area's font to SansSerif.

**Italic** This check box menu item changes the text area's style to italic.


**Bold** This check box menu item changes the text area's style to bold.

[Code Listing CS7-1](#) shows the code for the `TextEditor` class. The `main` method creates an instance of the class, which displays the text editor window. [Figure CS7-3](#) shows the window displaying text in various fonts and styles.

SansSerif font, plain style



Monospaced font, bold style



Serif font, bold and italic style




**Figure CS7-3 Text displayed in various fonts and styles**

[Figure CS7-3 Full Alternative Text](#)

## Code Listing CS7-1 (TextEditor.java)

```

1 import java.awt.*;
2 import java.awt.event.*;
3 import javax.swing.*;
4 import java.io.*;
5 import java.util.Scanner;
6
7 /**

```

```

8 The TextEditor class is a simple text editor.
9 */
10
11 public class TextEditor extends JFrame
12 {
13 // The following are fields for the menu system.
14 // First, the menu bar
15 private JMenuBar menuBar;
16
17 // The menus
18 private JMenu fileMenu;
19 private JMenu fontMenu;
20
21 // The menu items
22 private JMenuItem newItem;
23 private JMenuItem openItem;
24 private JMenuItem saveItem;
25 private JMenuItem saveAsItem;
26 private JMenuItem exitItem;
27
28 // The radio button menu items
29 private JRadioButtonMenuItem monoItem;
30 private JRadioButtonMenuItem serifItem;
31 private JRadioButtonMenuItem sansSerifItem;
32
33 // The checkbox menu items
34 private JCheckBoxMenuItem italicItem;
35 private JCheckBoxMenuItem boldItem;
36
37 private String filename; // To hold the file name
38 private JTextArea editorText; // To display the text
39 private final int NUM_LINES = 20; // Lines to display
40 private final int NUM_CHARS = 40; // Chars per line
41
42 /**
43 * Constructor
44 */
45
46 public TextEditor()
47 {
48 // Set the title.
49 setTitle("Text Editor");
50
51 // Specify what happens when the close
52 // button is clicked.
53 setDefaultCloseOperation(JFrame.EXIT_ON_CLOSE);
54
55 // Create the text area.
56 editorText = new JTextArea(NUM_LINES, NUM_CHARS);
57
58 // Turn line wrapping on.
59 editorText.setLineWrap(true);
60 editorText.setWrapStyleWord(true);
61
62 // Create a scroll pane and add the text area to it.
63 JScrollPane scrollPane = new JScrollPane(editorText);
64

```

```

65 // Add the scroll pane to the content pane.
66 add(scrollPane);
67
68 // Build the menu bar.
69 buildMenuBar();
70
71 // Pack and display the window.
72 pack();
73 setVisible(true);
74 }
75
76 /**
77 * The buildMenuBar method creates a menu bar and
78 * calls the createFileMenu method to create the
79 * file menu.
80 */
81
82 private void buildMenuBar()
83 {
84 // Build the file and font menus.
85 buildFileMenu();
86 buildFontMenu();
87
88 // Create the menu bar.
89 menuBar = new JMenuBar();
90
91 // Add the file and font menus to the menu bar.
92 menuBar.add(fileMenu);
93 menuBar.add(fontMenu);
94
95 // Set the menu bar for this frame.
96 setJMenuBar(menuBar);
97 }
98
99 /**
100 * The buildFileMenu method creates the file menu
101 * and populates it with its menu items.
102 */
103
104 private void buildFileMenu()
105 {
106 // Create the New menu item.
107 newItem = new JMenuItem("New");
108 newItem.setMnemonic(KeyEvent.VK_N);
109 newItem.addActionListener(new NewListener());
110
111 // Create the Open menu item.
112 openItem = new JMenuItem("Open");
113 openItem.setMnemonic(KeyEvent.VK_O);
114 openItem.addActionListener(new OpenListener());
115
116 // Create the Save menu item.
117 saveItem = new JMenuItem("Save");
118 saveItem.setMnemonic(KeyEvent.VK_S);
119 saveItem.addActionListener(new SaveListener());
120
121 // Create the Save As menu item.

```

```

122 saveAsItem = new JMenuItem("Save As");
123 saveAsItem.setMnemonic(KeyEvent.VK_A);
124 saveAsItem.addActionListener(new SaveListener());
125
126 // Create the Exit menu item.
127 exitItem = new JMenuItem("Exit");
128 exitItem.setMnemonic(KeyEvent.VK_X);
129 exitItem.addActionListener(new ExitListener());
130
131 // Create a menu for the items we just created.
132 fileMenu = new JMenu("File");
133 fileMenu.setMnemonic(KeyEvent.VK_F);
134
135 // Add the items and some separator bars to the menu.
136 fileMenu.add(newItem);
137 fileMenu.add(openItem);
138 fileMenu.addSeparator();// Separator bar
139 fileMenu.add(saveItem);
140 fileMenu.add(saveAsItem);
141 fileMenu.addSeparator();// Separator bar
142 fileMenu.add(exitItem);
143 }
144
145 /**
146 The buildFontMenu method creates the font menu
147 and populates it with its menu items.
148 */
149
150 private void buildFontMenu()
151 {
152 // Create the Monospaced menu item.
153 monoItem = new JRadioButtonMenuItem("Monospaced");
154 monoItem.addActionListener(new FontListener());
155
156 // Create the Serif menu item.
157 serifItem = new JRadioButtonMenuItem("Serif");
158 serifItem.addActionListener(new FontListener());
159
160 // Create the SansSerif menu item.
161 sansSerifItem =
162 new JRadioButtonMenuItem("SansSerif", true);
163 sansSerifItem.addActionListener(new FontListener());
164
165 // Group the radio button menu items.
166 ButtonGroup group = new ButtonGroup();
167 group.add(monoItem);
168 group.add(serifItem);
169 group.add(sansSerifItem);
170
171 // Create the Italic menu item.
172 italicItem = new JCheckBoxMenuItem("Italic");
173 italicItem.addActionListener(new FontListener());
174
175 // Create the Bold menu item.
176 boldItem = new JCheckBoxMenuItem("Bold");
177 boldItem.addActionListener(new FontListener());
178

```



```

179 // Create a menu for the items we just created.
180 fontMenu = new JMenu("Font");
181 fontMenu.setMnemonic(KeyEvent.VK_T);
182
183 // Add the items and some separator bars to the menu.
184 fontMenu.add(monoItem);
185 fontMenu.add(serifItem);
186 fontMenu.add(sansSerifItem);
187 fontMenu.addSeparator();// Separator bar
188 fontMenu.add(italicItem);
189 fontMenu.add(boldItem);
190 }
191
192 /**
193 Private inner class that handles the event that
194 is generated when the user selects New from
195 the file menu.
196 */
197
198 private class NewListener implements ActionListener
199 {
200 public void actionPerformed(ActionEvent e)
201 {
202 editorText.setText("");
203 filename = null;
204 }
205 }
206
207 /**
208 Private inner class that handles the event that
209 is generated when the user selects Open from
210 the file menu.
211 */
212
213 private class OpenListener implements ActionListener
214 {
215 public void actionPerformed(ActionEvent e)
216 {
217 int chooserStatus;
218
219 JFileChooser chooser = new JFileChooser();
220 chooserStatus = chooser.showOpenDialog(null);
221 if (chooserStatus == JFileChooser.APPROVE_OPTION)
222 {
223 // Get a reference to the selected file.
224 File selectedFile = chooser.getSelectedFile();
225
226 // Get the path of the selected file.
227 filename = selectedFile.getPath();
228
229 // Open the file.
230 if (!openFile(filename))
231 {
232 JOptionPane.showMessageDialog(null,
233 "Error reading " +
234 filename, "Error",
235 JOptionPane.ERROR_MESSAGE);

```

```

236 }
237 }
238 }
239
240 /**
241 The openFile method opens the file specified by
242 filename and reads its contents into the text
243 area. The method returns true if the file was
244 opened and read successfully, or false if an
245 error occurred.
246 @param filename The name of the file to open.
247 */
248
249 private boolean openFile(String filename)
250 {
251 boolean success;
252 String inputLine, editorString = "";
253
254 try
255 {
256 // Open the file.
257 File file = new File(filename);
258 Scanner inputFile = new Scanner(file);
259
260 // Read the file contents into the editor.
261 while (inputFile.hasNext())
262 {
263 // Read a line from the file.
264 inputLine = inputFile.nextLine();
265
266 // Append it to the string to display
267 // in the editor.
268 editorString = editorString +
269 inputLine + "\n";
270 }
271
272 // Display the string that was read from the
273 // file in the editor.
274 editorText.setText(editorString);
275
276 // Close the file.
277 inputFile.close();
278
279 // Indicate that everything went OK.
280 success = true;
281 }
282 catch (IOException e)
283 {
284 // Something went wrong.
285 success = false;
286 }
287
288 // Return our status.
289 return success;
290 }
291 }
292

```

```

293 /**
294 Private inner class that handles the event that
295 is generated when the user selects Save or Save
296 As from the file menu.
297 */
298
299 private class SaveListener implements ActionListener
300 {
301 public void actionPerformed(ActionEvent e)
302 {
303 int chooserStatus;
304
305 // If the user selected Save As, or the contents
306 // of the editor have not been saved, use a file
307 // chooser to get the file name. Otherwise, save
308 // the file under the current file name.
309
310 if (e.getActionCommand() == "Save As" ||
311 filename == null)
312 {
313 JFileChooser chooser = new JFileChooser();
314 chooserStatus = chooser.showSaveDialog(null);
315 if (chooserStatus == JFileChooser.APPROVE_OPTION)
316 {
317 // Get a reference to the selected file.
318 File selectedFile =
319 chooser.getSelectedFile();
320
321 // Get the path of the selected file.
322 filename = selectedFile.getPath();
323 }
324 }
325
326 // Save the file.
327 if (!saveFile(filename))
328 {
329 JOptionPane.showMessageDialog(null,
330 "Error saving " +
331 filename,
332 "Error",
333 JOptionPane.ERROR_MESSAGE);
334 }
335 }
336
337 /**
338 The saveFile method saves the contents of the
339 text area to a file. The method returns true if
340 the file was saved successfully, or false if an
341 error occurred.
342 @param filename The name of the file.
343 @return true if successful, false otherwise.
344 */
345
346 private boolean saveFile(String filename)
347 {
348 boolean success;
349 String editorString;

```

```

350 PrintWriter outputFile;
351
352 try
353 {
354 // Open the file.
355 outputFile = new PrintWriter(filename);
356
357 // Write the contents of the text area
358 // to the file.
359 editorString = editorText.getText();
360 outputFile.print(editorString);
361
362 // Close the file.
363 outputFile.close();
364
365 // Indicate that everything went OK.
366 success = true;
367 }
368 catch (IOException e)
369 {
370 // Something went wrong.
371 success = false;
372 }
373
374 // Return our status.
375 return success;
376 }
377 }
378
379 /**
380 Private inner class that handles the event that
381 is generated when the user selects Exit from
382 the file menu.
383 */
384
385 private class ExitListener implements ActionListener
386 {
387 public void actionPerformed(ActionEvent e)
388 {
389 System.exit(0);
390 }
391 }
392
393 /**
394 Private inner class that handles the event that
395 is generated when the user selects an item from
396 the font menu.
397 */
398
399 private class FontListener implements ActionListener
400 {
401 public void actionPerformed(ActionEvent e)
402 {
403 // Get the current font.
404 Font textFont = editorText.getFont();
405
406 // Retrieve the font name and size.

```

```

407 String fontName = textFont.getName();
408 int fontSize = textFont.getSize();
409
410 // Start with plain style.
411 int fontStyle = Font.PLAIN;
412
413 // Determine which font is selected.
414 if (monoItem.isSelected())
415 fontName = "Monospaced";
416 else if (serifItem.isSelected())
417 fontName = "Serif";
418 else if (sansSerifItem.isSelected())
419 fontName = "SansSerif";
420
421 // Determine whether italic is selected.
422 if (italicItem.isSelected())
423 fontStyle += Font.ITALIC;
424
425 // Determine whether bold is selected.
426 if (boldItem.isSelected())
427 fontStyle += Font.BOLD;
428
429 // Set the font as selected.
430 editorText.setFont(new Font(fontName,
431 fontStyle, fontSize));
432 }
433 }
434
435 /**
436 * main method
437 */
438
439 public static void main(String[] args)
440 {
441 TextEditor te = new TextEditor();
442 }
443 }

```

# Contents

1. [STARTING OUT WITH JAVA™ From Control Structures through Data Structures](#)
2. [Contents in Brief](#)
3. [Contents](#)
4. [Preface](#)
  1. [Control Structures First, Then Objects, Then Data Structures](#)
  2. [Changes in This Edition](#)
  3. [Organization of the Text](#)
  4. [Brief Overview of Each Chapter](#)
  5. [Features of the Text](#)
  6. [Supplements](#)
  7. [Reviewers for Previous Editions](#)
  8. [About the Authors](#)
5. [Chapter 1 Introduction to Computers and Java](#)
  1. [TOPICS](#)
  2. [1.1 Introduction](#)
  3. [1.2 Why Program?](#)
  4. [1.3 Computer Systems: Hardware and Software](#)
    1. [Hardware](#)
      1. [The CPU](#)
      2. [Main Memory](#)
      3. [Secondary Storage](#)
      4. [Input Devices](#)
      5. [Output Devices](#)
    2. [Software](#)
      1. [Checkpoint](#)
  5. [1.4 Programming Languages](#)
    1. [What Is a Program?](#)
    2. [A History of Java](#)
  6. [1.5 What Is a Program Made Of?](#)
    1. [Language Elements](#)
      1. [Key Words \(Reserved Words\)](#)
      2. [Programmer-Defined Names](#)
      3. [Operators](#)
      4. [Punctuation](#)
    2. [Lines and Statements](#)
    3. [Variables](#)
    4. [The Compiler and the Java Virtual Machine](#)
      1. [Portability](#)
    5. [Java Software Editions](#)
    6. [Compiling and Running a Java Program](#)

1. [Integrated Development Environments](#)
  1. [Checkpoint](#)
7. [1.6 The Programming Process](#)
  1. [Software Engineering](#)
    1. [Checkpoint](#)
8. [1.7 Object-Oriented Programming](#)
  1. [Checkpoint](#)
9. [Review Questions and Exercises](#)
  1. [Multiple Choice](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Predict the Result](#)
  5. [Short Answer](#)
10. [Programming Challenge](#)
  1. [1. Your First Java Program](#)
  2. [If You Are Using the JDK at the Command Prompt:](#)
  3. [If You Are Using an IDE:](#)
6. [Chapter 2 Java Fundamentals](#)
  1. [TOPICS](#)
  2. [2.1 The Parts of a Java Program](#)
    1. [Checkpoint](#)
  3. [2.2 The `print` and `println` Methods, and the Java API](#)
  4. [2.3 Variables and Literals](#)
    1. [Displaying Multiple Items with the `+` Operator](#)
    2. [Be Careful with Quotation Marks](#)
    3. [More about Literals](#)
    4. [Identifiers](#)
    5. [Class Names](#)
      1. [Checkpoint](#)
  5. [2.4 Primitive Data Types](#)
    1. [The Integer Data Types](#)
      1. [Integer Literals](#)
    2. [Floating-Point Data Types](#)
      1. [Floating-Point Literals](#)
      2. [Scientific and E Notation](#)
    3. [The `boolean` Data Type](#)
    4. [The `char` Data Type](#)
      1. [Unicode](#)
    5. [Variable Assignment and Initialization](#)
    6. [Variables Hold Only One Value at a Time](#)
      1. [Checkpoint](#)
  6. [2.5 Arithmetic Operators](#)
    1. [Integer Division](#)
    2. [Operator Precedence](#)
    3. [Grouping with Parentheses](#)

4. [The Math Class](#)
  1. [The Math.pow Method](#)
  2. [The Math.sqrt Method](#)
    1. [Checkpoint](#)
7. [2.6 Combined Assignment Operators](#)
  1. [Checkpoint](#)
8. [2.7 Conversion between Primitive Data Types](#)
  1. [Mixed Integer Operations](#)
  2. [Other Mixed Mathematical Expressions](#)
    1. [Checkpoint](#)
9. [2.8 Creating Named Constants with final](#)
10. [2.9 The String Class](#)
  1. [Objects Are Created from Classes](#)
  2. [The String Class](#)
  3. [Primitive Type Variables and Class Type Variables](#)
  4. [Creating a String Object](#)
    1. [Checkpoint](#)
11. [2.10 Scope](#)
12. [2.11 Comments](#)
13. [2.12 Programming Style](#)
14. [2.13 Reading Keyboard Input](#)
  1. [Reading a Character](#)
  2. [Mixing Calls to nextLine with Calls to Other Scanner Methods](#)
15. [2.14 Dialog Boxes](#)
  1. [Displaying Message Dialogs](#)
  2. [Displaying Input Dialogs](#)
  3. [An Example Program](#)
  4. [Converting String Input to Numbers](#)
    1. [Checkpoint](#)
16. [2.15 Common Errors to Avoid](#)
17. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Predict the Output](#)
  3. [Find the Error](#)
  4. [Algorithm Workbench](#)
  5. [Short Answer](#)
18. [Programming Challenges](#)
  1. [1. Name, Age, and Annual Income](#)
  2. [2. Name and Initials](#)
  3. [3. Personal Information](#)
  4. [4. Star Pattern](#)
  5. [5. Sales Prediction](#)
  6. [6. Land Calculation](#)
  7. [7. Sales Tax](#)
  8. [8. Cookie Calories](#)



9. [9. Miles-per-Gallon](#)
10. [10. Test Average](#)
11. [11. Circuit Board Profit](#)
12. [12. String Manipulator](#)
13. [13. Restaurant Bill](#)
14. [14. Male and Female Percentages](#)
15. [15. Stock Commission](#)
16. [16. Energy Drink Consumption](#)
17. [17. Ingredient Adjuster](#)
18. [18. Word Game](#)
19. [19. Stock Transaction Program](#)
20. [20. Planting Grapevines](#)
21. [21. Compound Interest](#)

## 7. [Chapter 3 Decision Structures](#)

1. [TOPICS](#)
2. [3.1 The `if` Statement](#)
  1. [Using Relational Operators to Form Conditions](#)
  2. [Putting It All Together](#)
  3. [Programming Style and the `if` Statement](#)
  4. [Be Careful with Semicolons](#)
  5. [Having Multiple Conditionally Executed Statements](#)
  6. [Flags](#)
  7. [Comparing Characters](#)
    1. [Checkpoint](#)
3. [3.2 The `if-else` Statement](#)
  1. [Checkpoint](#)
4. [3.3 Nested `if` Statements](#)
5. [3.4 The `if-else-if` Statement](#)
6. [3.5 Logical Operators](#)
  1. [The Precedence of Logical Operators](#)
  2. [Checking Numeric Ranges with Logical Operators](#)
    1. [Checkpoint](#)
7. [3.6 Comparing String Objects](#)
  1. [Ignoring Case in String Comparisons](#)
    1. [Checkpoint](#)
8. [3.7 More about Variable Declaration and Scope](#)
9. [3.8 The Conditional Operator \(Optional\)](#)
  1. [Checkpoint](#)
10. [3.9 The `switch` Statement](#)
  1. [Checkpoint](#)
11. [3.10 Displaying Formatted Output with `System.out.printf` and `String.format`](#)
  1. [Format Specifier Syntax](#)
  2. [Precision](#)
  3. [Specifying a Minimum Field Width](#)
    1. [Combining Minimum Field Width and Precision in the Same Format Specifier](#)

4. [Flags](#)
  1. [Comma Separators](#)
  2. [Padding Numbers with Leading Zeros](#)
  3. [Left-Justifying Numbers](#)
5. [Formatting String Arguments](#)
6. [The `String.format` Method](#)
  1. [Checkpoint](#)
12. [3.11 Common Errors to Avoid](#)
13. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
14. [Programming Challenges](#)
  1. [1. Roman Numerals](#)
  2. [2. Magic Dates](#)
  3. [3. Body Mass Index](#)
  4. [4. Test Scores and Grade](#)
  5. [5. Mass and Weight](#)
  6. [6. Time Calculator](#)
  7. [7. Sorted Names](#)
  8. [8. Software Sales](#)
  9. [9. Shipping Charges](#)
  10. [10. Fat Gram Calculator](#)
  11. [11. Running the Race](#)
  12. [12. The Speed of Sound](#)
  13. [13. Mobile Service Provider](#)
  14. [14. Mobile Service Provider, Part 2](#)
  15. [15. Bank Charges](#)
  16. [16. Book Club Points](#)
  17. [17. Wi-Fi Diagnostic Tree](#)
  18. [18. Restaurant Selector](#)
8. [Chapter 4 Loops and Files](#)
  1. [TOPICS](#)
  2. [4.1 The Increment and Decrement Operators](#)
    1. [The Difference between Postfix and Prefix Modes](#)
      1. [Checkpoint](#)
  3. [4.2 The `while` Loop](#)
    1. [The `while` Loop Is a Pretest Loop](#)
    2. [Infinite Loops](#)
    3. [Don't Forget the Braces with a Block of Statements](#)
    4. [Programming Style and the `while` Loop](#)
      1. [Checkpoint](#)
  4. [4.3 Using the `while` Loop for Input Validation](#)
    1. [Checkpoint](#)

5. [4.4 The do-while Loop](#)
6. [4.5 The for Loop](#)
  1. [The for Loop Is a Pretest Loop](#)
  2. [Avoid Modifying the Control Variable in the Body of the for Loop](#)
  3. [Other Forms of the Update Expression](#)
  4. [Declaring a Variable in the for Loop's Initialization Expression](#)
  5. [Creating a User Controlled for Loop](#)
  6. [Using Multiple Statements in the Initialization and Update Expressions](#)
    1. [Checkpoint](#)
7. [4.6 Running Totals and Sentinel Values](#)
  1. [Using a Sentinel Value](#)
    1. [Checkpoint](#)
8. [4.7 Nested Loops](#)
9. [4.8 The break and continue Statements \(Optional\)](#)
10. [4.9 Deciding Which Loop to Use](#)
11. [4.10 Introduction to File Input and Output](#)
  1. [Using the PrintWriter Class to Write Data to a File](#)
    1. [More about the PrintWriter Class's println Method](#)
    2. [The PrintWriter Class's print Method](#)
    3. [Adding a throws Clause to the Method Header](#)
    4. [An Example Program](#)
    5. [Review](#)
  2. [Appending Data to a File](#)
  3. [Specifying the File Location](#)
  4. [Reading Data from a File](#)
  5. [Reading Lines from a File with the nextLine Method](#)
  6. [Adding a throws Clause to the Method Header](#)
    1. [Detecting the End of a File](#)
    2. [Reading Primitive Values from a File](#)
    3. [Review](#)
  7. [Checking for a File's Existence](#)
    1. [Checkpoint](#)
12. [4.11 Generating Random Numbers with the Random Class](#)
  1. [Checkpoint](#)
13. [4.12 Common Errors to Avoid](#)
14. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
15. [Programming Challenges](#)
  1. [1. Sum of Numbers](#)
  2. [2. Distance Traveled](#)
  3. [3. Distance File](#)
  4. [4. Pennies for Pay](#)

5. [5. Letter Counter](#)
6. [6. File Letter Counter](#)
7. [7. Hotel Occupancy](#)
8. [8. Average Rainfall](#)
9. [9. Population](#)
10. [10. Largest and Smallest](#)
11. [11. Celsius to Fahrenheit Table](#)
12. [12. Bar Chart](#)
13. [13. File Head Display](#)
14. [14. Line Numbers](#)
15. [15. Uppercase File Converter](#)
16. [16. Budget Analysis](#)
17. [17. Random Number Guessing Game](#)
18. [18. Random Number Guessing Game Enhancement](#)
19. [19. ESP Game](#)
20. [20. Square Display](#)
21. [21. Dice Game](#)
22. [22. Slot Machine Simulation](#)
23. [23. Personal Web Page Generator](#)

## 9. [Chapter 5 Methods](#)

1. [TOPICS](#)
2. [5.1 Introduction to Methods](#)
  1. [void Methods and Value-Returning Methods](#)
  2. [Defining a void Method](#)
  3. [Calling a Method](#)
  4. [Layered Method Calls](#)
  5. [Using Documentation Comments with Methods](#)
    1. [Checkpoint](#)
3. [5.2 Passing Arguments to a Method](#)
  1. [Argument and Parameter Data Type Compatibility](#)
  2. [Parameter Variable Scope](#)
  3. [Passing Multiple Arguments](#)
  4. [Arguments Are Passed by Value](#)
  5. [Passing Object References to a Method](#)
  6. [Using the @param Tag in Documentation Comments](#)
    1. [Checkpoint](#)
4. [5.3 More about Local Variables](#)
  1. [Local Variable Lifetime](#)
  2. [Initializing Local Variables with Parameter Values](#)
5. [5.4 Returning a Value from a Method](#)
  1. [Defining a Value-Returning Method](#)
  2. [Calling a Value-Returning Method](#)
  3. [Using the @return Tag in Documentation Comments](#)
  4. [Returning a boolean Value](#)
  5. [Returning a Reference to an Object](#)

1. [Checkpoint](#)

6. [5.5 Problem Solving with Methods](#)

1. [Calling Methods That Throw Exceptions](#)

7. [5.6 Common Errors to Avoid](#)

8. [Review Questions and Exercises](#)

1. [Multiple Choice and True/False](#)

2. [Find the Error](#)

3. [Algorithm Workbench](#)

4. [Short Answer](#)

9. [Programming Challenges](#)

1. [1. showChar Method](#)

2. [2. Retail Price Calculator](#)

3. [3. Rectangle Area—Complete the Program](#)

4. [4. Paint Job Estimator](#)

5. [5. Falling Distance](#)

6. [6. Celsius Temperature Table](#)

7. [7. Test Average and Grade](#)

8. [8. Conversion Program](#)

9. [9. Distance Traveled Modification](#)

10. [10. Stock Profit](#)

11. [11. Multiple Stock Sales](#)

12. [12. Kinetic Energy](#)

13. [13. isPrime Method](#)

14. [14. Prime Number List](#)

15. [15. Even/Odd Counter](#)

16. [16. Present Value](#)

17. [17. Rock, Paper, Scissors Game](#)

18. [18. ESP Game](#)

10. [Chapter 6 A First Look at Classes](#)

1. [TOPICS](#)

2. [6.1 Objects and Classes](#)

1. [Classes: Where Objects Come From](#)

2. [Classes in the Java API](#)

3. [Primitive Variables vs. Objects](#)

1. [Checkpoint](#)

3. [6.2 Writing a Simple Class, Step by Step](#)

1. [Accessor and Mutator Methods](#)

2. [The Importance of Data Hiding](#)

3. [Avoiding Stale Data](#)

4. [Showing Access Specification in UML Diagrams](#)

5. [Data Type and Parameter Notation in UML Diagrams](#)

6. [Layout of Class Members](#)

1. [Checkpoint](#)

4. [6.3 Instance Fields and Methods](#)

1. [Checkpoint](#)

5. [6.4 Constructors](#)
  1. [Showing Constructors in a UML Diagram](#)
  2. [Uninitialized Local Reference Variables](#)
  3. [The Default Constructor](#)
  4. [Writing Your Own No-Arg Constructor](#)
  5. [The `String` Class Constructor](#)
    1. [Checkpoint](#)
6. [6.5 Passing Objects as Arguments](#)
7. [6.6 Overloading Methods and Constructors](#)
  1. [The `BankAccount` Class](#)
  2. [Overloaded Methods Make Classes More Useful](#)
8. [6.7 Scope of Instance Fields](#)
  1. [Shadowing](#)
    1. [Checkpoint](#)
9. [6.8 Packages and `import` Statements](#)
  1. [Explicit and Wildcard `import` Statements](#)
  2. [The `java.lang` Package](#)
  3. [Other API Packages](#)
10. [6.9 Focus on Object-Oriented Design: Finding the Classes and Their Responsibilities](#)
  1. [Finding the Classes](#)
    1. [Writing a Description of the Problem Domain](#)
    2. [Identifying All of the Nouns](#)
    3. [Refining the List of Nouns](#)
    4. [1. Some of the nouns really mean the same thing.](#)
    5. [2. Some nouns might represent items that we do not need to be concerned with in order to solve the problem.](#)
    6. [3. Some of the nouns might represent objects, not classes.](#)
    7. [4. Some of the nouns might represent simple values that can be stored in a primitive variable and do not require a class.](#)
  2. [Identifying a Class's Responsibilities](#)
    1. [The `Customer` class](#)
    2. [The `Car` Class](#)
    3. [The `ServiceQuote` Class](#)
  3. [This Is Only the Beginning](#)
    1. [Checkpoint](#)
11. [6.10 Common Errors to Avoid](#)
12. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
13. [Programming Challenges](#)
  1. [1. `Employee` Class](#)
  2. [2. `Car` Class](#)
  3. [3. Personal Information Class](#)

4. [4. RetailItem Class](#)
5. [5. Payroll Class](#)
6. [6. TestScores Class](#)
7. [7. Circle Class](#)
8. [8. Temperature Class](#)
9. [9. Days in a Month](#)
10. [10. A Game of Twenty-One](#)
11. [11. Freezing and Boiling Points](#)
12. [12. SavingsAccount Class](#)
13. [13. Deposit and Withdrawal Files](#)
14. [14. Patient Charges](#)
15. [15. Dice Game](#)
16. [16. Roulette Wheel Colors](#)
17. [17. Coin Toss Simulator](#)
18. [18. Tossing Coins for a Dollar](#)
19. [19. Fishing Game Simulation](#)

## 11. [Chapter 7 Arrays and the ArrayList Class](#)

1. [TOPICS](#)
2. [7.1 Introduction to Arrays](#)
  1. [Accessing Array Elements](#)
  2. [Inputting and Outputting Array Contents](#)
  3. [Java Performs Bounds Checking](#)
  4. [Watch Out for Off-by-One Errors](#)
  5. [Array Initialization](#)
  6. [Alternate Array Declaration Notation](#)
    1. [Checkpoint](#)
3. [7.2 Processing Array Elements](#)
  1. [Array Length](#)
  2. [The Enhanced for Loop](#)
    1. [The Enhanced for Loop versus the Traditional for Loop](#)
  3. [Letting the User Specify an Array's Size](#)
  4. [Reassigning Array Reference Variables](#)
  5. [Copying Arrays](#)
    1. [Checkpoint](#)
4. [7.3 Passing Arrays as Arguments to Methods](#)
  1. [Checkpoint](#)
5. [7.4 Some Useful Array Algorithms and Operations](#)
  1. [Comparing Arrays](#)
  2. [Summing the Values in a Numeric Array](#)
  3. [Getting the Average of the Values in a Numeric Array](#)
  4. [Finding the Highest and Lowest Values in a Numeric Array](#)
  5. [The SalesData Class](#)
  6. [Partially Filled Arrays](#)
  7. [Working with Arrays and Files](#)
6. [7.5 Returning Arrays from Methods](#)

7. [7.6 String Arrays](#)
  1. [Calling String Methods from an Array Element](#)
    1. [Checkpoint](#)
8. [7.7 Arrays of Objects](#)
  1. [Checkpoint](#)
9. [7.8 The Sequential Search Algorithm](#)
10. [7.9 Two-Dimensional Arrays](#)
  1. [Initializing a Two-Dimensional Array](#)
  2. [The `length` Field in a Two-Dimensional Array](#)
  3. [Displaying All the Elements of a Two-Dimensional Array](#)
  4. [Summing All the Elements of a Two-Dimensional Array](#)
  5. [Summing the Rows of a Two-Dimensional Array](#)
  6. [Summing the Columns of a Two-Dimensional Array](#)
  7. [Passing Two-Dimensional Arrays to Methods](#)
  8. [Ragged Arrays](#)
11. [7.10 Arrays with Three or More Dimensions](#)
  1. [Checkpoint](#)
12. [7.11 The Selection Sort and the Binary Search Algorithms](#)
  1. [The Selection Sort Algorithm](#)
  2. [The Binary Search Algorithm](#)
    1. [Checkpoint](#)
13. [7.12 Command-Line Arguments and Variable-Length Argument Lists](#)
  1. [Command-Line Arguments](#)
  2. [Variable-Length Argument Lists](#)
14. [7.13 The `ArrayList` Class](#)
  1. [Creating and Using an `ArrayList` Object](#)
  2. [Using the Enhanced `for` Loop with an `ArrayList`](#)
  3. [The `ArrayList` Class's `toString` method](#)
  4. [Removing an Item from an `ArrayList`](#)
  5. [Inserting an Item](#)
  6. [Replacing an Item](#)
  7. [Capacity](#)
  8. [Storing Your Own Objects in an `ArrayList`](#)
  9. [Using the Diamond Operator for Type Inference](#)
    1. [Checkpoint](#)
15. [7.14 Common Errors to Avoid](#)
16. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
17. [Programming Challenges](#)
  1. [1. Rainfall Class](#)
  2. [2. Payroll Class](#)
  3. [3. Charge Account Validation](#)



4. [4. Charge Account Modification](#)
5. [5. Larger Than  \$n\$](#)
6. [6. Driver's License Exam](#)
7. [7. Magic 8 Ball](#)
8. [8. Grade Book](#)
9. [9. Grade Book Modification](#)
10. [10. Average Steps Taken](#)
11. [11. Array Operations](#)
12. [12. 1994 Gas Prices](#)
13. [13. Sorted List of 1994 Gas Prices](#)
14. [14. Name Search](#)
15. [15. Population Data](#)
16. [16. World Series Champions](#)
17. [17. 2D Array Operations](#)
18. [18. Phone Book `ArrayList`](#)
19. [19. Trivia Game](#)
20. [20. Lo Shu Magic Square](#)

## 12. [Chapter 8 A Second Look at Classes and Objects](#)

1. [TOPICS](#)
2. [8.1 Static Class Members](#)
  1. [A Quick Review of Instance Fields and Instance Methods](#)
  2. [Static Members](#)
  3. [Static Fields](#)
  4. [Static Methods](#)
    1. [Checkpoint](#)
3. [8.2 Passing Objects as Arguments to Methods](#)
4. [8.3 Returning Objects from Methods](#)
5. [8.4 The `toString` Method](#)
6. [8.5 Writing an `equals` Method](#)
7. [8.6 Methods That Copy Objects](#)
  1. [Copy Constructors](#)
8. [8.7 Aggregation](#)
  1. [Aggregation in UML Diagrams](#)
  2. [Security Issues with Aggregate Classes](#)
    1. [Perform Deep Copies When Creating Field Objects](#)
    2. [Return Copies of Field Objects, Not the Originals](#)
  3. [Avoid Using `null` References](#)
9. [8.8 The `this` Reference Variable](#)
  1. [Using `this` to Overcome Shadowing](#)
  2. [Using `this` to Call an Overloaded Constructor from Another Constructor](#)
    1. [Checkpoint](#)
10. [8.9 Enumerated Types](#)
  1. [Enumerated Types Are Specialized Classes](#)
  2. [Switching On an Enumerated Type](#)
    1. [Checkpoint](#)

11. [8.10 Garbage Collection](#)
  1. [The finalize Method](#)
12. [8.11 Focus on Object-Oriented Design: Class Collaboration](#)
  1. [Determining Class Collaborations with CRC Cards](#)
13. [8.12 Common Errors to Avoid](#)
14. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
15. [Programming Challenges](#)
  1. [1. Area Class](#)
  2. [2. BankAccount Class Copy Constructor](#)
  3. [3. Carpet Calculator](#)
  4. [4. LandTract Class](#)
  5. [5. Month Class](#)
  6. [6. CashRegister Class](#)
  7. [7. Sales Receipt File](#)
  8. [8. Parking Ticket Simulator](#)
  9. [9. Geometry Calculator](#)
  10. [10. Car Instrument Simulator](#)
  11. [11. First to One Game](#)
  12. [12. Heads or Tails Game](#)
13. [Chapter 9 Text Processing and More about Wrapper Classes](#)
  1. [TOPICS](#)
  2. [9.1 Introduction to Wrapper Classes](#)
  3. [9.2 Character Testing and Conversion with the Character Class](#)
    1. [Character Case Conversion](#)
      1. [Checkpoint](#)
  4. [9.3 More String Methods](#)
    1. [Searching for Substrings](#)
      1. [The startsWith and endsWith Methods](#)
      2. [The regionMatches Methods](#)
      3. [Finding Characters with the indexOf and lastIndexOf Methods](#)
      4. [Finding Substrings with the indexOf and lastIndexOf Methods](#)
    2. [Extracting Substrings](#)
      1. [The substring Methods](#)
      2. [The getChars and toCharArray Methods](#)
    3. [Methods That Return a Modified String](#)
    4. [The Static valueOf Methods](#)
      1. [Checkpoint](#)
  5. [9.4 The StringBuilder Class](#)
    1. [The StringBuilder Constructors](#)
    2. [Other StringBuilder Methods](#)
      1. [The append Methods](#)

2. [The insert Methods](#)
  3. [The replace Method](#)
  4. [The delete, deleteCharAt, and setCharAt Methods](#)
3. [The toString Method](#)
  1. [Checkpoint](#)
6. [9.5 Tokenizing Strings](#)
7. [9.6 Wrapper Classes for the Numeric Data Types](#)
  1. [The Static toString Methods](#)
  2. [The toBinaryString, toHexString, and toOctalString Methods](#)
  3. [The MIN\\_VALUE and MAX\\_VALUE Constants](#)
  4. [Autoboxing and Unboxing](#)
    1. [Checkpoint](#)
8. [9.7 Focus on Problem Solving: The TestScoreReader Class](#)
9. [9.8 Common Errors to Avoid](#)
10. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
11. [Programming Challenges](#)
  1. [1. Backward String](#)
  2. [2. Word Counter](#)
  3. [3. Sentence Capitalizer](#)
  4. [4. Vowels and Consonants](#)
  5. [5. Password Verifier](#)
  6. [6. Telemarketing Phone Number List](#)
  7. [7. Check Writer](#)
  8. [8. Sum of Numbers in a String](#)
  9. [9. Sum of Digits in a String](#)
  10. [10. Word Counter](#)
  11. [11. Sales Analysis](#)
  12. [12. Miscellaneous String Operations](#)
  13. [13. Alphabetic Telephone Number Translator](#)
  14. [14. Word Separator](#)
  15. [15. Pig Latin](#)
  16. [16. Morse Code Converter](#)
  17. [17. Lottery Statistics](#)
  18. [18. Gas Prices](#)
14. [Chapter 10 Inheritance](#)
  1. [TOPICS](#)
  2. [10.1 What Is Inheritance?](#)
    1. [Generalization and Specialization](#)
    2. [Inheritance and the “Is a” Relationship](#)
      1. [Fields:](#)
      2. [Methods:](#)

3. [Inheritance in UML Diagrams](#)
4. [The Superclass's Constructor](#)
5. [Inheritance Does Not Work in Reverse](#)
  1. [Checkpoint](#)
3. [10.2 Calling the Superclass Constructor](#)
  1. [When the Superclass Has No Default or No-Arg Constructors](#)
  2. [Summary of Constructor Issues in Inheritance](#)
    1. [Checkpoint](#)
4. [10.3 Overriding Superclass Methods](#)
  1. [Overloading versus Overriding](#)
  2. [Preventing a Method from Being Overridden](#)
    1. [Checkpoint](#)
5. [10.4 Protected Members](#)
  1. [Package Access](#)
    1. [Checkpoint](#)
6. [10.5 Chains of Inheritance](#)
  1. [Class Hierarchies](#)
7. [10.6 The `Object` Class](#)
  1. [Checkpoint](#)
8. [10.7 Polymorphism](#)
  1. [Polymorphism and Dynamic Binding](#)
  2. [The "Is-a" Relationship Does Not Work in Reverse](#)
  3. [The `instanceof` Operator](#)
    1. [Checkpoint](#)
9. [10.8 Abstract Classes and Abstract Methods](#)
  1. [Abstract Classes in UML](#)
    1. [Checkpoint](#)
10. [10.9 Interfaces](#)
  1. [An Interface is a Contract](#)
  2. [Fields in Interfaces](#)
  3. [Implementing Multiple Interfaces](#)
  4. [Interfaces in UML](#)
  5. [Default Methods](#)
  6. [Polymorphism and Interfaces](#)
    1. [Checkpoint](#)
11. [10.10 Anonymous Inner Classes](#)
12. [10.11 Functional Interfaces and Lambda Expressions](#)
13. [10.12 Common Errors to Avoid](#)
14. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
15. [Programming Challenges](#)
  1. [1. `Employee` and `ProductionWorker` Classes](#)

2. [2. ShiftSupervisor Class](#)
3. [3. TeamLeader Class](#)
4. [4. Essay Class](#)
5. [5. Course Grades](#)
6. [6. Analyzable Interface](#)
7. [7. Person and Customer Classes](#)
8. [8. PreferredCustomer Class](#)
9. [9. BankAccount and SavingsAccount Classes](#)
10. [10. Ship, CruiseShip, and CargoShip Classes](#)
15. [Chapter 11 Exceptions and Advanced File I/O](#)
  1. [TOPICS](#)
  2. [11.1 Handling Exceptions](#)
    1. [Exception Classes](#)
    2. [Handling an Exception](#)
    3. [Retrieving the Default Error Message](#)
    4. [Polymorphic References to Exceptions](#)
    5. [Using Multiple `catch` Clauses to Handle Multiple Exceptions](#)
      1. [Using Exception Handlers to Recover from Errors](#)
      2. [Handle Each Exception Only Once in a `try` Statement](#)
    6. [The `finally` Clause](#)
    7. [The Stack Trace](#)
    8. [Handling Multiple Exceptions with One `catch` Clause](#)
    9. [When an Exception Is Not Caught](#)
    10. [Checked and Unchecked Exceptions](#)
      1. [Checkpoint](#)
  3. [11.2 Throwing Exceptions](#)
    1. [Creating Your Own Exception Classes](#)
    2. [Using the `@exception` Tag in Documentation Comments](#)
      1. [Checkpoint](#)
  4. [11.3 Advanced Topics: Binary Files, Random Access Files, and Object Serialization](#)
    1. [Binary Files](#)
      1. [Writing Data to a Binary File](#)
      2. [Reading Data from a Binary File](#)
      3. [Writing and Reading Strings](#)
      4. [Appending Data to an Existing Binary File](#)
    2. [Random Access Files](#)
      1. [Reading and Writing with the `RandomAccessFile` Class](#)
      2. [The File Pointer](#)
    3. [Object Serialization](#)
    4. [Serializing Aggregate Objects](#)
      1. [Checkpoint](#)
  5. [11.4 Common Errors to Avoid](#)
  6. [Review Questions and Exercises](#)
    1. [Multiple Choice and True/False](#)
    2. [Find the Error](#)

3. [Algorithm Workbench](#)
4. [Short Answer](#)
7. [Programming Challenges](#)
  1. [1. TestScores Class](#)
  2. [2. TestScores Class Custom Exception](#)
  3. [3. RetailItem Exceptions](#)
  4. [4. Month Class Exceptions](#)
  5. [5. Payroll Class Exceptions](#)
  6. [6. FileArray Class](#)
  7. [7. File Encryption Filter](#)
  8. [8. File Decryption Filter](#)
  9. [9. TestScores Modification for Serialization](#)
  10. [10. Exception Project](#)
16. [Chapter 12 JavaFX: GUI Programming and Basic Controls](#)
  1. [TOPICS](#)
  2. [12.1 Graphical User Interfaces](#)
    1. [Event-Driven GUI Programs](#)
      1. [Checkpoint](#)
  3. [12.2 Introduction to JavaFX](#)
    1. [Controls](#)
    2. [Stages and Scenes](#)
    3. [The Application Class](#)
      1. [Checkpoint](#)
  4. [12.3 Creating Scenes](#)
    1. [Creating Controls](#)
    2. [Creating Layout Containers](#)
    3. [Creating a Scene Object](#)
    4. [Adding the Scene Object to the Stage](#)
    5. [Setting the Size of the Scene](#)
    6. [Aligning Controls in an HBox Layout Container](#)
      1. [Checkpoint](#)
  5. [12.4 Displaying Images](#)
    1. [Loading Images from an Internet Location](#)
    2. [Setting the Size of an Image](#)
    3. [Preserving the Image's Aspect Ratio](#)
    4. [Changing an ImageView's Image](#)
      1. [Checkpoint](#)
  6. [12.5 More about the HBox, VBox, and GridPane Layout Containers](#)
    1. [The HBox Layout Container](#)
      1. [Spacing](#)
      2. [Padding](#)
    2. [The VBox Layout Container](#)
    3. [The GridPane Layout Container](#)
    4. [Using Multiple Layout Containers in the Same Screen](#)
      1. [Checkpoint](#)

7. [12.6 Button Controls and Events](#)
  1. [Handling Events](#)
  2. [Writing Event Handlers](#)
  3. [Registering an Event Handler](#)
    1. [Checkpoint](#)
8. [12.7 Reading Input with `TextField` Controls](#)
  1. [Checkpoint](#)
9. [12.8 Using Anonymous Inner Classes and Lambda Expressions to Handle Events](#)
  1. [Using Anonymous Inner Classes to Create Event Handlers](#)
  2. [Using Lambda Expressions to Create Event Handlers](#)
    1. [Checkpoint](#)
10. [12.9 The `BorderPane` Layout Container](#)
  1. [Checkpoint](#)
11. [12.10 The `ObservableList` Interface](#)
12. [12.11 Common Errors to Avoid](#)
13. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
14. [Programming Challenges](#)
  1. [1. Latin Translator](#)
  2. [2. Name Formatter](#)
  3. [3. Tip, Tax, and Total](#)
  4. [4. Property Tax](#)
  5. [5. Heads or Tails](#)
  6. [6. Dice Simulator](#)
  7. [7. Travel Expenses](#)
  8. [8. Joe's Automotive](#)
  9. [9. Tic-Tac-Toe Simulator](#)
  10. [10. Slot Machine Simulation](#)
17. [Chapter 13 JavaFX: Advanced Controls](#)
  1. [TOPICS](#)
  2. [13.1 Styling JavaFX Applications with CSS](#)
    1. [Type Selector Names](#)
    2. [Style Properties](#)
    3. [Applying a Stylesheet to a JavaFX Application](#)
    4. [Applying Styles to the Root Node](#)
    5. [Specifying Multiple Selectors in the Same Style Definition](#)
    6. [Working with Colors](#)
      1. [Named Colors](#)
    7. [Creating a Custom Style Class Name](#)
    8. [ID Selectors](#)
    9. [Inline Style Rules](#)
      1. [Checkpoint](#)

3. [13.2 RadioButton Controls](#)
  1. [Determining in Code Whether a RadioButton Is Selected](#)
  2. [Selecting a RadioButton in Code](#)
  3. [Responding to RadioButton Clicks](#)
    1. [Checkpoint](#)
4. [13.3 CheckBox Controls](#)
  1. [Determining in Code Whether a CheckBox Is Selected](#)
  2. [Selecting a CheckBox in Code](#)
  3. [Responding to CheckBox Clicks](#)
    1. [Checkpoint](#)
5. [13.4 ListView Controls](#)
  1. [Retrieving the Selected Item](#)
  2. [Retrieving the Index of the Selected Item](#)
  3. [Responding to Item Selection with an Event Handler](#)
  4. [Adding Items versus Setting Items](#)
  5. [Initializing a ListView with an Array or an ArrayList](#)
  6. [Selection Modes](#)
  7. [Retrieving Multiple Selected Items](#)
  8. [Working With the Elements of an ObservableList](#)
  9. [Converting an ObservableList to an Array](#)
  10. [Using Code to Select an Item in a ListView](#)
  11. [ListView Orientation](#)
  12. [Creating ListViews of Objects Other Than String](#)
    1. [Checkpoint](#)
6. [13.5 ComboBox Controls](#)
  1. [Retrieving the Selected Item](#)
  2. [Responding to ComboBox Item Selection with an Event Handler](#)
  3. [Editable ComboBoxes](#)
    1. [Checkpoint](#)
7. [13.6 Slider Controls](#)
  1. [Checkpoint](#)
8. [13.7 TextArea Controls](#)
  1. [Checkpoint](#)
9. [13.8 Menus](#)
  1. [Assigning Mnemonics to Menu Items](#)
    1. [Checkpoint](#)
10. [13.9 The FileChooser Class](#)
  1. [Displaying a FileChooser Dialog Box](#)
    1. [Checkpoint](#)
11. [13.10 Using Console Output to Debug a GUI Application](#)
12. [13.11 Common Errors to Avoid](#)
13. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)



#### 4. [Short Answer](#)

### 14. [Programming Challenges](#)

1. [1. Dorm and Meal Plan Calculator](#)
2. [2. Skateboard Designer](#)
3. [3. Conference Registration System](#)
4. [4. Smartphone Packages](#)
5. [5. Shopping Cart System](#)

### 18. [Chapter 14 JavaFX: Graphics, Effects, and Media](#)

#### 1. [TOPICS](#)

#### 2. [14.1 Drawing Shapes](#)

1. [The Screen Coordinate System](#)
2. [The Shape Class and Its Subclasses](#)
3. [The Line Class](#)
4. [Changing the Stroke Color](#)
5. [The Circle Class](#)
6. [The Rectangle Class](#)
7. [The Ellipse Class](#)
8. [The Arc Class](#)
9. [The Polygon Class](#)
10. [The Polyline Class](#)
11. [The Text Class](#)
12. [Rotating Nodes](#)
13. [Scaling Nodes](#)

##### 1. [Checkpoint](#)

#### 3. [14.2 Animation](#)

1. [The TranslateTransition Class](#)
2. [The RotateTransition Class](#)
3. [The ScaleTransition Class](#)
4. [The StrokeTransition Class](#)
5. [The FillTransition Class](#)
6. [The FadeTransition Class](#)
7. [Controlling the Animation](#)
8. [Specifying an Interpolator](#)

##### 1. [Checkpoint](#)

#### 4. [14.3 Effects](#)

1. [The DropShadow Class](#)
2. [The InnerShadow Class](#)
3. [The ColorAdjust Class](#)
4. [The BoxBlur, GaussianBlur, and MotionBlur Classes](#)
  1. [The BoxBlur Class](#)
  2. [The GaussianBlur Class](#)
  3. [The MotionBlur Class](#)
5. [The SepiaTone Class](#)
6. [The Glow Class](#)
7. [The Reflection Class](#)

8. [Combining Effects](#)
  1. [Checkpoint](#)
5. [14.4 Playing Sound Files](#)
  1. [Registering an `EndOfMedia` Event Handler](#)
    1. [Checkpoint](#)
6. [14.5 Playing Videos](#)
  1. [Checkpoint](#)
7. [14.6 Handling Key Events](#)
  1. [Using an Anonymous Inner Class to Register a Key Event Handler to the Scene](#)
  2. [Using a Lambda Expression to Register a Key Event Handler to the Scene](#)
    1. [Checkpoint](#)
8. [14.7 Handling Mouse Events](#)
9. [14.8 Common Errors to Avoid](#)
10. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Algorithm Workbench](#)
  3. [Short Answer](#)
11. [Programming Challenges](#)
  1. [1. This Old House](#)
  2. [2. Tree Age](#)
  3. [3. Hollywood Star](#)
  4. [4. Vehicle Outline](#)
  5. [5. Solar System](#)
  6. [6. Opacity Adjuster](#)
  7. [7. Color Adjuster](#)
  8. [8. Mouse Rollover](#)
  9. [9. Coin Toss](#)
  10. [10. Lunar Lander](#)
  11. [11. Change for a Dollar Game](#)
  12. [12. Rock, Paper, Scissors Game](#)
19. [Chapter 15 Recursion](#)
  1. [TOPICS](#)
  2. [15.1 Introduction to Recursion](#)
  3. [15.2 Solving Problems with Recursion](#)
    1. [Direct and Indirect Recursion](#)
      1. [Checkpoint](#)
  4. [15.3 Examples of Recursive Methods](#)
    1. [Summing a Range of Array Elements with Recursion](#)
    2. [Drawing Concentric Circles](#)
    3. [The Fibonacci Series](#)
    4. [Finding the Greatest Common Divisor](#)
  5. [15.4 A Recursive Binary Search Method](#)
  6. [15.5 The Towers of Hanoi](#)
  7. [15.6 Common Errors to Avoid](#)
  8. [Review Questions and Exercises](#)

1. [Multiple Choice and True/False](#)
2. [Find the Error](#)
3. [Algorithm Workbench](#)
4. [Short Answer](#)
9. [Programming Challenges](#)
  1. [1. Recursive Multiplication](#)
  2. [2. isMember Method](#)
  3. [3. String Reverser](#)
  4. [4. maxElement Method](#)
  5. [5. Palindrome Detector](#)
  6. [6. Character Counter](#)
  7. [7. Recursive Power Method](#)
  8. [8. Sum of Numbers](#)
  9. [9. Ackermann's Function](#)
  10. [10. Recursive Population Class](#)
20. [Chapter 16 Sorting, Searching, and Algorithm Analysis](#)
  1. [TOPICS](#)
  2. [16.1 Introduction to Sorting Algorithms](#)
    1. [The Bubble Sort](#)
    2. [Using the Bubble Sort to Sort Objects](#)
    3. [The Selection Sort](#)
    4. [Using the Selection Sort to Sort Objects](#)
    5. [The Insertion Sort](#)
    6. [Using the Insertion Sort to Sort Objects](#)
    7. [The Quicksort Algorithm](#)
    8. [Using Quicksort to Sort Objects](#)
      1. [Checkpoint](#)
  3. [16.2 Introduction to Search Algorithms](#)
    1. [The Sequential Search Algorithm](#)
    2. [Average Case Analysis of the Sequential Search](#)
    3. [The Binary Search](#)
    4. [Efficiency of the Binary Search](#)
    5. [A Recursive Binary Search](#)
      1. [Checkpoint](#)
  4. [16.3 Analysis of Algorithms](#)
    1. [Computational Problems and Basic Steps](#)
    2. [Complexity of Algorithms](#)
    3. [Worst Case Complexity of Algorithms](#)
    4. [Average Case Complexity](#)
    5. [Asymptotic Complexity and the Big O Notation](#)
      1. [Checkpoint](#)
  5. [16.4 Common Errors to Avoid](#)
  6. [Review Questions and Exercises](#)
    1. [Multiple Choice and True/False](#)
    2. [Find the Error](#)

3. [Algorithm Workbench](#)

4. [Short Answer](#)

7. [Programming Challenges](#)

1. [1. Sorting Objects with the Insertion Sort Algorithm](#)

2. [2. Sorting Objects with the Quicksort Algorithm](#)

3. [3. Searching for Objects with the Binary Search Algorithm](#)

4. [4. Charge Account Validation](#)

5. [5. Charge Account Validation Modification](#)

6. [6. Name Search](#)

7. [7. Search Benchmarks](#)

8. [8. Sorting Benchmarks](#)

9. [9. Efficient Computation of Fibonacci Numbers](#)

21. [Chapter 17 Generics](#)

1. [TOPICS](#)

2. [17.1 Introduction to Generics](#)

1. [Using the Diamond Operator for Type Inference](#)

1. [Checkpoint](#)

3. [17.2 Writing a Generic Class](#)

1. [Only Reference Types Can Be Passed to Type Parameters](#)

1. [Autoboxing and Unboxing](#)

2. [Instantiating a Generic Class without Specifying a Type Argument](#)

3. [Commonly Used Type Parameter Names](#)

1. [Checkpoint](#)

4. [17.3 Passing Objects of a Generic Class to a Method](#)

1. [Constraining a Type Parameter](#)

2. [Defining a Type Parameter in a Method Header](#)

3. [The `extends` Key Word Constrains a Type to an Upper Bound](#)

4. [The `super` Key Word Constrains a Type to a Lower Bound](#)

1. [Checkpoint](#)

5. [17.4 Writing Generic Methods](#)

1. [Checkpoint](#)

6. [17.5 Constraining a Type Parameter in a Generic Class](#)

1. [Checkpoint](#)

7. [17.6 Inheritance and Generic Classes](#)

1. [Checkpoint](#)

8. [17.7 Defining Multiple Type Parameters](#)

9. [17.8 Generics and Interfaces](#)

1. [Constraining a Type Parameter to a Type That Implements an Interface](#)

1. [Checkpoint](#)

10. [17.9 Erasure](#)

1. [Checkpoint](#)

11. [17.10 Restrictions on the Use of Generic Types](#)

12. [17.11 Common Errors to Avoid](#)

13. [Review Questions and Exercises](#)

1. [Multiple Choice and True/False](#)

2. [Find the Error](#)
3. [Algorithm Workbench](#)
4. [Short Answer](#)

#### 14. [Programming Challenges](#)

1. [1. MyList Class](#)
2. [2. MyList Modification](#)
3. [3. PointList Class](#)
4. [4. Generic Insertion Sort](#)
5. [5. Generic Binary Search](#)
6. [6. Highest and Lowest Elements](#)
7. [7. Number Analyzer](#)
8. [8. Generic Random Permutation](#)

#### 22. [Chapter 18 Collections and the Stream API](#)

1. [TOPICS](#)
2. [18.1 Introduction to the Java Collections Framework](#)
  1. [Lists, Sets, and Maps](#)
  2. [Java Functional Interfaces](#)
  3. [JCF Interfaces](#)
  4. [Iterating a Collection](#)
  5. [The `Iterable` Interface and the `forEach` method](#)
  6. [Iterators](#)
  7. [The Enhanced `For` Loop](#)
  8. [The `Collection` Interface](#)
    1. [Checkpoint](#)
3. [18.2 Lists](#)
  1. [The `List` Interface](#)
  2. [The `ArrayList` and `LinkedList` Classes](#)
  3. [Creating and Initializing Lists](#)
  4. [List Iterators](#)
  5. [Methods of the `LinkedList` Class](#)
  6. [Using an Interface Variable to Reference a `Collection` Object](#)
    1. [Checkpoint](#)
4. [18.3 Sets](#)
  1. [The `Set` Interface](#)
  2. [The `HashSet` Class](#)
  3. [Implementing Your Own `hashCode` Method](#)
  4. [The `LinkedHashSet` Class](#)
  5. [The `SortedSet` Interface and the `TreeSet` Class](#)
  6. [Using a Comparator to Compare Objects](#)
    1. [Checkpoint](#)
5. [18.4 Maps](#)
  1. [The `HashMap` Class](#)
  2. [The `LinkedHashMap` Class](#)
  3. [The `SortedMap` Interface and the `TreeMap` Class](#)
    1. [Checkpoint](#)

6. [18.5 The Collections Class](#)
7. [18.6 The Stream API](#)
  1. [The Stream Package](#)
  2. [Creating a Stream](#)
  3. [Terminal Operations](#)
  4. [Intermediate Operations on Streams](#)
  5. [Intermediate Stream Operations are Lazy](#)
  6. [The Stream `map` Operation](#)
  7. [The Stream `reduce` Methods](#)
  8. [Stream Collectors](#)
    1. [Checkpoint](#)
8. [18.7 Common Errors to Avoid](#)
9. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
10. [Programming Challenges](#)
  1. [1. Word Set](#)
  2. [2. Instructor Set](#)
  3. [3. EmployeeMap Class](#)
  4. [4. Stock Map](#)
  5. [5. Prime Number List](#)
  6. [6. Dealing Cards](#)
  7. [7. Word Frequency Count](#)
  8. [8. groupingBy Collector](#)
23. [Chapter 19 Linked Lists](#)
  1. [TOPICS](#)
  2. [19.1 Introduction to Linked Lists](#)
    1. [Creating Lists and Adding Nodes](#)
    2. [Removing a Node](#)
    3. [Traversing Linked Lists](#)
      1. [Checkpoint](#)
  3. [19.2 Operations on Linked Lists](#)
    1. [The `isEmpty` Method](#)
    2. [The `size` Method](#)
    3. [The `add` Methods](#)
    4. [The `remove` Methods](#)
    5. [The `toString` Method](#)
    6. [A Graphical Interface for the Linked List Program](#)
  4. [19.3 Doubly Linked and Circularly Linked Lists](#)
    1. [Adding a Node to a Doubly Linked List](#)
    2. [Removing a Node from a Doubly Linked List](#)
    3. [Circularly Linked Lists](#)
      1. [Checkpoint](#)

5. [19.4 Recursion on Linked Lists](#)
  1. [A Recursive Implementation of `size`](#)
  2. [A Recursive Implementation of the `add\(String e, Node list\)` Method](#)
  3. [A Recursive Implementation of `add\(int index, String e, Node list\)`](#)
  4. [A Recursive Implementation of the `remove` Methods](#)
6. [19.5 Common Errors to Avoid](#)
7. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
8. [Programming Challenges](#)
  1. [1. Generic Linked Lists](#)
  2. [2. Linked List Sorting and Reversing](#)
  3. [3. Recursive Linked List Operations](#)
  4. [4. Top Ten Gamers](#)
  5. [5. Wedding Lottery](#)
  6. [6. Wedding Lottery II](#)
  7. [7. Adjacency List Input](#)
24. [Chapter 20 Stacks and Queues](#)
  1. [TOPICS](#)
  2. [20.1 Stacks and Their Applications](#)
    1. [Examples and Applications of Stacks](#)
    2. [Stack Operations](#)
    3. [Stacks of Primitive Types](#)
  3. [20.2 Array Implementation of Stacks](#)
    1. [The Stack Push Operation](#)
    2. [The Stack `empty` Method](#)
    3. [The Stack `peek` and `pop` Methods](#)
    4. [Stacks of Objects](#)
  4. [20.3 Linked Implementation of Stacks](#)
    1. [Implementation of Stack Methods](#)
      1. [Checkpoint](#)
  5. [20.4 Queues and Their Applications](#)
  6. [20.5 Array Implementation of Queues](#)
    1. [GUI Front End for the Queue Demo Program](#)
  7. [20.6 Linked List Implementation of Queues](#)
    1. [Queue Initialization and Enqueuing of Items](#)
    2. [Dequeuing Items](#)
      1. [Checkpoint](#)
  8. [20.7 Generic Implementation of Stacks and Queues](#)
  9. [20.8 Queues and Breadth-First Search](#)
  10. [20.9 Common Errors to Avoid](#)
  11. [Review Questions and Exercises](#)
    1. [Multiple Choice and True/False](#)

2. [Find the Error](#)
3. [Algorithm Workbench](#)
4. [Short Answer](#)

12. [Programming Challenges](#)

1. [1. Double-Ended Queue](#)
2. [2. Array-Based Deque](#)
3. [3. Prefix Expressions](#)
4. [4. Properly Nested Delimiters](#)
5. [5. Tracing Genealogies](#)

25. [Chapter 21 Binary Trees, AVL Trees, and Priority Queues](#)

1. [TOPICS](#)

2. [21.1 Binary Trees and Their Applications](#)

1. [Binary Tree Concepts](#)
2. [Applications of Binary Trees](#)
3. [Representing Nodes of Binary Trees](#)
4. [Traversing a Binary Tree](#)
  1. [Checkpoint](#)
5. [Graphical Display of Binary Trees](#)
6. [Class Implementation of Binary Trees](#)

3. [21.2 Binary Search Trees](#)

1. [Adding a Value to a Binary Search Tree](#)
2. [Removing a Value from a Binary Search Tree](#)
3. [Implementation of Binary Search Trees](#)
4. [Graphical Display of Binary Search Trees](#)
5. [A Program for Testing the Binary Search Tree Class](#)

4. [21.3 AVL Trees](#)

1. [Adding New Elements to AVL Trees](#)
2. [Implementation of AVL Trees](#)

5. [21.4 Priority Queues](#)

1. [Applications of Priority Queues](#)
2. [Using Comparators with Priority Queues](#)
3. [Analysis of Heapsort](#)
4. [Implementation of Priority Queues](#)
5. [Binary Trees with the Heap Order Property](#)
6. [Complete Binary Trees and Heaps](#)
7. [The Depth of Complete Binary Trees](#)
8. [Storing a Complete Binary Tree in an Array](#)
9. [Adding an Item to a Heap](#)
10. [Removing the Minimum Element](#)

6. [21.5 Common Errors to Avoid](#)

7. [Review Questions and Exercises](#)

1. [Multiple Choice and True/False](#)
2. [Find the Error](#)
3. [Short Answer](#)
4. [Algorithm Workbench](#)



8. [Programming Challenges](#)
  1. [1. AVL Trees](#)
  2. [2. Family Tree](#)
  3. [3. Balanced Search Tree](#)
  4. [4. Prefix Calculator](#)
  5. [5. Levels in an AVL Tree](#)
  6. [6. Binary Tree Input and Display](#)
  7. [7. Breadth-First Traversal of Binary trees](#)

26. [Chapter 22 Databases](#)

1. [TOPICS](#)
2. [22.1 Introduction to Database Management Systems](#)
  1. [JDBC](#)
  2. [SQL](#)
  3. [Using a DBMS](#)
  4. [Java DB and Apache Derby](#)
  5. [Creating the `CoffeeDB` Database](#)
  6. [Connecting to the `CoffeeDB` Database](#)
  7. [Connecting to a Password-Protected Database](#)
    1. [Checkpoint](#)
3. [22.2 Tables, Rows, and Columns](#)
  1. [Column Data Types](#)
  2. [Primary Keys](#)
    1. [Checkpoint](#)
4. [22.3 Introduction to the SQL `SELECT` Statement](#)
  1. [Passing an SQL Statement to the DBMS](#)
    1. [Getting a Row from the `ResultSet` Object](#)
    2. [Getting Columns in a `ResultSet` Row](#)
    3. [More about the `SELECT` Statement](#)
  2. [Specifying Search Criteria with the `WHERE` Clause](#)
    1. [String Comparisons in a `SELECT` Statement](#)
    2. [Using the `LIKE` Operator](#)
    3. [Using `AND` and `OR`](#)
  3. [Sorting the Results of a `SELECT` Query](#)
  4. [Mathematical Functions](#)
    1. [Checkpoint](#)
5. [22.4 Inserting Rows](#)
  1. [Inserting Rows with JDBC](#)
    1. [Checkpoint](#)
6. [22.5 Updating and Deleting Existing Rows](#)
  1. [Updating Rows with JDBC](#)
  2. [Deleting Rows with the `DELETE` Statement](#)
  3. [Deleting Rows with JDBC](#)
    1. [Checkpoint](#)
7. [22.6 Creating and Deleting Tables](#)
  1. [Removing a Table with the `DROP TABLE` Statement](#)

1. [Checkpoint](#)
8. [22.7 Creating a New Database with JDBC](#)
9. [22.8 Scrollable Result Sets](#)
  1. [ResultSet Navigation Methods](#)
10. [22.9 Result Set Metadata](#)
11. [22.10 Relational Data](#)
  1. [Joining Data from Multiple Tables](#)
  2. [An Order Entry System](#)
12. [22.11 Advanced Topics](#)
  1. [Transactions](#)
  2. [Stored Procedures](#)
13. [22.12 Common Errors to Avoid](#)
14. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
15. [Programming Challenges](#)
  1. [1. Customer Insertor](#)
  2. [2. Customer Updater](#)
  3. [3. Unpaid Order Sum](#)
  4. [4. Population Database](#)
  5. [5. Personnel Database Creator](#)
  6. [6. Employee Insertor](#)
  7. [7. Employee Updater](#)
  8. [8. PhoneBook Database](#)
27. [Chapter 23 A First Look at GUI Applications with Swing](#)
  1. [TOPICS](#)
  2. [23.1 Introduction](#)
    1. [The JFC, AWT, and Swing](#)
    2. [Event-Driven Programming](#)
    3. [The javax.swing and java.awt Packages](#)
  3. [23.2 Creating Windows](#)
    1. [Using Inheritance to Extend the JFrame Class](#)
    2. [Equipping GUI Classes with a main Method](#)
    3. [Adding Components to a Window](#)
      1. [Content Panes and Panels](#)
    4. [Handling Events with Action Listeners](#)
      1. [Writing Event Listener Classes as Private Inner Classes](#)
      2. [Event Listeners Must Implement an Interface](#)
      3. [Registering an Event Listener Object](#)
    5. [Writing an Event Listener for the KiloConverter Class](#)
    6. [Background and Foreground Colors](#)
      1. [Changing the Background Color of a JFrame Object's Content Pane](#)
    7. [The ActionEvent Object](#)

1. [The `getActionCommand` Method](#)
  2. [The `getSource` Method](#)
    1. [Checkpoint](#)
4. [23.3 Layout Managers](#)
  1. [Adding a Layout Manager to a Container](#)
  2. [The `FlowLayout` Manager](#)
    1. [Adjusting the `FlowLayout` Alignment](#)
    2. [Adjusting the `FlowLayout` Component Gaps](#)
  3. [The `BorderLayout` Manager](#)
    1. [Nesting Panels Inside a Container's Regions](#)
  4. [The `GridLayout` Manager](#)
    1. [Checkpoint](#)
5. [23.4 Radio Buttons and Check Boxes](#)
  1. [Radio Buttons](#)
    1. [Grouping with the `ButtonGroup` class](#)
    2. [Responding to Radio Button Events](#)
    3. [Determining in Code Whether a Radio Button Is Selected](#)
    4. [Selecting a Radio Button in Code](#)
  2. [Check Boxes](#)
    1. [Responding to Check Box Events](#)
    2. [Determining in Code Whether a Check Box Is Selected](#)
    3. [Selecting a Check Box in Code](#)
      1. [Checkpoint](#)
6. [23.5 Borders](#)
  1. [Empty Borders](#)
  2. [Line Borders](#)
  3. [Titled Borders](#)
    1. [Checkpoint](#)
7. [23.6 Focus on Problem Solving: Extending Classes from `JPanel`](#)
  1. [The Brandi's Bagel House Application](#)
  2. [The `GreetingPanel` Class](#)
  3. [The `BagelPanel` Class](#)
  4. [The `ToppingPanel` Class](#)
  5. [The `CoffeePanel` Class](#)
  6. [Putting It All Together](#)
8. [23.7 Splash Screens](#)
9. [23.8 Using Console Output to Debug a GUI Application](#)
10. [23.9 Common Errors to Avoid](#)
11. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
12. [Programming Challenges](#)
  1. [1. Retail Price Calculator](#)

2. [2. Monthly Sales Tax](#)
3. [3. Property Tax](#)
4. [4. Travel Expenses](#)
5. [5. Theater Revenue](#)
6. [6. Joe's Automotive](#)
7. [7. Long Distance Calls](#)
8. [8. Latin Translator](#)
9. [9. MPG Calculator](#)
10. [10. Celsius to Fahrenheit](#)

## 28. [Chapter 24 Advanced Swing GUI Applications](#)

1. [TOPICS](#)
2. [24.1 The Swing and AWT Class Hierarchy](#)
3. [24.2 Read-Only Text Fields](#)
4. [24.3 Lists](#)
  1. [Selection Modes](#)
  2. [Responding to List Events](#)
  3. [Retrieving the Selected Item](#)
  4. [Placing a Border around a List](#)
  5. [Adding a Scroll Bar to a List](#)
  6. [Adding Items to an Existing `JList` Component](#)
  7. [Multiple Selection Lists](#)
    1. [Single Interval Selection Mode](#)
    2. [Multiple Interval Selection Mode](#)
5. [24.4 Combo Boxes](#)
  1. [Retrieving the Selected Item](#)
  1. [Editable Combo Boxes](#)
  1. [Checkpoint](#)
6. [24.5 Displaying Images in Labels and Buttons](#)
  1. [Checkpoint](#)
7. [24.6 Mnemonics and Tool Tips](#)
  1. [Mnemonics](#)
  2. [Tool Tips](#)
    1. [Checkpoint](#)
8. [24.7 File Choosers and Color Choosers](#)
  1. [File Choosers](#)
    1. [Displaying a File Chooser Dialog Box](#)
  2. [Color Choosers](#)
9. [24.8 Menus](#)
  1. [Checkpoint](#)
10. [24.9 More about Text Components: Text Areas and Fonts](#)
  1. [Text Areas](#)
  2. [Fonts](#)
    1. [Checkpoint](#)
11. [24.10 Sliders](#)
  1. [Checkpoint](#)

12. [24.11 Look and Feel](#)
13. [24.12 Common Errors to Avoid](#)
14. [Review Questions and Exercises](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
15. [Programming Challenges](#)
  1. [1. Scrollable Tax Calculator](#)
  2. [2. Image Viewer](#)
  3. [3. Dorm and Meal Plan Calculator](#)
  4. [4. Skateboard Designer](#)
  5. [5. Shopping Cart System](#)
  6. [6. Cell Phone Packages](#)
  7. [7. Shade Designer](#)
  8. [8. Conference Registration System](#)
  9. [9. Dice Simulator](#)
  10. [10. Card Dealer](#)
  11. [11. Tic Tac Toe Simulator](#)
29. [Chapter 25 Applets and More](#)
  1. [TOPICS](#)
  2. [25.1 Introduction to Applets](#)
    1. [Checkpoint](#)
  3. [25.2 A Brief Introduction to HTML](#)
    1. [Hypertext](#)
    2. [Markup Language](#)
    3. [Document Structure Tags](#)
    4. [Text Formatting Tags](#)
    5. [Creating Breaks in Text](#)
    6. [Inserting Links](#)
      1. [Checkpoint](#)
  4. [25.3 Creating Applets with Swing](#)
    1. [Running an Applet](#)
      1. [Running an Applet with `appletviewer`](#)
    2. [Handling Events in an Applet](#)
      1. [Checkpoint](#)
  5. [25.4 Using AWT for Portability](#)
    1. [Checkpoint](#)
  6. [25.5 Drawing Shapes](#)
    1. [The `XY` Coordinate System](#)
    2. [Graphics Objects](#)
    3. [The `repaint` Method](#)
    4. [Drawing on Panels](#)
      1. [Checkpoint](#)
  7. [25.6 Handling Mouse Events](#)

1. [Handling Mouse Events](#)

1. [Using Adapter Classes](#)

1. [Checkpoint](#)

8. [25.7 Timer Objects](#)

1. [Checkpoint](#)

9. [25.8 Playing Audio](#)

1. [Using an AudioClip Object](#)

2. [Playing Audio in an Application](#)

1. [Checkpoint](#)

10. [25.9 Common Errors to Avoid](#)

11. [Review Questions and Exercises](#)

1. [Multiple Choice and True/False](#)

2. [Find the Error](#)

3. [Algorithm Workbench](#)

4. [Short Answer](#)

12. [Programming Challenges](#)

1. [1. FollowMe Applet](#)

2. [2. House Applet](#)

3. [3. WatchMe Applet](#)

4. [4. Thermometer Applet](#)

5. [5. Polygon Drawer](#)

6. [6. GridFiller Applet](#)

7. [7. DrinkMachine Applet](#)

8. [8. Stopwatch Applet](#)

9. [9. Slideshow Application](#)

30. [Index](#)

1. [A](#)

2. [B](#)

3. [C](#)

4. [D](#)

5. [E](#)

6. [F](#)

7. [G](#)

8. [H](#)

9. [I](#)

10. [J](#)

11. [K](#)

12. [L](#)

13. [M](#)

14. [N](#)

15. [O](#)

16. [P](#)

17. [Q](#)

18. [R](#)

19. [S](#)

20. [T](#)
21. [U](#)
22. [V](#)
23. [W](#)
24. [Z](#)
31. [APPENDIX A The ASCII/Unicode Characters\\*](#)
32. [APPENDIX B Operator Precedence and Associativity\\*](#)
33. [APPENDIX C Java Key Words\\*](#)
34. [APPENDIX D Installing the JDK and Using the JDK Documentation](#)
35. [APPENDIX E Using the javadoc Utility](#)
  1. [Writing Documentation Comments](#)
  2. [The @param Tag](#)
  3. [The @return Tag](#)
  4. [The @exception and @throws Tags](#)
  5. [Running javadoc](#)
  6. [The @author and @version Tags](#)
  7. [Embedding HTML in Documentation Comments](#)
  8. [The -public, -private, and -protected Command Line Options](#)
36. [APPENDIX F More about the Math Class](#)
37. [APPENDIX G Packages](#)
38. [APPENDIX H Working with Records and Random Access Files\\*](#)
39. [APPENDIX I Configuring Java DB and Installing Apache Derby](#)
  1. [Configuring Java DB with Java 7 or Java 8](#)
    1. [Setting the CLASSPATH Environment Variable](#)
      1. [Windows 10](#)
      2. [Windows 8](#)
      3. [Windows 7](#)
  2. [Configuring Derby with Java 9](#)
    1. [Setting the CLASSPATH Environment Variable](#)
      1. [Windows 10](#)
      2. [Windows 8](#)
      3. [Windows 7](#)
40. [APPENDIX J The QuickSort Algorithm](#)
41. [APPENDIX K Named Colors](#)
42. [APPENDIX L Answers to Checkpoints\\*](#)
  1. [Chapter 1](#)
  2. [Chapter 2](#)
  3. [Chapter 3](#)
  4. [Chapter 4](#)
  5. [Chapter 5](#)
  6. [Chapter 6](#)
  7. [Chapter 7](#)
  8. [Chapter 8](#)
  9. [Chapter 9](#)
  10. [Chapter 10](#)

11. [Chapter 11](#)
12. [Chapter 12](#)
13. [Chapter 13](#)
14. [Chapter 14](#)
15. [Chapter 15](#)
16. [Chapter 16](#)
17. [Chapter 17](#)
18. [Chapter 18](#)
19. [Chapter 19](#)
20. [Chapter 20](#)
21. [Chapter 21](#)
22. [Chapter 22 \(Available on the Book's Companion Website\)](#)
23. [Chapter 23 \(Available on the Book's Companion Website\)](#)
24. [Chapter 24 \(Available on the Book's Companion Website\)](#)
25. [Chapter 25 \(Available on the Book's Companion Website\)](#)
43. [APPENDIX M Answers to Odd-Numbered Review Questions\\*](#)
  1. [Chapter 1](#)
    1. [Multiple Choice](#)
    2. [Find the Error](#)
    3. [Algorithm Workbench](#)
    4. [Predict the Result](#)
    5. [Short Answer](#)
  2. [Chapter 2](#)
    1. [Multiple Choice and True/False](#)
    2. [Predict the Output](#)
    3. [Find the Error](#)
    4. [Algorithm Workbench](#)
    5. [Short Answer](#)
  3. [Chapter 3](#)
    1. [Multiple Choice and True/False](#)
    2. [Find the Error](#)
    3. [Algorithm Workbench](#)
    4. [Short Answer](#)
  4. [Chapter 4](#)
    1. [Multiple Choice and True/False](#)
    2. [Find the Error](#)
    3. [Algorithm Workbench](#)
    4. [Short Answer](#)
  5. [Chapter 5](#)
    1. [Multiple Choice and True/False](#)
    2. [Find the Error](#)
    3. [Algorithm Workbench](#)
    4. [Short Answer](#)
  6. [Chapter 6](#)
    1. [Multiple Choice and True/False](#)



2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
7. [Chapter 7](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
8. [Chapter 8](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
9. [Chapter 9](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
10. [Chapter 10](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
11. [Chapter 11](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
12. [Chapter 12](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
13. [Chapter 13](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
14. [Chapter 14](#)
  1. [Multiple Choice and True/False](#)
  2. [Algorithm Workbench](#)
  3. [Short Answer](#)
15. [Chapter 15](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)

3. [Algorithm Workbench](#)
4. [Short Answer](#)
16. [Chapter 16](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
17. [Chapter 17](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
18. [Chapter 18](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
19. [Chapter 19](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
20. [Chapter 20](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
21. [Chapter 21](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Short Answer](#)
  4. [Algorithm Workbench](#)
22. [Chapter 22 \(Available on the Book's Companion Website\)](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
23. [Chapter 23 \(Available on the Book's Companion Website\)](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
24. [Chapter 24 \(Available on the Book's Companion Website\)](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)

3. [Algorithm Workbench](#)
4. [Short Answer](#)
25. [Chapter 25 \(Available on the Book's Companion Website\)](#)
  1. [Multiple Choice and True/False](#)
  2. [Find the Error](#)
  3. [Algorithm Workbench](#)
  4. [Short Answer](#)
44. [CASE STUDY 1 Calculating Sales Commission\\*](#)
  1. [Program Design](#)
45. [CASE STUDY 2 The `Amortization` Class\\*](#)
  1. [Calculations](#)
  2. [Report Requirements](#)
  3. [The Main Program](#)
46. [CASE STUDY 3 The `PinTester` Class\\*](#)
47. [CASE STUDY 4 Parallel Arrays\\*](#)
48. [CASE STUDY 5 The `FeetInches` Class\\*](#)
  1. [More about the `toString` Method](#)
  2. [The `add` Method](#)
  3. [The `equals` Method](#)
  4. [The `copy` Method](#)
  5. [The Copy Constructor](#)
49. [CASE STUDY 6 The `SerialNumber` Class\\*](#)
50. [CASE STUDY 7 A Simple Text Editor Application\\*](#)
  1. [File Menu](#)
  2. [Font Menu](#)

## List of Illustrations

1. [Figure P-1 Chapter dependencies](#)
2. [Figure 1-1 The organization of a computer system](#) 1-1a (TL) iko/Shutterstock, 1-1b (TL) Nikita Rogul/Shutterstock, 1-1c (TLC) Feng Yu/Shutterstock, 1-1d (TLC) Chiyacat/Shutterstock, 1-1e Elkostas/Shutterstock, 1-1f (TLB) tkemot/Shutterstock, 1-1g (BC) Vitaly Korovin/Shutterstock, 1-1h (TRB) Lusoimages/Shutterstock, 1-1i (TRC) jocic/Shutterstock, 1-1j (TR) Best Pictures here/Shutterstock, 1-1k (MC) Peter Guess/Shutterstock, 1-1l (TC) Aquila/Shutterstock
3. [Figure 1-2 The organization of the CPU](#)
4. [Figure 1-3 Memory bytes and their addresses](#)
5. [Figure 1-4 A variable name represents a location in memory](#)
6. [Figure 1-5 Program development process](#)
7. [Figure 1-6 Java byte code may be run on any computer with a Java Virtual Machine](#)
8. [Figure 1-7 An integrated development environment \(IDE\)](#)
9. [Figure 1-8 Screen produced by the pay-calculating algorithm](#)
10. [Figure 1-9 Data is passed among procedures](#)
11. [Figure 1-10 An object contains data and procedures](#)

12. [Figure 1-11 Code outside the object interacts with the object's methods](#)
13. [Figure 2-1 The main method header](#)
14. [Figure 2-2 A console window](#)
15. [Figure 2-3 Relationship among the System class, the out object, and the print and println methods](#)
16. [Figure 2-4 Characters and how they are stored in memory](#)
17. [Figure 2-5 Precedence illustrated](#)
18. [Figure 2-6 Primitive data type ranking](#)
19. [Figure 2-7 A primitive type variable holds the data with which it is associated](#)
20. [Figure 2-8 A String class variable can hold the address of a String object](#)
21. [Figure 2-9 The name variable holds the address of a String object](#)
22. [Figure 2-10 Documentation generated by javadoc](#)
23. [Figure 2-11 Indentation](#)
24. [Figure 2-12 The parts of the statement](#)
25. [Figure 2-13 The keyboard variable references a Scanner object](#)
26. [Figure 2-14 A message dialog and an input dialog](#)
27. [Figure 2-15 Message dialog](#)
28. [Figure 2-16 Input dialog](#)
29. [Figure 2-17 Dialog boxes displayed by the NamesDialog program](#)
30. [Figure 2-18 Dialog boxes displayed by PayrollDialog.java](#)
31. [Figure 2-19 Dialog boxes](#)
32. [Figure 3-1 Sequence structure](#)
33. [Figure 3-2 Simple decision structure logic](#)
34. [Figure 3-3 Three-action decision structure logic](#)
35. [Figure 3-4 Logic of the if statement](#)
36. [Figure 3-5 Do not prematurely terminate an if statement with a semicolon](#)
37. [Figure 3-6 An if statement missing its braces](#)
38. [Figure 3-7 Logic of the if-else statement](#)
39. [Figure 3-8 Logic of nested if statements](#)
40. [Figure 3-9 Alignment of if and else clauses](#)
41. [Figure 3-10 Nested decision structure to determine a grade](#)
42. [Figure 3-11 Interaction with the NestedDecision program](#)
43. [Figure 3-12 Interaction with the NestedDecision program](#)
44. [Figure 3-13 Interaction with the NestedDecision program](#)
45. [Figure 3-14 The name1 and name2 variables reference different String objects](#)
46. [Figure 3-15 String comparison of "Mary" and "Mark"](#)
47. [Figure 3-16 A multiple alternative decision structure](#)
48. [Figure 3-17 The value of the sales variable is displayed in the place of the %f format specifier](#)
49. [Figure 3-18 The format specifiers and their corresponding arguments](#)
50. [Figure 3-19 The format specifiers and their corresponding arguments](#)
51. [Figure 3-20 The number is displayed in a field that is 20 spaces wide](#)
52. [Figure 3-21 Format specifier that pads with leading zeros](#)
53. [Figure 3-22 Output of Code Listing 3-21](#)
54. [Figure 3-23 Troubleshooting a bad Wi-Fi connection](#)
55. [Figure 4-1 Logic of a while loop](#)
56. [Figure 4-2 The while loop](#)

- 57. [Figure 4-3 Logic of the example while loop](#)
- 58. [Figure 4-4 Input validation logic](#)
- 59. [Figure 4-5 Interaction with the SoccerTeams program](#)
- 60. [Figure 4-6 Logic of the do-while loop](#)
- 61. [Figure 4-7 Sequence of events in the for loop](#)
- 62. [Figure 4-8 Logic of the for loop](#)
- 63. [Figure 4-9 Sequence of events with the for loop in Code Listing 4-7](#)
- 64. [Figure 4-10 Logic of the for loop in Code Listing 4-7](#)
- 65. [Figure 4-11 Logic for calculating a running total](#)
- 66. [Figure 4-12 Interaction with the TotalSales program](#)
- 67. [Figure 4-13 File contents displayed in Notepad](#)
- 68. [Figure 4-14 Contents of file displayed in Notepad](#)
- 69. [Figure 4-15 Contents of the file displayed in Notepad](#)
- 70. [Figure 4-16 File with three lines](#)
- 71. [Figure 4-17 Initial read position](#)
- 72. [Figure 4-18 Read position after first line is read](#)
- 73. [Figure 4-19 Logic of reading a file until the end is reached](#)
- 74. [Figure 4-20 Contents of Numbers.txt](#)
- 75. [Figure 5-1 Using methods to divide and conquer a problem](#)
- 76. [Figure 5-2 The header and body of a main method](#)
- 77. [Figure 5-3 Parts of the method header](#)
- 78. [Figure 5-4 Branching in the SimpleMethod.java program](#)
- 79. [Figure 5-5 Interaction with the CreditCard.java program](#)
- 80. [Figure 5-6 Interaction with the CreditCard.java program](#)
- 81. [Figure 5-7 Passing 5 to the displayValue method](#)
- 82. [Figure 5-8 Multiple arguments passed into multiple parameters](#)
- 83. [Figure 5-9 Passing a reference as an argument](#)
- 84. [Figure 5-10 Both name and str reference the same object](#)
- 85. [Figure 5-11 Before line 36 executes, both name and str reference the same object](#)
- 86. [Figure 5-12 After line 36 executes, name and str reference different objects](#)
- 87. [Figure 5-13 Documentation for the showSum method in TwoArgs2.java](#)
- 88. [Figure 5-14 Return type in the method header](#)
- 89. [Figure 5-15 Arguments passed to sum and a value returned](#)
- 90. [Figure 5-16 Documentation for the sum method in ValueReturn.java](#)
- 91. [Figure 5-17 Interaction with the CupConverter program](#)
- 92. [Figure 5-18 The fullName method returning a reference to a String object](#)
- 93. [Figure 5-19 Interaction with the SalesReport.java program](#)
- 94. [Figure 6-1 A blueprint and houses built from the blueprint](#)
- 95. [Figure 6-2 Objects created by the ObjectDemo program](#)
- 96. [Figure 6-3 Memory allocation](#)
- 97. [Figure 6-4 Values assigned to the variables](#)
- 98. [Figure 6-5 Two objects referenced by variables](#)
- 99. [Figure 6-6 The kite and string metaphor](#)
- 100. [Figure 6-7 The rand variable references a Random object](#)
- 101. [Figure 6-8 General layout of a UML diagram for a class](#)

- 102. [Figure 6-9 UML diagram for the Rectangle class](#)
- 103. [Figure 6-10 Header for the setLength method](#)
- 104. [Figure 6-11 The box variable references a Rectangle class object](#)
- 105. [Figure 6-12 The state of the box object after the setLength method executes](#)
- 106. [Figure 6-13 State of the box object](#)
- 107. [Figure 6-14 UML diagram for the Rectangle class](#)
- 108. [Figure 6-15 UML diagram for the Rectangle class with parameter and data type notation](#)
- 109. [Figure 6-16 Typical layout of class members](#)
- 110. [Figure 6-17 UML diagram](#)
- 111. [Figure 6-18 Interaction with the RoomAreas.java program](#)
- 112. [Figure 6-19 The kitchen, bedroom, and den variables reference Rectangle objects](#)
- 113. [Figure 6-20 States of the objects after data has been stored in them](#)
- 114. [Figure 6-21 Fill in the boxes for each field](#)
- 115. [Figure 6-22 UML diagram for the Rectangle class showing the constructor](#)
- 116. [Figure 6-23 UML diagram for the CellPhone class](#)
- 117. [Figure 6-24 UML diagram for the Die class](#)
- 118. [Figure 6-25 UML diagram for the Dealer class](#)
- 119. [Figure 6-26 UML diagram for the Player class](#)
- 120. [Figure 6-27 UML diagram for the BankAccount class](#)
- 121. [Figure 6-28 Interaction with the AccountTest program](#)
- 122. [Figure 6-29 UML diagram for the Customer class](#)
- 123. [Figure 6-30 UML diagram for the Car class](#)
- 124. [Figure 6-31 UML diagram for the ServiceQuote class](#)
- 125. [Figure 7-1 Variable declarations and their memory allocations](#)
- 126. [Figure 7-2 The numbers array](#)
- 127. [Figure 7-3 Subscripts for the numbers array](#)
- 128. [Figure 7-4 Contents of the array after 20 is assigned to numbers\[0\] and 30 is assigned to numbers\[3\]](#)
- 129. [Figure 7-5 Contents of the hours array](#)
- 130. [Figure 7-6 Annotated loop](#)
- 131. [Figure 7-7 The numbers variable references a ten-element array](#)
- 132. [Figure 7-8 The numbers variable references a five-element array](#)
- 133. [Figure 7-9 Both array1 and array2 reference the same array](#)
- 134. [Figure 7-10 An array passed as an argument](#)
- 135. [Figure 7-11 Interaction with the Sales.java program](#)
- 136. [Figure 7-12 UML diagram for the Grader class](#)
- 137. [Figure 7-13 Array reference return type](#)
- 138. [Figure 7-14 The names variable references a String array](#)
- 139. [Figure 7-15 An uninitialized String array](#)
- 140. [Figure 7-16 The accounts variable references an array of references](#)
- 141. [Figure 7-17 Each element of the array references an object](#)
- 142. [Figure 7-18 Rows and columns](#)
- 143. [Figure 7-19 Declaration of a two-dimensional array](#)
- 144. [Figure 7-20 Subscripts for each element of the scores array](#)
- 145. [Figure 7-21 Division and quarter data stored in the sales array](#)

146. [Figure 7-22 The numbers array](#)  
147. [Figure 7-23 The numbers array is an array of arrays](#)  
148. [Figure 7-24 A three-dimensional array](#)  
149. [Figure 7-25 Values in an array](#)  
150. [Figure 7-26 Values in array after first swap](#)  
151. [Figure 7-27 Values in array after second swap](#)  
152. [Figure 7-28 Values in array after third swap](#)  
153. [Figure 7-29 Values in array after fourth swap](#)  
154. [Figure 7-30 Values in array after fifth swap](#)  
155. [Figure 7-31 Lo Shu Magic Square](#)  
156. [Figure 7-32 Row, column, and diagonal sums in the Lo Shu Magic Square](#)  
157. [Figure 8-1 All instances of the class share the static field](#)  
158. [Figure 8-2 Interaction with the MetricDemo.java program](#)  
159. [Figure 8-3 Passing a reference as an argument](#)  
160. [Figure 8-4 Both box and r reference the same object](#)  
161. [Figure 8-5 Interaction with the ReturnObject.java program](#)  
162. [Figure 8-6 The getAccount method header](#)  
163. [Figure 8-7 The getAccount method returns a reference to a BankAccount object](#)  
164. [Figure 8-8 UML diagram for the Stock class](#)  
165. [Figure 8-9 The if statement tests the contents of the reference variables, not the contents of the objects the variables reference](#)  
166. [Figure 8-10 Both variables reference the same object](#)  
167. [Figure 8-11 UML diagram for the Instructor class](#)  
168. [Figure 8-12 UML diagram for the TextBook class](#)  
169. [Figure 8-13 UML diagram for the Course class](#)  
170. [Figure 8-14 UML diagram showing aggregation](#)  
171. [Figure 8-15 The workDay variable references the Day.WEDNESDAY object](#)  
172. [Figure 8-16 The Day enumerated data type and the ordinal positions of its enum constants](#)  
173. [Figure 8-17 Both account1 and account2 reference the same object](#)  
174. [Figure 8-18 The object is only referenced by the account2 variable](#)  
175. [Figure 8-19 The object is no longer referenced](#)  
176. [Figure 8-20 CRC card](#)  
177. [Figure 8-21 UML diagram for Programming Challenge 3](#)  
178. [Figure 9-1 Interaction with the CharacterTest.java program](#)  
179. [Figure 9-2 Interaction with the CharacterTest.java program](#)  
180. [Figure 9-3 Interaction with the CustomerNumber.java program](#)  
181. [Figure 9-4 Interaction with the CustomerNumber.java program](#)  
182. [Figure 9-5 The fullName and lastName variables reference separate objects](#)  
183. [Figure 9-6 The String object containing “George” is no longer referenced](#)  
184. [Figure 9-7 Microsoft Excel spreadsheet](#)  
185. [Figure 10-1 Bumblebees and grasshoppers are specialized versions of an insect](#)  
186. [Figure 10-2 UML diagram for the GradedActivity class](#)  
187. [Figure 10-3 Interaction with the GradeDemo.java program](#)  
188. [Figure 10-4 Interaction with the GradeDemo.java program](#)  
189. [Figure 10-5 UML diagram for the FinalExam class](#)



190. [Figure 10-6 FinalExam class header](#)

191. [Figure 10-7 Interaction with the FinalExamDemo.java program](#)

192. [Figure 10-8 UML diagram showing inheritance](#)

193. [Figure 10-9 UML diagram for the Rectangle class](#)

194. [Figure 10-10 UML diagram for the Rectangle and Cube classes](#)

195. [Figure 10-11 The GradedActivity and CurvedActivity classes](#)

196. [Figure 10-12 Interaction with the ProtectedDemo.java program](#)

197. [Figure 10-13 UML diagram for the GradedActivity2 class](#)

198. [Figure 10-14 A chain of inheritance](#)

199. [Figure 10-15 The GradedActivity, PassFailActivity, and PassFailExam classes](#)

200. [Figure 10-16 Class hierarchy](#)

201. [Figure 10-17 The line of inheritance from Object to PassFailExam](#)

202. [Figure 10-18 Rectangle and Cube classes](#)

203. [Figure 10-19 UML diagram for the Student class](#)

204. [Figure 10-20 Realization relationship in a UML diagram](#)

205. [Figure 10-21 Creating an instance of an anonymous inner class](#)

206. [Figure 11-1 Part of the exception class hierarchy](#)

207. [Figure 11-2 Interaction with the OpenFile.java program \(assume that BadFile.txt does not exist\)](#)

208. [Figure 11-3 Interaction with the OpenFile.java program \(assume that GoodFile.txt does exist\)](#)

209. [Figure 11-4 Sequence of events with an exception](#)

210. [Figure 11-5 Sequence of events with no exception](#)

211. [Figure 11-6 Interaction with the ExceptionMessage.java program \(assume that BadFile.txt does not exist\)](#)

212. [Figure 11-7 Dialog box displayed by the SalesReport.java program when no error occurs](#)

213. [Figure 11-8 Dialog box displayed by the SalesReport.java program when the file cannot be found](#)

214. [Figure 11-9 Interaction with the SalesReport2.java program](#)

215. [Figure 11-10 Dialog boxes displayed by the SalesReport2.java program \(Oracle Corporate Counsel\)](#)

216. [Figure 11-11 Inheritance hierarchy for the NumberFormatException class](#)

217. [Figure 11-12 UML diagram for the BankAccount class](#)

218. [Figure 11-13 The number 1297 expressed as characters](#)

219. [Figure 11-14 The number 1297 as a binary number, as it is stored in memory](#)

220. [Figure 11-15 Sequential access versus random access](#)

221. [Figure 11-16 Layout of the Letters.dat file](#)

222. [Figure 12-1 A command line interface](#)

223. [Figure 12-2 A window in a graphical user interface](#)

224. [Figure 12-3 Interaction with a program in a text environment](#)

225. [Figure 12-4 A GUI application](#)

226. [Figure 12-5 Controls in a GUI](#)

227. [Figure 12-6 Nodes in a scene graph](#)

228. [Figure 12-7 Window displayed by MyFirstGUI.java](#)

229. [Figure 12-8 Examples of layout containers](#)

230. [Figure 12-9 Nodes in the scene graph](#)

231. [Figure 12-10 Window displayed by the HelloWorld application](#)



232. [Figure 12-11 Window displayed by the modified HelloWorld application](#)  
233. [Figure 12-12 Window displayed by the HelloWorld2 application](#)  
234. [Figure 12-13 The ImageDemo application](#)  
235. [Figure 12-14 Scene graph for the ImageDemo program in Code Listing 12-4](#)  
236. [Figure 12-15 The HBoxImages application](#)  
237. [Figure 12-16 Scene graph for the HBoxImages program in Code Listing 12-5](#)  
238. [Figure 12-17 Images in an HBox with 10 pixels spacing](#)  
239. [Figure 12-18 The HBoxImagesWithPadding application](#)  
240. [Figure 12-19 The VBoxImagesWithPadding application](#)  
241. [Figure 12-20 Scene graph for the VBoxImagesWithPadding program in Code Listing 12-7](#)  
242. [Figure 12-21 Column and row indexes](#)  
243. [Figure 12-22 The GridPaneDemo application](#)  
244. [Figure 12-23 A GridPane with visible grid lines](#)  
245. [Figure 12-24 The GridPaneImages application](#)  
246. [Figure 12-25 Scene graph for the GridPaneImages program in Code Listing 12-9](#)  
247. [Figure 12-26 A window with nested layout containers](#)  
248. [Figure 12-27 The layout](#)  
249. [Figure 12-28 A window with a Button control](#)  
250. [Figure 12-29 The ButtonDemo application](#)  
251. [Figure 12-30 A Button control firing an\(ActionEvent\)](#)  
252. [Figure 12-31 The EventDemo application](#)  
253. [Figure 12-32 Layout of the controls in the Kilometer Converter application](#)  
254. [Figure 12-33 Scene graph for the Kilometer Converter application](#)  
255. [Figure 12-34 The Kilometer Converter application](#)  
256. [Figure 12-35 Typical GUI organization](#)  
257. [Figure 12-36 The regions of a BorderPane layout container](#)  
258. [Figure 12-37 The BorderPaneDemo1 GUI](#)  
259. [Figure 12-38 The tic-tac-toe application](#)  
260. [Figure 12-39 Slot machine application](#)  
261. [Figure 13-1 Output of CSSDemo1.java](#)  
262. [Figure 13-2 Output of CSSDemo2.java](#)  
263. [Figure 13-3 Output of CSSDemo3.java](#)  
264. [Figure 13-4 Example Label with styles applied](#)  
265. [Figure 13-5 Output of CSSDemo9.java](#)  
266. [Figure 13-6 RadioButton controls](#)  
267. [Figure 13-7 The Metric Converter application](#)  
268. [Figure 13-8 The Metric Converter application with conversions displayed](#)  
269. [Figure 13-9 The RadioButtonEvent application](#)  
270. [Figure 13-10 CheckBox controls](#)  
271. [Figure 13-11 The PizzaToppings application](#)  
272. [Figure 13-12 The PizzaToppings application running](#)  
273. [Figure 13-13 ListView examples](#)  
274. [Figure 13-14 The ListViewDemo1.java application](#)  
275. [Figure 13-15 The ListViewDemo2.java application](#)  
276. [Figure 13-16 The ListViewDemo3.java application](#)

277. [Figure 13-17 Selection modes](#)  
278. [Figure 13-18 The ListViewDemo4.java application](#)  
279. [Figure 13-19 A horizontal ListView](#)  
280. [Figure 13-20 The ListViewDemo5.java application](#)  
281. [Figure 13-21 The ListViewDemo6.java application](#)  
282. [Figure 13-22 The ComboBox control](#)  
283. [Figure 13-23 The ComboBoxDemo1.java application](#)  
284. [Figure 13-24 The ComboBoxDemo2.java application](#)  
285. [Figure 13-25 Uneditable versus Editable ComboBoxes](#)  
286. [Figure 13-26 A horizontal and a vertical Slider](#)  
287. [Figure 13-27 The SliderDemo.java application](#)  
288. [Figure 13-28 The TextArea control](#)  
289. [Figure 13-29 Example menu system](#)  
290. [Figure 13-30 The SimpleMenu.java application](#)  
291. [Figure 13-31 Menu system in the TextMenu.java program](#)  
292. [Figure 13-32 The TextMenu.java application](#)  
293. [Figure 13-33 The File menu with the Alt+F mnemonic](#)  
294. [Figure 13-34 A file chooser dialog box for opening a file](#)  
295. [Figure 13-35 A file chooser dialog box for saving a file](#)  
296. [Figure 13-36 Messages displayed to the console during the application's execution](#)  
297. [Figure 14-1 Various pixel locations in a 640 by 480 window](#)  
298. [Figure 14-2 Output of Triangle.java](#)  
299. [Figure 14-3 Coordinates of the triangle's corners](#)  
300. [Figure 14-4 Output of BullsEye.java](#)  
301. [Figure 14-5 Output of CheckerBoard.java](#)  
302. [Figure 14-6 An ellipse's center point, x-radius, and y-radius](#)  
303. [Figure 14-7 Output of Orbit.java](#)  
304. [Figure 14-8 Arc properties](#)  
305. [Figure 14-9 Types of arcs](#)  
306. [Figure 14-10 Output of PieChart.java](#)  
307. [Figure 14-11 Vertices of the octagon displayed by Octagon.java](#)  
308. [Figure 14-12 Output of Octagon.java](#)  
309. [Figure 14-13 Output of PolylineDemo.java](#)  
310. [Figure 14-14 Output of StopSign.java](#)  
311. [Figure 14-15 Output of RotateBox.java](#)  
312. [Figure 14-16 Output of RotateStopSign.java](#)  
313. [Figure 14-17 Output of ScaleText.java](#)  
314. [Figure 14-18 Output of BallDrop.java](#)  
315. [Figure 14-19 A line rotating from 0 degrees to 90 degrees](#)  
316. [Figure 14-20 A line starting at 90 degrees, rotating by an additional 45 degrees](#)  
317. [Figure 14-21 Output of RotateImage.java](#)  
318. [Figure 14-22 An animation in which text is scaled to 5 times its original size](#)  
319. [Figure 14-23 Fade transition example](#)  
320. [Figure 14-24 Output of DropShadowDemo.java](#)  
321. [Figure 14-25 Drop shadow with X and Y offsets of 10](#)

322. [Figure 14-26 A rectangle with an inner shadow](#)  
323. [Figure 14-27 Comparison of normal text and text with a BoxBlur effect](#)  
324. [Figure 14-28 BlurBox width and height set to 10 pixels](#)  
325. [Figure 14-29 Comparison of normal text and text with a GaussianBlur effect](#)  
326. [Figure 14-30 Comparison of normal text and text with a MotionBlur effect](#)  
327. [Figure 14-31 Motion blur applied to an image](#)  
328. [Figure 14-32 An image without, and with, the Glow effect](#)  
329. [Figure 14-33 The Reflection effect applied to a Text object](#)  
330. [Figure 14-34 Applying DropShadow and Reflection effects to a Text object](#)  
331. [Figure 14-35 Applying DropShadow, sepiaTone, and Reflection effects to an image](#)  
332. [Figure 14-36 Window displayed by SoundPlayer.java](#)  
333. [Figure 14-37 The VideoDemo.java application](#)  
334. [Figure 14-38 The VideoPlayer.java application](#)  
335. [Figure 14-39 The KeyPressedDemo.java application after 5 keys have been pressed](#)  
336. [Figure 14-40 The MoveBall.java application](#)  
337. [Figure 14-41 The MouseEventDemo.java application](#)  
338. [Figure 14-42 The MouseMovedDemo.java application](#)  
339. [Figure 14-43 Hollywood star](#)  
340. [Figure 15-1 First two calls of the method](#)  
341. [Figure 15-2 Total of six calls to the message method](#)  
342. [Figure 15-3 Control returns to the point after the recursive method call](#)  
343. [Figure 15-4 Interaction with the FactorialDemo.java program](#)  
344. [Figure 15-5 Recursive calls to the factorial method \(Oracle Corporate Counsel\)](#)  
345. [Figure 15-6 Circles applet](#)  
346. [Figure 15-7 The pegs and discs in the Towers of Hanoi game](#)  
347. [Figure 15-8 Steps for moving three pegs](#)  
348. [Figure 16-1 An array](#)  
349. [Figure 16-2 Elements 0 and 1 are exchanged](#)  
350. [Figure 16-3 Elements 1 and 2 are exchanged](#)  
351. [Figure 16-4 Elements 4 and 5 are exchanged](#)  
352. [Figure 16-5 Elements 3 and 4 are exchanged](#)  
353. [Figure 16-6 The array after the third pass](#)  
354. [Figure 16-7 The array with all elements sorted](#)  
355. [Figure 16-8 Values in an array](#)  
356. [Figure 16-9 Values in the array after the first swap](#)  
357. [Figure 16-10 Values in the array after the second swap](#)  
358. [Figure 16-11 Values in the array after the third swap](#)  
359. [Figure 16-12 Values in the array after the fourth swap](#)  
360. [Figure 16-13 Values in the array after the fifth swap](#)  
361. [Figure 16-14 An unsorted array](#)  
362. [Figure 16-15 The third element must be moved](#)  
363. [Figure 16-16 The fourth element must be moved](#)  
364. [Figure 16-17 The fifth element must be moved](#)  
365. [Figure 16-18 The sixth element must be moved](#)  
366. [Figure 16-19 All of the elements are in the correct position](#)

367. [Figure 16-20 Sublists and pivot](#)  
368. [Figure 17-1 Passing a type argument to a type parameter](#)  
369. [Figure 17-2 The type parameter T in the Point class](#)  
370. [Figure 17-3 Type parameter defined in a method header](#)  
371. [Figure 17-4 Generic subclass header](#)  
372. [Figure 18-1 The main JCF interfaces](#)  
373. [Figure 19-1 Adding nodes to a linked list](#)  
374. [Figure 19-2 Inserting a node in the middle of a linked list](#)  
375. [Figure 19-3 Removing a node from the middle of a linked list](#)  
376. [Figure 19-4 A graphical interface for testing the linked list class](#)  
377. [Figure 19-5 A node for a doubly linked list](#)  
378. [Figure 19-6 A doubly linked list](#)  
379. [Figure 19-7 The first step in inserting a node into a doubly linked list](#)  
380. [Figure 19-8 The adjacency relation for nodes of a linked list](#)  
381. [Figure 19-9 An example of a digraph](#)  
382. [Figure 20-1 Pushing items onto a stack](#)  
383. [Figure 20-2 A linked stack after pushing two items](#)  
384. [Figure 20-3 The linked stack of Figure 20-2 after push of the string "Carol"](#)  
385. [Figure 20-4 Use of an array-based queue](#)  
386. [Figure 20-5 Use of an array as a circular buffer](#)  
387. [Figure 20-6 Screenshot of GUI front end for the queue program](#)  
388. [Figure 20-7 A linked list representing a queue with three elements](#)  
389. [Figure 21-1 A binary tree](#)  
390. [Figure 21-2 Left and right subtrees](#)  
391. [Figure 21-3 Binary tree representation of expressions](#)  
392. [Figure 21-4 A binary tree constructed with the Node class](#)  
393. [Figure 21-5 Partitioning a view to display a binary tree](#)  
394. [Figure 21-6 GUI display of a binary tree](#)  
395. [Figure 21-7 A binary search tree](#)  
396. [Figure 21-8 Before and after removal of a node with one child](#)  
397. [Figure 21-9 Removing a node with two children](#)  
398. [Figure 21-10 The tree of Figure 21-9 after removing 90](#)  
399. [Figure 21-11 A binary tree after adding some values](#)  
400. [Figure 21-12 The binary tree of Figure 21-11 after removing 100](#)  
401. [Figure 21-13 Examples of AVL trees](#)  
402. [Figure 21-14 An example of a non-AVL tree](#)  
403. [Figure 21-15 A single right rotation is performed at a node with an LL imbalance](#)  
404. [Figure 21-16 A double LR rotation is performed at a node with an LR imbalance](#)  
405. [Figure 21-17 A node with an LL imbalance](#)  
406. [Figure 21-18 Effect of a single right rotation](#)  
407. [Figure 21-19 A node with an LR imbalance](#)  
408. [Figure 21-20 Effect of a double LR rotation on the tree of Figure 21-19](#)  
409. [Figure 21-21 Sample interaction with AVL program](#)  
410. [Figure 21-22 A binary tree with the heap order property](#)  
411. [Figure 21-23 A complete binary tree that is also a heap](#)

- 112. [Figure 21-24 Storing a complete binary tree in an array](#)
- 113. [Figure 21-25 Adding a new value may violate the heap order property](#)
- 114. [Figure 21-26 The heap of Figure 21-23 after replacing the root with the deepest rightmost leaf](#)
- 115. [Figure 21-27 A rooted tree](#)
- 116. [Figure 21-28 Screen shot of the TreeView example program](#)
- 117. [Figure 22-1 A Java application interacts with a DBMS, which manipulates data](#)
- 118. [Figure 22-2 A Java application uses the JDBC classes to interact with a DBMS](#)
- 119. [Figure 22-3 The Coffee database table](#)
- 120. [Figure 22-5 A ResultSet object contains the results of an SQL query](#)
- 121. [Figure 22-6 ResultSet rows and columns](#)
- 122. [Figure 22-7 The cursor is initially positioned before the first row](#)
- 123. [Figure 22-8 The next method moves the cursor forward](#)
- 124. [Figure 22-16 The DBViewer application](#)
- 125. [Figure 22-17 Entity relationship diagram](#)
- 126. [Figure 22-18 Customer table displayed in the DBViewer application](#)
- 127. [Figure 22-19 Order data entered](#)
- 128. [Figure 22-20 Order information viewed in DBViewer](#)
- 129. [Figure 23-1 Various GUI components](#)
- 130. [Figure 23-2 Window displayed by ShowWindow.java](#)
- 131. [Figure 23-3 Window displayed by the EmbeddedMain class](#)
- 132. [Figure 23-4 Sketch of the Kilometer Converter window](#)
- 133. [Figure 23-5 A panel is added to the content pane](#)
- 134. [Figure 23-6 Kilometer Converter window](#)
- 135. [Figure 23-7 Components in the Kilometer Converter window](#)
- 136. [Figure 23-8 A class with an inner class](#)
- 137. [Figure 23-9 A JButton component firing an action event](#)
- 138. [Figure 23-10 Windows displayed by the KiloConverter class](#)
- 139. [Figure 23-11 The window produced by the ColorWindow class](#)
- 140. [Figure 23-12 Output of EventObjectWindow class](#)
- 141. [Figure 23-13 The window displayed by the FlowWindow class](#)
- 142. [Figure 23-14 The arrangements of the buttons after resizing](#)
- 143. [Figure 23-15 Left, center, and right alignment](#)
- 144. [Figure 23-16 The regions of a BorderLayout manager](#)
- 145. [Figure 23-17 The window displayed by the BorderWindow class](#)
- 146. [Figure 23-18 The window resized](#)
- 147. [Figure 23-19 Nesting JPanel objects inside each region](#)
- 148. [Figure 23-20 Window displayed by the BorderPanelWindow class](#)
- 149. [Figure 23-21 The GridLayout manager divides a container into cells](#)
- 150. [Figure 23-22 Window displayed by the GridWindow class](#)
- 151. [Figure 23-23 Window displayed by the GridPanelWindow class](#)
- 152. [Figure 23-24 Radio buttons](#)
- 153. [Figure 23-25 Metric Converter window](#)
- 154. [Figure 23-26 Window and dialog boxes displayed by the MetricConverter class](#)
- 155. [Figure 23-27 Check boxes](#)
- 156. [Figure 23-28 Window displayed by the ColorCheckBoxWindow class](#)



- 157. [Figure 23-29 A group of check boxes with a titled border](#)
- 158. [Figure 23-30 Sketch of the Order Calculator window](#)
- 159. [Figure 23-31 Appearance of the GreetingPanel component](#)
- 160. [Figure 23-32 Appearance of the BagelPanel component](#)
- 161. [Figure 23-33 Appearance of the ToppingPanel component](#)
- 162. [Figure 23-34 Appearance of the CoffeePanel component](#)
- 163. [Figure 23-35 Placement of the custom panels](#)
- 164. [Figure 23-36 The Order Calculator window](#)
- 165. [Figure 23-37 The subtotal, tax, and total displayed](#)
- 166. [Figure 23-38 Splash screen for the Brandi's Bagel House application](#)
- 167. [Figure 23-39 Messages displayed to the console during the application's execution](#)
- 168. [Figure 24-1 Part of the Swing and AWT class hierarchy](#)
- 169. [Figure 24-2 A window with three read-only text fields](#)
- 170. [Figure 24-3 Read-only text fields with data displayed](#)
- 171. [Figure 24-4 A JList component](#)
- 172. [Figure 24-5 Selection modes](#)
- 173. [Figure 24-6 Window displayed by the ListWindow class](#)
- 174. [Figure 24-7 List with a line border](#)
- 175. [Figure 24-8 List component with a scroll bar](#)
- 176. [Figure 24-9 List component with scroll bars](#)
- 177. [Figure 24-10 The window displayed by the MultipleIntervalSelection class](#)
- 178. [Figure 24-11 A combo box](#)
- 179. [Figure 24-12 The window displayed by the ComboBoxWindow class](#)
- 180. [Figure 24-13 An editable combo box](#)
- 181. [Figure 24-14 Labels displaying an image icon](#)
- 182. [Figure 24-15 Buttons displaying an image icon](#)
- 183. [Figure 24-16 Window displayed by the MyCatImage class](#)
- 184. [Figure 24-17 Button with mnemonic X](#)
- 185. [Figure 24-18 Radio buttons and check boxes with mnemonics assigned](#)
- 186. [Figure 24-19 Button with tool tip displayed](#)
- 187. [Figure 24-20 A file chooser dialog box for opening a file](#)
- 188. [Figure 24-21 A save file dialog box](#)
- 189. [Figure 24-22 A color chooser dialog box](#)
- 190. [Figure 24-23 Example menu system](#)
- 191. [Figure 24-24 Window displayed by the MenuWindow class](#)
- 192. [Figure 24-25 Sketch of the MenuWindow class's menu system](#)
- 193. [Figure 24-26 The window with the File menu and Text menu opened](#)
- 194. [Figure 24-27 Text areas with and without scroll bars](#)
- 195. [Figure 24-28 Examples of fonts](#)
- 196. [Figure 24-29 Examples of serif plain, bold, italic, and bold plus italic](#)
- 197. [Figure 24-30 A horizontal and a vertical slider](#)
- 198. [Figure 24-31 Window displayed by the TempConverterWindow class](#)
- 199. [Figure 24-32 Metal, Motif, and Windows looks and feels](#)
- 500. [Figure 24-33 Dice simulator](#)
- 501. [Figure 24-34 Card dealer](#)

502. [Figure 24-35 The Tic Tac Toe application](#)

503. [Figure 25-1 Applets are transmitted along with Web pages](#)

504. [Figure 25-2 A Web page with an applet](#)

505. [Figure 25-3 Web page with a title only](#)

506. [Figure 25-4 Web page produced by BasicWebPage2.html](#)

507. [Figure 25-5 Header levels](#)

508. [Figure 25-6 Web page produced by BasicWebPage3.html](#)

509. [Figure 25-7 Line breaks in an HTML document](#)

510. [Figure 25-8 Paragraph breaks in an HTML document](#)

511. [Figure 25-9 Horizontal rules in a Web page](#)

512. [Figure 25-10 Web page produced by BasicWebPage4.html](#)

513. [Figure 25-11 Web page produced by LinkDemo.html](#)

514. [Figure 25-12 The Web page produced by SimpleApplet.html](#)

515. [Figure 25-13 Security warning in Internet Explorer](#)

516. [Figure 25-14 Applet executed by appletviewer](#)

517. [Figure 25-15 Web page produced by TempConverter.html](#)

518. [Figure 25-16 Web page running the AWTTempConverter applet](#)

519. [Figure 25-17 X and Y coordinates on a 300 pixel wide by 200 pixel high component](#)

520. [Figure 25-18 LineDemo applet](#)

521. [Figure 25-19 RectangleDemo applet](#)

522. [Figure 25-20 An oval and its enclosing rectangle](#)

523. [Figure 25-21 OvalDemo applet](#)

524. [Figure 25-22 An arc](#)

525. [Figure 25-23 ArcDemo applet](#)

526. [Figure 25-24 Points of each vertex in an octagon](#)

527. [Figure 25-25 PolygonDemo applet](#)

528. [Figure 25-26 GraphicStringDemo applet](#)

529. [Figure 25-27 GraphicsWindow applet](#)

530. [Figure 25-28 GraphicsWindow applet with all graphics selected](#)

531. [Figure 25-29 MouseEvents applet](#)

532. [Figure 25-30 DrawBoxes applet](#)

533. [Figure 25-31 BouncingBall applet](#)

534. [Figure 25-32 AudioDemo2 applet](#)

535. [Figure 25-33 House drawing](#)

536. [Figure 25-34 Eyes following the mouse cursor](#)

537. [Figure D-1 The Java SE documentation site](#)

538. [Figure D-2 Java SE API documentation](#)

539. [Figure E-1 BankAccount class documentation \(Oracle Corporate Counsel\)](#)

540. [Figure E-2 A documentation file for multiple classes \(Oracle Corporate Counsel\)](#)

541. [Figure G-1 Example directory structure for storing a package on windows](#)

542. [Figure G-2 employees directory added to the mypackages directory](#)

543. [Figure G-3 Directory structure showing the contents of the hourly directory](#)

544. [Figure I-1 The Apache Derby Download Page](#)

545. [Figure I-2 Select the latest bin distribution](#)

546. [Figure J-1 Sub lists and pivot](#)

- 547. [Figure CS1-1 Interaction with the HalsCommission program](#)
- 548. [Figure CS2-1 UML diagram for the Amortization class](#)
- 549. [Figure CS2-2 Interaction with the LoanReport program](#)
- 550. [Figure CS4-1 The relationship between the names and addresses arrays](#)
- 551. [Figure CS5-1 UML diagram for the FeetInches class](#)
- 552. [Figure CS5-2 The if statement tests the contents of the reference variables, not the contents of the objects the variables reference](#)
- 553. [Figure CS5-3 Both variables reference the same object](#)
- 554. [Figure CS6-1 UML diagram for the SerialNumber class](#)
- 555. [Figure CS7-1 The text editor window](#)
- 556. [Figure CS7-2 Menu system for the text editor application](#)
- 557. [Figure CS7-3 Text displayed in various fonts and styles](#)

## List of Tables

- 1. [Table 1-1 Programming languages](#)
- 2. [Table 1-2 The common elements of a programming language](#)
- 3. [Table 1-3 The Java key words](#)
- 4. [Table 2-1 Special characters](#)
- 5. [Table 2-2 Common escape sequences](#)
- 6. [Table 2-3 Literals](#)
- 7. [Table 2-4 Some variable names](#)
- 8. [Table 2-5 Primitive data types for numeric data](#)
- 9. [Table 2-6 Floating-point representations](#)
- 10. [Table 2-7 Arithmetic operators](#)
- 11. [Table 2-8 Precedence of arithmetic operators \(highest to lowest\)](#)
- 12. [Table 2-9 Associativity of arithmetic operators](#)
- 13. [Table 2-10 Some expressions and their values](#)
- 14. [Table 2-11 More expressions and their values](#)
- 15. [Table 2-12 Various assignment statements \(assume x = 6 in each statement\)](#)
- 16. [Table 2-13 Combined assignment operators](#)
- 17. [Table 2-14 Example uses of cast operators](#)
- 18. [Table 2-15 A few String class methods](#)
- 19. [Table 2-16 Block comments](#)
- 20. [Table 2-17 Some of the Scanner class methods](#)
- 21. [Table 2-18 Methods for converting strings to numbers](#)
- 22. [Table 3-1 Relational operators](#)
- 23. [Table 3-2 boolean expressions using relational operators](#)
- 24. [Table 3-3 Other examples of if statements](#)
- 25. [Table 3-4 Logical operators](#)
- 26. [Table 3-5 boolean expressions using logical operators](#)
- 27. [Table 3-6 Truth table for the && operator](#)
- 28. [Table 3-7 Truth table for the || operator](#)
- 29. [Table 3-8 Truth table for the ! operator](#)



30. [Table 3-9 Logical operators in order of precedence](#)
31. [Table 3-10 Precedence of all operators discussed so far](#)
32. [Table 4-1 Some of the Random class's methods](#)
33. [Table 6-1 Summary of the private and public access specifiers for class members](#)
34. [Table 6-2 A few of the standard Java packages](#)
35. [Table 8-1 The Stock class methods](#)
36. [Table 9-1 Some static Character class methods for testing char values](#)
37. [Table 9-2 Some Character class methods for case conversion](#)
38. [Table 9-3 String methods that search for a substring](#)
39. [Table 9-4 String methods for getting a character or substring's location](#)
40. [Table 9-5 String methods for extracting substrings](#)
41. [Table 9-6 Methods that return a modified copy of a String object](#)
42. [Table 9-7 Some of the String class's valueOf methods](#)
43. [Table 9-8 StringBuilder constructors](#)
44. [Table 9-9 Methods that are common to the String and StringBuilder classes](#)
45. [Table 9-10 The StringBuilder class's delete, deleteCharAt, and setCharAt methods](#)
46. [Table 9-11 Wrapper classes for the numeric primitive data types](#)
47. [Table 9-12 Morse code](#)
48. [Table 10-1 CurvedActivity class fields](#)
49. [Table 10-2 CurvedActivity class methods](#)
50. [Table 10-3 Accessibility from within the class's package](#)
51. [Table 10-4 Accessibility from outside the class's package](#)
52. [Table 11-1 Some of the DataOutputStream methods](#)
53. [Table 11-2 Some of the DataInputStream methods](#)
54. [Table 12-1 Three of the most simple layout containers](#)
55. [Table 12-2 enum constants of the Pos type](#)
56. [Table 12-3 BorderPane constructors](#)
57. [Table 12-4 A few of the ObservableList methods](#)
58. [Table 13-1 Some of the JavaFX nodes and their corresponding CSS type selectors](#)
59. [Table 13-2 A few CSS style properties for the Labeled class \(inherited by the Label and Button classes\)](#)
60. [Table 13-3 A few CSS style properties for the Region class \(inherited by the Label and Button classes\)](#)
61. [Table 13-4 A few of the named colors](#)
62. [Table 13-5 ListView methods to select items](#)
63. [Table 13-6 Some of the ComboBox methods](#)
64. [Table 13-7 Some of the Slider class methods](#)
65. [Table 13-8 Some of the TextArea class methods](#)
66. [Table 13-9 Skateboard products](#)
67. [Table 13-10 Optional preconference workshops](#)
68. [Table 14-1 Some of the Shape subclasses](#)
69. [Table 14-2 The Line class constructors](#)
70. [Table 14-3 The Circle class constructors](#)
71. [Table 14-4 The Rectangle class constructors](#)
72. [Table 14-5 The Ellipse class constructors](#)

- 73. [Table 14-6 The Arc class constructors](#)
- 74. [Table 14-7 Arc types](#)
- 75. [Table 14-8 The Polygon class constructors](#)
- 76. [Table 14-9 The PolyLine class constructors](#)
- 77. [Table 14-10 The Text class constructors](#)
- 78. [Table 14-11 Some of the transition classes](#)
- 79. [Table 14-12 TranslateTransition class constructors](#)
- 80. [Table 14-13 RotateTransition class constructors](#)
- 81. [Table 14-14 ScaleTransition class constructors](#)
- 82. [Table 14-15 StrokeTransition class constructors](#)
- 83. [Table 14-16 FillTransition class constructors](#)
- 84. [Table 14-17 FadeTransition class constructors](#)
- 85. [Table 14-18 A few methods inherited from the Animation class](#)
- 86. [Table 14-19 Interpolator constants](#)
- 87. [Table 14-20 Some of the effect classes](#)
- 88. [Table 14-21 Some of the DropShadow class methods](#)
- 89. [Table 14-22 Some of the ColorAdjust class methods](#)
- 90. [Table 14-23 Some of the Reflection class methods](#)
- 91. [Table 14-24 Some of the MediaPlayer methods](#)
- 92. [Table 14-25 Mouse Event Types](#)
- 93. [Table 17-1 Common Type Parameter Names](#)
- 94. [Table 18-1 Methods of the Iterator<E> interface](#)
- 95. [Table 18-2 Methods of the Collection<E> interface](#)
- 96. [Table 18-3 Methods of the List<E> interface](#)
- 97. [Table 18-4 Methods of the ListIterator<E> interface](#)
- 98. [Table 18-5 Methods of the LinkedList<E> class](#)
- 99. [Table 18-6 HashSet constructors](#)
- 100. [Table 18-7 SortedSet<E> methods](#)
- 101. [Table 18-8 Abstract methods of the Comparator<T> interface](#)
- 102. [Table 18-9 TreeSet<E> constructors](#)
- 103. [Table 18-10 Methods specified by the Map<K,V> interface](#)
- 104. [Table 18-11 HashMap<K, V> constructors](#)
- 105. [Table 18-12 LinkedHashMap constructors](#)
- 106. [Table 18-13 TreeMap<K, V> constructors](#)
- 107. [Table 18-14 Some of the static methods in the Collections class](#)
- 108. [Table 18-15 Stream reduction operators](#)
- 109. [Table 18-16 Stream<T> intermediate operators](#)
- 110. [Table 19-1 Representing a digraph by adjacency lists](#)
- 111. [Table 19-2 An adjacency list using integers to represent digraph nodes](#)
- 112. [Table 19-3 Storing digraph information in a file to simplify parsing](#)
- 113. [Table 20-1 Stack class methods](#)
- 114. [Table 20-2 Some methods in the Queue interface](#)
- 115. [Table 20-3 Some methods of the File class](#)
- 116. [Table 21-1 PriorityQueue class methods](#)
- 117. [Table 22-1 The Coffee database table](#)

- 118. [Table 22-2 A few of the SQL data types](#)
- 119. [Table 22-3 A few of the ResultSet methods](#)
- 120. [Table 22-4 SQL relational operators](#)
- 121. [Table 22-5 Rows inserted into the Customer table](#)
- 122. [Table 22-6 A few ResultSetMetaData methods](#)
- 123. [Table 23-1 Some GUI components](#)
- 124. [Table 23-2 Label, text field, and button controls](#)
- 125. [Table 23-3 Color constants](#)
- 126. [Table 23-4 ActionEvent methods](#)
- 127. [Table 23-5 Layout managers](#)
- 128. [Table 23-6 Borders produced by the BorderLayout class](#)
- 129. [Table 24-1 Look and feel class names](#)
- 130. [Table 24-2 Skateboard products](#)
- 131. [Table 24-3 Optional preconference workshops](#)
- 132. [Table 25-1 Several AWT classes](#)
- 133. [Table 25-2 Some of the Graphics class methods](#)
- 134. [Table 25-3 Methods required by the MouseListener interface](#)
- 135. [Table 25-4 Methods required by the MouseMotionListener interface](#)
- 136. [Table 25-5 Timer class methods](#)
- 137. [Table F-1 Several Math class methods](#)
- 138. [Table F-2 The E and PI constants](#)
- 139. [Table G-1 The effect of access specifiers](#)
- 140. [Table G-2 A few of the standard Java packages](#)
- 141. [Table CS1-1 Sales commission rates](#)
- 142. [Table CS1-2 Variables](#)
- 143. [Table CS2-1 Amortization class fields](#)
- 144. [Table CS2-2 Amortization class methods](#)
- 145. [Table CS3-1 PinTester class fields](#)
- 146. [Table CS3-2 PinTester class methods](#)
- 147. [Table CS5-1 The FeetInches class methods](#)
- 148. [Table CS6-1 SerialNumber class's methods](#)

# Landmarks

- 1. [Contents in Brief](#)
- 2. [Frontmatter](#)
- 3. [Start of Content](#)
- 4. [backmatter](#)
- 5. [List of Illustrations](#)
- 6. [List of Tables](#)

- 1. [i](#)
- 2. [ii](#)
- 3. [iii](#)

- 4. [iv](#)
- 5. [v](#)
- 6. [vi](#)
- 7. [vii](#)
- 8. [viii](#)
- 9. [ix](#)
- 10. [x](#)
- 11. [xi](#)
- 12. [xii](#)
- 13. [xiii](#)
- 14. [xiv](#)
- 15. [xv](#)
- 16. [xvi](#)
- 17. [xvii](#)
- 18. [xviii](#)
- 19. [xix](#)
- 20. [xx](#)
- 21. [xxi](#)
- 22. [xxii](#)
- 23. [xxiii](#)
- 24. [xxiv](#)
- 25. [xxv](#)
- 26. [xxvi](#)
- 27. [xxvii](#)
- 28. [xxviii](#)
- 29. [xxix](#)
- 30. [xxx](#)
- 31. [xxxi](#)
- 32. [xxxii](#)
- 33. [xxxiii](#)
- 34. [xxxiv](#)
- 35. [1](#)
- 36. [2](#)
- 37. [3](#)
- 38. [4](#)
- 39. [5](#)
- 40. [6](#)
- 41. [7](#)
- 42. [8](#)
- 43. [9](#)
- 44. [10](#)
- 45. [11](#)
- 46. [12](#)
- 47. [13](#)
- 48. [14](#)

- 49. [15](#)
- 50. [16](#)
- 51. [17](#)
- 52. [18](#)
- 53. [19](#)
- 54. [20](#)
- 55. [21](#)
- 56. [22](#)
- 57. [23](#)
- 58. [24](#)
- 59. [25](#)
- 60. [26](#)
- 61. [27](#)
- 62. [28](#)
- 63. [29](#)
- 64. [30](#)
- 65. [31](#)
- 66. [32](#)
- 67. [33](#)
- 68. [34](#)
- 69. [35](#)
- 70. [36](#)
- 71. [37](#)
- 72. [38](#)
- 73. [39](#)
- 74. [40](#)
- 75. [41](#)
- 76. [42](#)
- 77. [43](#)
- 78. [44](#)
- 79. [45](#)
- 80. [46](#)
- 81. [47](#)
- 82. [48](#)
- 83. [49](#)
- 84. [50](#)
- 85. [51](#)
- 86. [52](#)
- 87. [53](#)
- 88. [54](#)
- 89. [55](#)
- 90. [56](#)
- 91. [57](#)
- 92. [58](#)
- 93. [59](#)

94. [60](#)  
95. [61](#)  
96. [62](#)  
97. [63](#)  
98. [64](#)  
99. [65](#)  
100. [66](#)  
101. [67](#)  
102. [68](#)  
103. [69](#)  
104. [70](#)  
105. [71](#)  
106. [72](#)  
107. [73](#)  
108. [74](#)  
109. [75](#)  
110. [76](#)  
111. [77](#)  
112. [78](#)  
113. [79](#)  
114. [80](#)  
115. [81](#)  
116. [82](#)  
117. [83](#)  
118. [84](#)  
119. [85](#)  
120. [86](#)  
121. [87](#)  
122. [88](#)  
123. [89](#)  
124. [90](#)  
125. [91](#)  
126. [92](#)  
127. [93](#)  
128. [94](#)  
129. [95](#)  
130. [96](#)  
131. [97](#)  
132. [98](#)  
133. [99](#)  
134. [100](#)  
135. [101](#)  
136. [102](#)  
137. [103](#)  
138. [104](#)

139. [105](#)  
140. [106](#)  
141. [107](#)  
142. [108](#)  
143. [109](#)  
144. [110](#)  
145. [111](#)  
146. [112](#)  
147. [113](#)  
148. [114](#)  
149. [115](#)  
150. [116](#)  
151. [117](#)  
152. [118](#)  
153. [119](#)  
154. [120](#)  
155. [121](#)  
156. [122](#)  
157. [123](#)  
158. [124](#)  
159. [125](#)  
160. [126](#)  
161. [127](#)  
162. [128](#)  
163. [129](#)  
164. [130](#)  
165. [131](#)  
166. [132](#)  
167. [133](#)  
168. [134](#)  
169. [135](#)  
170. [136](#)  
171. [137](#)  
172. [138](#)  
173. [139](#)  
174. [140](#)  
175. [141](#)  
176. [142](#)  
177. [143](#)  
178. [144](#)  
179. [145](#)  
180. [146](#)  
181. [147](#)  
182. [148](#)  
183. [149](#)

184. [150](#)  
185. [151](#)  
186. [152](#)  
187. [153](#)  
188. [154](#)  
189. [155](#)  
190. [156](#)  
191. [157](#)  
192. [158](#)  
193. [159](#)  
194. [160](#)  
195. [161](#)  
196. [162](#)  
197. [163](#)  
198. [164](#)  
199. [165](#)  
200. [166](#)  
201. [167](#)  
202. [168](#)  
203. [169](#)  
204. [170](#)  
205. [171](#)  
206. [172](#)  
207. [173](#)  
208. [174](#)  
209. [175](#)  
210. [176](#)  
211. [177](#)  
212. [178](#)  
213. [179](#)  
214. [180](#)  
215. [181](#)  
216. [182](#)  
217. [183](#)  
218. [184](#)  
219. [185](#)  
220. [186](#)  
221. [187](#)  
222. [188](#)  
223. [189](#)  
224. [190](#)  
225. [191](#)  
226. [192](#)  
227. [193](#)  
228. [194](#)



229. [195](#)  
230. [196](#)  
231. [197](#)  
232. [198](#)  
233. [199](#)  
234. [200](#)  
235. [201](#)  
236. [202](#)  
237. [203](#)  
238. [204](#)  
239. [205](#)  
240. [206](#)  
241. [207](#)  
242. [208](#)  
243. [209](#)  
244. [210](#)  
245. [211](#)  
246. [212](#)  
247. [213](#)  
248. [214](#)  
249. [215](#)  
250. [216](#)  
251. [217](#)  
252. [218](#)  
253. [219](#)  
254. [220](#)  
255. [221](#)  
256. [222](#)  
257. [223](#)  
258. [224](#)  
259. [225](#)  
260. [226](#)  
261. [227](#)  
262. [228](#)  
263. [229](#)  
264. [230](#)  
265. [231](#)  
266. [232](#)  
267. [233](#)  
268. [234](#)  
269. [235](#)  
270. [236](#)  
271. [237](#)  
272. [238](#)  
273. [239](#)

274. [240](#)  
275. [241](#)  
276. [242](#)  
277. [243](#)  
278. [244](#)  
279. [245](#)  
280. [246](#)  
281. [247](#)  
282. [248](#)  
283. [249](#)  
284. [250](#)  
285. [251](#)  
286. [252](#)  
287. [253](#)  
288. [254](#)  
289. [255](#)  
290. [256](#)  
291. [257](#)  
292. [258](#)  
293. [259](#)  
294. [260](#)  
295. [261](#)  
296. [262](#)  
297. [263](#)  
298. [264](#)  
299. [265](#)  
300. [266](#)  
301. [267](#)  
302. [268](#)  
303. [269](#)  
304. [270](#)  
305. [271](#)  
306. [272](#)  
307. [273](#)  
308. [274](#)  
309. [275](#)  
310. [276](#)  
311. [277](#)  
312. [278](#)  
313. [279](#)  
314. [280](#)  
315. [281](#)  
316. [282](#)  
317. [283](#)  
318. [284](#)

319. [285](#)  
320. [286](#)  
321. [287](#)  
322. [288](#)  
323. [289](#)  
324. [290](#)  
325. [291](#)  
326. [292](#)  
327. [293](#)  
328. [294](#)  
329. [295](#)  
330. [296](#)  
331. [297](#)  
332. [298](#)  
333. [299](#)  
334. [300](#)  
335. [301](#)  
336. [302](#)  
337. [303](#)  
338. [304](#)  
339. [305](#)  
340. [306](#)  
341. [307](#)  
342. [308](#)  
343. [309](#)  
344. [310](#)  
345. [311](#)  
346. [312](#)  
347. [313](#)  
348. [314](#)  
349. [315](#)  
350. [316](#)  
351. [317](#)  
352. [318](#)  
353. [319](#)  
354. [320](#)  
355. [321](#)  
356. [322](#)  
357. [323](#)  
358. [324](#)  
359. [325](#)  
360. [326](#)  
361. [327](#)  
362. [328](#)  
363. [329](#)

364. [330](#)  
365. [331](#)  
366. [332](#)  
367. [333](#)  
368. [334](#)  
369. [335](#)  
370. [336](#)  
371. [337](#)  
372. [338](#)  
373. [339](#)  
374. [340](#)  
375. [341](#)  
376. [342](#)  
377. [343](#)  
378. [344](#)  
379. [345](#)  
380. [346](#)  
381. [347](#)  
382. [348](#)  
383. [349](#)  
384. [350](#)  
385. [351](#)  
386. [352](#)  
387. [353](#)  
388. [354](#)  
389. [355](#)  
390. [356](#)  
391. [357](#)  
392. [358](#)  
393. [359](#)  
394. [360](#)  
395. [361](#)  
396. [362](#)  
397. [363](#)  
398. [364](#)  
399. [365](#)  
400. [366](#)  
401. [367](#)  
402. [368](#)  
403. [369](#)  
404. [370](#)  
405. [371](#)  
406. [372](#)  
407. [373](#)  
408. [374](#)

109. [375](#)  
110. [376](#)  
111. [377](#)  
112. [378](#)  
113. [379](#)  
114. [380](#)  
115. [381](#)  
116. [382](#)  
117. [383](#)  
118. [384](#)  
119. [385](#)  
120. [386](#)  
121. [387](#)  
122. [388](#)  
123. [389](#)  
124. [390](#)  
125. [391](#)  
126. [392](#)  
127. [393](#)  
128. [394](#)  
129. [395](#)  
130. [396](#)  
131. [397](#)  
132. [398](#)  
133. [399](#)  
134. [400](#)  
135. [401](#)  
136. [402](#)  
137. [403](#)  
138. [404](#)  
139. [405](#)  
140. [406](#)  
141. [407](#)  
142. [408](#)  
143. [409](#)  
144. [410](#)  
145. [411](#)  
146. [412](#)  
147. [413](#)  
148. [414](#)  
149. [415](#)  
150. [416](#)  
151. [417](#)  
152. [418](#)  
153. [419](#)

154. [420](#)  
155. [421](#)  
156. [422](#)  
157. [423](#)  
158. [424](#)  
159. [425](#)  
160. [426](#)  
161. [427](#)  
162. [428](#)  
163. [429](#)  
164. [430](#)  
165. [431](#)  
166. [432](#)  
167. [433](#)  
168. [434](#)  
169. [435](#)  
170. [436](#)  
171. [437](#)  
172. [438](#)  
173. [439](#)  
174. [440](#)  
175. [441](#)  
176. [442](#)  
177. [443](#)  
178. [444](#)  
179. [445](#)  
180. [446](#)  
181. [447](#)  
182. [448](#)  
183. [449](#)  
184. [450](#)  
185. [451](#)  
186. [452](#)  
187. [453](#)  
188. [454](#)  
189. [455](#)  
190. [456](#)  
191. [457](#)  
192. [458](#)  
193. [459](#)  
194. [460](#)  
195. [461](#)  
196. [462](#)  
197. [463](#)  
198. [464](#)

499. [465](#)  
500. [466](#)  
501. [467](#)  
502. [468](#)  
503. [469](#)  
504. [470](#)  
505. [471](#)  
506. [472](#)  
507. [473](#)  
508. [474](#)  
509. [475](#)  
510. [476](#)  
511. [477](#)  
512. [478](#)  
513. [479](#)  
514. [480](#)  
515. [481](#)  
516. [482](#)  
517. [483](#)  
518. [484](#)  
519. [485](#)  
520. [486](#)  
521. [487](#)  
522. [488](#)  
523. [489](#)  
524. [490](#)  
525. [491](#)  
526. [492](#)  
527. [493](#)  
528. [494](#)  
529. [495](#)  
530. [496](#)  
531. [497](#)  
532. [498](#)  
533. [499](#)  
534. [500](#)  
535. [501](#)  
536. [502](#)  
537. [503](#)  
538. [504](#)  
539. [505](#)  
540. [506](#)  
541. [507](#)  
542. [508](#)  
543. [509](#)

544. [510](#)  
545. [511](#)  
546. [512](#)  
547. [513](#)  
548. [514](#)  
549. [515](#)  
550. [516](#)  
551. [517](#)  
552. [518](#)  
553. [519](#)  
554. [520](#)  
555. [521](#)  
556. [522](#)  
557. [523](#)  
558. [524](#)  
559. [525](#)  
560. [526](#)  
561. [527](#)  
562. [528](#)  
563. [529](#)  
564. [530](#)  
565. [531](#)  
566. [532](#)  
567. [533](#)  
568. [534](#)  
569. [535](#)  
570. [536](#)  
571. [537](#)  
572. [538](#)  
573. [539](#)  
574. [540](#)  
575. [541](#)  
576. [542](#)  
577. [543](#)  
578. [544](#)  
579. [545](#)  
580. [546](#)  
581. [547](#)  
582. [548](#)  
583. [549](#)  
584. [550](#)  
585. [551](#)  
586. [552](#)  
587. [553](#)  
588. [554](#)



589. [555](#)  
590. [556](#)  
591. [557](#)  
592. [558](#)  
593. [559](#)  
594. [560](#)  
595. [561](#)  
596. [562](#)  
597. [563](#)  
598. [564](#)  
599. [565](#)  
500. [566](#)  
501. [567](#)  
502. [568](#)  
503. [569](#)  
504. [570](#)  
505. [571](#)  
506. [572](#)  
507. [573](#)  
508. [574](#)  
509. [575](#)  
510. [576](#)  
511. [577](#)  
512. [578](#)  
513. [579](#)  
514. [580](#)  
515. [581](#)  
516. [582](#)  
517. [583](#)  
518. [584](#)  
519. [585](#)  
520. [586](#)  
521. [587](#)  
522. [588](#)  
523. [589](#)  
524. [590](#)  
525. [591](#)  
526. [592](#)  
527. [593](#)  
528. [594](#)  
529. [595](#)  
530. [596](#)  
531. [597](#)  
532. [598](#)  
533. [599](#)

534. [600](#)  
535. [601](#)  
536. [602](#)  
537. [603](#)  
538. [604](#)  
539. [605](#)  
540. [606](#)  
541. [607](#)  
542. [608](#)  
543. [609](#)  
544. [610](#)  
545. [611](#)  
546. [612](#)  
547. [613](#)  
548. [614](#)  
549. [615](#)  
550. [616](#)  
551. [617](#)  
552. [618](#)  
553. [619](#)  
554. [620](#)  
555. [621](#)  
556. [622](#)  
557. [623](#)  
558. [624](#)  
559. [625](#)  
560. [626](#)  
561. [627](#)  
562. [628](#)  
563. [629](#)  
564. [630](#)  
565. [631](#)  
566. [632](#)  
567. [633](#)  
568. [634](#)  
569. [635](#)  
570. [636](#)  
571. [637](#)  
572. [638](#)  
573. [639](#)  
574. [640](#)  
575. [641](#)  
576. [642](#)  
577. [643](#)  
578. [644](#)

579. [645](#)  
580. [646](#)  
581. [647](#)  
582. [648](#)  
583. [649](#)  
584. [650](#)  
585. [651](#)  
586. [652](#)  
587. [653](#)  
588. [654](#)  
589. [655](#)  
590. [656](#)  
591. [657](#)  
592. [658](#)  
593. [659](#)  
594. [660](#)  
595. [661](#)  
596. [662](#)  
597. [663](#)  
598. [664](#)  
599. [665](#)  
700. [666](#)  
701. [667](#)  
702. [668](#)  
703. [669](#)  
704. [670](#)  
705. [671](#)  
706. [672](#)  
707. [673](#)  
708. [674](#)  
709. [675](#)  
710. [676](#)  
711. [677](#)  
712. [678](#)  
713. [679](#)  
714. [680](#)  
715. [681](#)  
716. [682](#)  
717. [683](#)  
718. [684](#)  
719. [685](#)  
720. [686](#)  
721. [687](#)  
722. [688](#)  
723. [689](#)

724. [690](#)  
725. [691](#)  
726. [692](#)  
727. [693](#)  
728. [694](#)  
729. [695](#)  
730. [696](#)  
731. [697](#)  
732. [698](#)  
733. [699](#)  
734. [700](#)  
735. [701](#)  
736. [702](#)  
737. [703](#)  
738. [704](#)  
739. [705](#)  
740. [706](#)  
741. [707](#)  
742. [708](#)  
743. [709](#)  
744. [710](#)  
745. [711](#)  
746. [712](#)  
747. [713](#)  
748. [714](#)  
749. [715](#)  
750. [716](#)  
751. [717](#)  
752. [718](#)  
753. [719](#)  
754. [720](#)  
755. [721](#)  
756. [722](#)  
757. [723](#)  
758. [724](#)  
759. [725](#)  
760. [726](#)  
761. [727](#)  
762. [728](#)  
763. [729](#)  
764. [730](#)  
765. [731](#)  
766. [732](#)  
767. [733](#)  
768. [734](#)

769. [735](#)  
770. [736](#)  
771. [737](#)  
772. [738](#)  
773. [739](#)  
774. [740](#)  
775. [741](#)  
776. [742](#)  
777. [743](#)  
778. [744](#)  
779. [745](#)  
780. [746](#)  
781. [747](#)  
782. [748](#)  
783. [749](#)  
784. [750](#)  
785. [751](#)  
786. [752](#)  
787. [753](#)  
788. [754](#)  
789. [755](#)  
790. [756](#)  
791. [757](#)  
792. [758](#)  
793. [759](#)  
794. [760](#)  
795. [761](#)  
796. [762](#)  
797. [763](#)  
798. [764](#)  
799. [765](#)  
300. [766](#)  
301. [767](#)  
302. [768](#)  
303. [769](#)  
304. [770](#)  
305. [771](#)  
306. [772](#)  
307. [773](#)  
308. [774](#)  
309. [775](#)  
310. [776](#)  
311. [777](#)  
312. [778](#)  
313. [779](#)

314. [780](#)  
315. [781](#)  
316. [782](#)  
317. [783](#)  
318. [784](#)  
319. [785](#)  
320. [786](#)  
321. [787](#)  
322. [788](#)  
323. [789](#)  
324. [790](#)  
325. [791](#)  
326. [792](#)  
327. [793](#)  
328. [794](#)  
329. [795](#)  
330. [796](#)  
331. [797](#)  
332. [798](#)  
333. [799](#)  
334. [800](#)  
335. [801](#)  
336. [802](#)  
337. [803](#)  
338. [804](#)  
339. [805](#)  
340. [806](#)  
341. [807](#)  
342. [808](#)  
343. [809](#)  
344. [810](#)  
345. [811](#)  
346. [812](#)  
347. [813](#)  
348. [814](#)  
349. [815](#)  
350. [816](#)  
351. [817](#)  
352. [818](#)  
353. [819](#)  
354. [820](#)  
355. [821](#)  
356. [822](#)  
357. [823](#)  
358. [824](#)

359. [825](#)  
360. [826](#)  
361. [827](#)  
362. [828](#)  
363. [829](#)  
364. [830](#)  
365. [831](#)  
366. [832](#)  
367. [833](#)  
368. [834](#)  
369. [835](#)  
370. [836](#)  
371. [837](#)  
372. [838](#)  
373. [839](#)  
374. [840](#)  
375. [841](#)  
376. [842](#)  
377. [843](#)  
378. [844](#)  
379. [845](#)  
380. [846](#)  
381. [847](#)  
382. [848](#)  
383. [849](#)  
384. [850](#)  
385. [851](#)  
386. [852](#)  
387. [853](#)  
388. [854](#)  
389. [855](#)  
390. [856](#)  
391. [857](#)  
392. [858](#)  
393. [859](#)  
394. [860](#)  
395. [861](#)  
396. [862](#)  
397. [863](#)  
398. [864](#)  
399. [865](#)  
400. [866](#)  
401. [867](#)  
402. [868](#)  
403. [869](#)

004. [870](#)  
005. [871](#)  
006. [872](#)  
007. [873](#)  
008. [874](#)  
009. [875](#)  
010. [876](#)  
011. [877](#)  
012. [878](#)  
013. [879](#)  
014. [880](#)  
015. [881](#)  
016. [882](#)  
017. [883](#)  
018. [884](#)  
019. [885](#)  
020. [886](#)  
021. [887](#)  
022. [888](#)  
023. [889](#)  
024. [890](#)  
025. [891](#)  
026. [892](#)  
027. [893](#)  
028. [894](#)  
029. [895](#)  
030. [896](#)  
031. [897](#)  
032. [898](#)  
033. [899](#)  
034. [900](#)  
035. [901](#)  
036. [902](#)  
037. [903](#)  
038. [904](#)  
039. [905](#)  
040. [906](#)  
041. [907](#)  
042. [908](#)  
043. [909](#)  
044. [910](#)  
045. [911](#)  
046. [912](#)  
047. [913](#)  
048. [914](#)



49. [915](#)  
50. [916](#)  
51. [917](#)  
52. [918](#)  
53. [919](#)  
54. [920](#)  
55. [921](#)  
56. [922](#)  
57. [923](#)  
58. [924](#)  
59. [925](#)  
60. [926](#)  
61. [927](#)  
62. [928](#)  
63. [929](#)  
64. [930](#)  
65. [931](#)  
66. [932](#)  
67. [933](#)  
68. [934](#)  
69. [935](#)  
70. [936](#)  
71. [937](#)  
72. [938](#)  
73. [939](#)  
74. [940](#)  
75. [941](#)  
76. [942](#)  
77. [943](#)  
78. [944](#)  
79. [945](#)  
80. [946](#)  
81. [947](#)  
82. [948](#)  
83. [949](#)  
84. [950](#)  
85. [951](#)  
86. [952](#)  
87. [953](#)  
88. [954](#)  
89. [955](#)  
90. [956](#)  
91. [957](#)  
92. [958](#)  
93. [959](#)

994. [960](#)  
995. [961](#)  
996. [962](#)  
997. [963](#)  
998. [964](#)  
999. [965](#)  
1000. [966](#)  
1001. [967](#)  
1002. [968](#)  
1003. [969](#)  
1004. [970](#)  
1005. [971](#)  
1006. [972](#)  
1007. [973](#)  
1008. [974](#)  
1009. [975](#)  
1010. [976](#)  
1011. [977](#)  
1012. [978](#)  
1013. [979](#)  
1014. [980](#)  
1015. [981](#)  
1016. [982](#)  
1017. [983](#)  
1018. [984](#)  
1019. [985](#)  
1020. [986](#)  
1021. [987](#)  
1022. [988](#)  
1023. [989](#)  
1024. [990](#)  
1025. [991](#)  
1026. [992](#)  
1027. [993](#)  
1028. [994](#)  
1029. [995](#)  
1030. [996](#)  
1031. [997](#)  
1032. [998](#)  
1033. [999](#)  
1034. [1000](#)  
1035. [1001](#)  
1036. [1002](#)  
1037. [1003](#)  
1038. [1004](#)

139. [1005](#)  
140. [1006](#)  
141. [1007](#)  
142. [1008](#)  
143. [1009](#)  
144. [1010](#)  
145. [1011](#)  
146. [1012](#)  
147. [1013](#)  
148. [1014](#)  
149. [1015](#)  
150. [1016](#)  
151. [1017](#)  
152. [1018](#)  
153. [1019](#)  
154. [1020](#)  
155. [1021](#)  
156. [1022](#)  
157. [1023](#)  
158. [1024](#)  
159. [1025](#)  
160. [1026](#)  
161. [1027](#)  
162. [1028](#)  
163. [1029](#)  
164. [1030](#)  
165. [1031](#)  
166. [1032](#)  
167. [1033](#)  
168. [1034](#)  
169. [1035](#)  
170. [1036](#)  
171. [1037](#)  
172. [1038](#)  
173. [1039](#)  
174. [1040](#)  
175. [1041](#)  
176. [1042](#)  
177. [1043](#)  
178. [1044](#)  
179. [1045](#)  
180. [1046](#)  
181. [1047](#)  
182. [1048](#)  
183. [1049](#)

984. [1050](#)  
985. [1051](#)  
986. [1052](#)  
987. [1053](#)  
988. [1054](#)  
989. [1055](#)  
990. [1056](#)  
991. [1057](#)  
992. [1058](#)  
993. [1059](#)  
994. [1060](#)  
995. [1061](#)  
996. [1062](#)  
997. [1063](#)  
998. [1064](#)  
999. [1065](#)  
1000. [1066](#)  
1001. [1067](#)  
1002. [1068](#)  
1003. [1069](#)  
1004. [1070](#)  
1005. [1071](#)  
1006. [1072](#)  
1007. [1073](#)  
1008. [1074](#)  
1009. [1075](#)  
1010. [1076](#)  
1011. [1077](#)  
1012. [1078](#)  
1013. [1079](#)  
1014. [1080](#)  
1015. [1081](#)  
1016. [1082](#)  
1017. [1083](#)  
1018. [1084](#)  
1019. [1085](#)  
1020. [1086](#)  
1021. [1087](#)  
1022. [1088](#)  
1023. [1089](#)  
1024. [1090](#)  
1025. [1091](#)  
1026. [1092](#)  
1027. [1093](#)  
1028. [1094](#)

129. [1095](#)  
130. [1096](#)  
131. [1097](#)  
132. [1098](#)  
133. [1099](#)  
134. [1100](#)  
135. [1101](#)  
136. [1102](#)  
137. [1103](#)  
138. [1104](#)  
139. [1105](#)  
140. [1106](#)  
141. [1107](#)  
142. [1108](#)  
143. [1109](#)  
144. [1110](#)  
145. [1111](#)  
146. [1112](#)  
147. [1113](#)  
148. [1114](#)  
149. [1115](#)  
150. [1116](#)  
151. [1117](#)  
152. [1118](#)  
153. [1119](#)  
154. [1120](#)  
155. [1121](#)  
156. [1122](#)  
157. [1123](#)  
158. [1124](#)  
159. [1125](#)  
160. [1126](#)  
161. [1127](#)  
162. [1128](#)  
163. [1129](#)  
164. [1130](#)  
165. [1131](#)  
166. [1132](#)  
167. [1133](#)  
168. [1134](#)  
169. [1135](#)  
170. [1136](#)  
171. [1137](#)  
172. [1138](#)  
173. [1139](#)

174. [1140](#)  
175. [1141](#)  
176. [1142](#)  
177. [1143](#)  
178. [1144](#)  
179. [1145](#)  
180. [1146](#)  
181. [1147](#)  
182. [1148](#)  
183. [1149](#)  
184. [1150](#)  
185. [1151](#)  
186. [1152](#)  
187. [1153](#)  
188. [1154](#)  
189. [1155](#)  
190. [1156](#)  
191. [1157](#)  
192. [1158](#)  
193. [1159](#)  
194. [1160](#)  
195. [1161](#)  
196. [1162](#)  
197. [1163](#)  
198. [1164](#)  
199. [1165](#)  
200. [1166](#)  
201. [1167](#)  
202. [1168](#)  
203. [1169](#)  
204. [1170](#)  
205. [1171](#)  
206. [1172](#)  
207. [1173](#)  
208. [1174](#)  
209. [1175](#)  
210. [1176](#)  
211. [1177](#)  
212. [1178](#)  
213. [1179](#)  
214. [1180](#)  
215. [1181](#)  
216. [1182](#)  
217. [1183](#)  
218. [1184](#)

219. [1185](#)  
220. [1186](#)  
221. [1187](#)  
222. [1188](#)  
223. [1189](#)  
224. [1190](#)  
225. [1191](#)  
226. [1192](#)  
227. [1193](#)  
228. [1194](#)  
229. [1195](#)  
230. [1196](#)  
231. [1197](#)  
232. [1198](#)  
233. [1199](#)  
234. [1200](#)  
235. [1201](#)  
236. [1202](#)  
237. [1203](#)  
238. [1204](#)  
239. [1205](#)  
240. [1206](#)  
241. [1207](#)  
242. [1208](#)  
243. [1209](#)  
244. [1210](#)  
245. [1211](#)  
246. [1212](#)  
247. [1213](#)  
248. [1214](#)  
249. [1215](#)  
250. [1216](#)  
251. [1217](#)  
252. [1218](#)  
253. [1219](#)  
254. [1220](#)  
255. [1221](#)  
256. [1222](#)  
257. [1223](#)  
258. [1224](#)  
259. [1225](#)  
260. [1226](#)  
261. [1227](#)  
262. [1228](#)  
263. [1229](#)

264. [1230](#)  
265. [1231](#)  
266. [1232](#)  
267. [1233](#)  
268. [1234](#)  
269. [1235](#)  
270. [1236](#)  
271. [1237](#)  
272. [1238](#)  
273. [1239](#)  
274. [1240](#)  
275. [1241](#)  
276. [1242](#)  
277. [1243](#)  
278. [1244](#)  
279. [1245](#)  
280. [1246](#)  
281. [1247](#)  
282. [1248](#)  
283. [1249](#)  
284. [1250](#)  
285. [1251](#)  
286. [1252](#)  
287. [1253](#)  
288. [1254](#)  
289. [1255](#)  
290. [1256](#)  
291. [1257](#)  
292. [1258](#)  
293. [1259](#)  
294. [1260](#)  
295. [1261](#)  
296. [1262](#)  
297. [1263](#)  
298. [1264](#)  
299. [1265](#)  
300. [1266](#)  
301. [1267](#)  
302. [1268](#)  
303. [1269](#)  
304. [1270](#)  
305. [1271](#)  
306. [1272](#)  
307. [1273](#)  
308. [1274](#)



309. [1275](#)  
310. [1276](#)  
311. [1277](#)  
312. [1278](#)  
313. [1279](#)  
314. [1280](#)  
315. [1281](#)  
316. [1282](#)  
317. [1283](#)  
318. [1284](#)  
319. [1285](#)  
320. [1286](#)  
321. [1287](#)  
322. [1288](#)  
323. [1289](#)  
324. [1290](#)  
325. [1291](#)  
326. [1292](#)  
327. [1293](#)  
328. [1294](#)  
329. [1295](#)  
330. [1296](#)  
331. [1297](#)  
332. [1298](#)  
333. [1299](#)  
334. [1300](#)  
335. [1301](#)  
336. [1302](#)  
337. [1303](#)  
338. [1304](#)  
339. [1305](#)  
340. [1306](#)  
341. [1307](#)  
342. [1308](#)  
343. [1309](#)  
344. [1310](#)  
345. [1311](#)  
346. [1312](#)  
347. [1313](#)  
348. [1314](#)  
349. [1315](#)  
350. [1316](#)  
351. [1317](#)  
352. [1318](#)  
353. [1319](#)

354. [1320](#)  
355. [1321](#)  
356. [1322](#)  
357. [1323](#)  
358. [1324](#)  
359. [1325](#)  
360. [1326](#)  
361. [1327](#)  
362. [1328](#)  
363. [1329](#)  
364. [1330](#)  
365. [1331](#)  
366. [1332](#)  
367. [1333](#)  
368. [1334](#)  
369. [1335](#)  
370. [1336](#)  
371. [1337](#)  
372. [1338](#)  
373. [1339](#)  
374. [1340](#)  
375. [1341](#)  
376. [1342](#)  
377. [1343](#)  
378. [1344](#)  
379. [1345](#)  
380. [1346](#)  
381. [1347](#)  
382. [1348](#)  
383. [1349](#)  
384. [1350](#)  
385. [1351](#)  
386. [1352](#)  
387. [22-1](#)  
388. [22-2](#)  
389. [22-3](#)  
390. [22-4](#)  
391. [22-5](#)  
392. [22-6](#)  
393. [22-7](#)  
394. [22-8](#)  
395. [22-9](#)  
396. [22-10](#)  
397. [22-11](#)  
398. [22-12](#)

399. [22-13](#)  
400. [22-14](#)  
401. [22-15](#)  
402. [22-16](#)  
403. [22-17](#)  
404. [22-18](#)  
405. [22-19](#)  
406. [22-20](#)  
407. [22-21](#)  
408. [22-22](#)  
409. [22-23](#)  
410. [22-24](#)  
411. [22-25](#)  
412. [22-26](#)  
413. [22-27](#)  
414. [22-28](#)  
415. [22-29](#)  
416. [22-30](#)  
417. [22-31](#)  
418. [22-32](#)  
419. [22-33](#)  
420. [22-34](#)  
421. [22-35](#)  
422. [22-36](#)  
423. [22-37](#)  
424. [22-38](#)  
425. [22-39](#)  
426. [22-40](#)  
427. [22-41](#)  
428. [22-42](#)  
429. [22-43](#)  
430. [22-44](#)  
431. [22-45](#)  
432. [22-46](#)  
433. [22-47](#)  
434. [22-48](#)  
435. [22-49](#)  
436. [22-50](#)  
437. [22-51](#)  
438. [22-52](#)  
439. [22-53](#)  
440. [22-54](#)  
441. [22-55](#)  
442. [22-56](#)  
443. [22-57](#)

144. [22-58](#)  
145. [22-59](#)  
146. [22-60](#)  
147. [22-61](#)  
148. [22-62](#)  
149. [22-63](#)  
150. [22-64](#)  
151. [22-65](#)  
152. [22-66](#)  
153. [22-67](#)  
154. [22-68](#)  
155. [22-69](#)  
156. [22-70](#)  
157. [22-71](#)  
158. [22-72](#)  
159. [22-73](#)  
160. [22-74](#)  
161. [22-75](#)  
162. [22-76](#)  
163. [22-77](#)  
164. [22-78](#)  
165. [22-79](#)  
166. [22-80](#)  
167. [22-81](#)  
168. [22-82](#)  
169. [23-1](#)  
170. [23-2](#)  
171. [23-3](#)  
172. [23-4](#)  
173. [23-5](#)  
174. [23-6](#)  
175. [23-7](#)  
176. [23-8](#)  
177. [23-9](#)  
178. [23-10](#)  
179. [23-11](#)  
180. [23-12](#)  
181. [23-13](#)  
182. [23-14](#)  
183. [23-15](#)  
184. [23-16](#)  
185. [23-17](#)  
186. [23-18](#)  
187. [23-19](#)  
188. [23-20](#)

489. [23-21](#)  
490. [23-22](#)  
491. [23-23](#)  
492. [23-24](#)  
493. [23-25](#)  
494. [23-26](#)  
495. [23-27](#)  
496. [23-28](#)  
497. [23-29](#)  
498. [23-30](#)  
499. [23-31](#)  
500. [23-32](#)  
501. [23-33](#)  
502. [23-34](#)  
503. [23-35](#)  
504. [23-36](#)  
505. [23-37](#)  
506. [23-38](#)  
507. [23-39](#)  
508. [23-40](#)  
509. [23-41](#)  
510. [23-42](#)  
511. [23-43](#)  
512. [23-44](#)  
513. [23-45](#)  
514. [23-46](#)  
515. [23-47](#)  
516. [23-48](#)  
517. [23-49](#)  
518. [23-50](#)  
519. [23-51](#)  
520. [23-52](#)  
521. [23-53](#)  
522. [23-54](#)  
523. [23-55](#)  
524. [23-56](#)  
525. [23-57](#)  
526. [23-58](#)  
527. [23-59](#)  
528. [23-60](#)  
529. [23-61](#)  
530. [23-62](#)  
531. [23-63](#)  
532. [23-64](#)  
533. [23-65](#)

534. [23-66](#)  
535. [23-67](#)  
536. [23-68](#)  
537. [23-69](#)  
538. [23-70](#)  
539. [23-71](#)  
540. [23-72](#)  
541. [23-73](#)  
542. [23-74](#)  
543. [23-75](#)  
544. [23-76](#)  
545. [23-77](#)  
546. [23-78](#)  
547. [23-79](#)  
548. [23-80](#)  
549. [23-81](#)  
550. [23-82](#)  
551. [23-83](#)  
552. [23-84](#)  
553. [23-85](#)  
554. [23-86](#)  
555. [23-87](#)  
556. [23-88](#)  
557. [24-1](#)  
558. [24-2](#)  
559. [24-3](#)  
560. [24-4](#)  
561. [24-5](#)  
562. [24-6](#)  
563. [24-7](#)  
564. [24-8](#)  
565. [24-9](#)  
566. [24-10](#)  
567. [24-11](#)  
568. [24-12](#)  
569. [24-13](#)  
570. [24-14](#)  
571. [24-15](#)  
572. [24-16](#)  
573. [24-17](#)  
574. [24-18](#)  
575. [24-19](#)  
576. [24-20](#)  
577. [24-21](#)  
578. [24-22](#)

579. [24-23](#)  
580. [24-24](#)  
581. [24-25](#)  
582. [24-26](#)  
583. [24-27](#)  
584. [24-28](#)  
585. [24-29](#)  
586. [24-30](#)  
587. [24-31](#)  
588. [24-32](#)  
589. [24-33](#)  
590. [24-34](#)  
591. [24-35](#)  
592. [24-36](#)  
593. [24-37](#)  
594. [24-38](#)  
595. [24-39](#)  
596. [24-40](#)  
597. [24-41](#)  
598. [24-42](#)  
599. [24-43](#)  
500. [24-44](#)  
501. [24-45](#)  
502. [24-46](#)  
503. [24-47](#)  
504. [24-48](#)  
505. [24-49](#)  
506. [24-50](#)  
507. [24-51](#)  
508. [24-52](#)  
509. [24-53](#)  
510. [24-54](#)  
511. [24-55](#)  
512. [24-56](#)  
513. [24-57](#)  
514. [24-58](#)  
515. [24-59](#)  
516. [24-60](#)  
517. [24-61](#)  
518. [24-62](#)  
519. [24-63](#)  
520. [24-64](#)  
521. [24-65](#)  
522. [24-66](#)  
523. [24-67](#)

524. [24-68](#)  
525. [25-1](#)  
526. [25-2](#)  
527. [25-3](#)  
528. [25-4](#)  
529. [25-5](#)  
530. [25-6](#)  
531. [25-7](#)  
532. [25-8](#)  
533. [25-9](#)  
534. [25-10](#)  
535. [25-11](#)  
536. [25-12](#)  
537. [25-13](#)  
538. [25-14](#)  
539. [25-15](#)  
540. [25-16](#)  
541. [25-17](#)  
542. [25-18](#)  
543. [25-19](#)  
544. [25-20](#)  
545. [25-21](#)  
546. [25-22](#)  
547. [25-23](#)  
548. [25-24](#)  
549. [25-25](#)  
550. [25-26](#)  
551. [25-27](#)  
552. [25-28](#)  
553. [25-29](#)  
554. [25-30](#)  
555. [25-31](#)  
556. [25-32](#)  
557. [25-33](#)  
558. [25-34](#)  
559. [25-35](#)  
560. [25-36](#)  
561. [25-37](#)  
562. [25-38](#)  
563. [25-39](#)  
564. [25-40](#)  
565. [25-41](#)  
566. [25-42](#)  
567. [25-43](#)  
568. [25-44](#)



569. [25-45](#)  
570. [25-46](#)  
571. [25-47](#)  
572. [25-48](#)  
573. [25-49](#)  
574. [25-50](#)  
575. [25-51](#)  
576. [25-52](#)  
577. [25-53](#)  
578. [25-54](#)  
579. [25-55](#)  
580. [25-56](#)  
581. [25-57](#)  
582. [25-58](#)  
583. [25-59](#)  
584. [25-60](#)  
585. [25-61](#)  
586. [25-62](#)  
587. [25-63](#)  
588. [25-64](#)  
589. [25-65](#)  
590. [25-66](#)  
591. [25-67](#)  
592. [25-68](#)  
593. [25-69](#)  
594. [25-70](#)  
595. [25-71](#)  
596. [25-72](#)  
597. [25-73](#)  
598. [25-74](#)  
599. [1353](#)  
700. [1354](#)  
701. [1355](#)  
702. [1356](#)  
703. [1357](#)  
704. [1358](#)  
705. [1359](#)  
706. [1360](#)  
707. [1361](#)  
708. [1362](#)  
709. [1363](#)  
710. [1364](#)  
711. [1365](#)  
712. [1366](#)  
713. [1367](#)

714. [1368](#)  
715. [1369](#)  
716. [1370](#)  
717. [1371](#)  
718. [1372](#)  
719. [1373](#)  
720. [1374](#)  
721. [A-1](#)  
722. [B-1](#)  
723. [B-2](#)  
724. [C-1](#)  
725. [D-1](#)  
726. [D-2](#)  
727. [D-3](#)  
728. [D-4](#)  
729. [E-1](#)  
730. [E-2](#)  
731. [E-3](#)  
732. [E-4](#)  
733. [E-5](#)  
734. [E-6](#)  
735. [E-7](#)  
736. [F-1](#)  
737. [F-2](#)  
738. [F-3](#)  
739. [G-1](#)  
740. [G-2](#)  
741. [G-3](#)  
742. [G-4](#)  
743. [G-5](#)  
744. [G-6](#)  
745. [G-7](#)  
746. [G-8](#)  
747. [G-9](#)  
748. [G-10](#)  
749. [G-11](#)  
750. [G-12](#)  
751. [H-1](#)  
752. [H-2](#)  
753. [H-3](#)  
754. [H-4](#)  
755. [H-5](#)  
756. [H-6](#)  
757. [H-7](#)  
758. [H-8](#)

759. [H-9](#)  
760. [H-10](#)  
761. [H-11](#)  
762. [H-12](#)  
763. [H-13](#)  
764. [H-14](#)  
765. [I-1](#)  
766. [I-2](#)  
767. [I-3](#)  
768. [I-4](#)  
769. [I-5](#)  
770. [I-6](#)  
771. [I-7](#)  
772. [J-1](#)  
773. [J-2](#)  
774. [J-3](#)  
775. [J-4](#)  
776. [J-5](#)  
777. [K-1](#)  
778. [K-2](#)  
779. [L-1](#)  
780. [L-2](#)  
781. [L-3](#)  
782. [L-4](#)  
783. [L-5](#)  
784. [L-6](#)  
785. [L-7](#)  
786. [L-8](#)  
787. [L-9](#)  
788. [L-10](#)  
789. [L-11](#)  
790. [L-12](#)  
791. [L-13](#)  
792. [L-14](#)  
793. [L-15](#)  
794. [L-16](#)  
795. [L-17](#)  
796. [L-18](#)  
797. [L-19](#)  
798. [L-20](#)  
799. [L-21](#)  
800. [L-22](#)  
801. [L-23](#)  
802. [L-24](#)  
803. [L-25](#)

304. [L-26](#)  
305. [L-27](#)  
306. [L-28](#)  
307. [L-29](#)  
308. [L-30](#)  
309. [L-31](#)  
310. [L-32](#)  
311. [L-33](#)  
312. [L-34](#)  
313. [L-35](#)  
314. [L-36](#)  
315. [L-37](#)  
316. [L-38](#)  
317. [L-39](#)  
318. [L-40](#)  
319. [L-41](#)  
320. [L-42](#)  
321. [M-1](#)  
322. [M-2](#)  
323. [M-3](#)  
324. [M-4](#)  
325. [M-5](#)  
326. [M-6](#)  
327. [M-7](#)  
328. [M-8](#)  
329. [M-9](#)  
330. [M-10](#)  
331. [M-11](#)  
332. [M-12](#)  
333. [M-13](#)  
334. [M-14](#)  
335. [M-15](#)  
336. [M-16](#)  
337. [M-17](#)  
338. [M-18](#)  
339. [M-19](#)  
340. [M-20](#)  
341. [M-21](#)  
342. [M-22](#)  
343. [M-23](#)  
344. [M-24](#)  
345. [M-25](#)  
346. [M-26](#)  
347. [M-27](#)  
348. [M-28](#)

349. [M-29](#)  
350. [M-30](#)  
351. [M-31](#)  
352. [M-32](#)  
353. [M-33](#)  
354. [M-34](#)  
355. [M-35](#)  
356. [M-36](#)  
357. [M-37](#)  
358. [M-38](#)  
359. [M-39](#)  
360. [M-40](#)  
361. [M-41](#)  
362. [M-42](#)  
363. [M-43](#)  
364. [CS1-1](#)  
365. [CS1-2](#)  
366. [CS1-3](#)  
367. [CS1-4](#)  
368. [CS2-1](#)  
369. [CS2-2](#)  
370. [CS2-3](#)  
371. [CS2-4](#)  
372. [CS2-5](#)  
373. [CS2-6](#)  
374. [CS2-7](#)  
375. [CS2-8](#)  
376. [CS2-9](#)  
377. [CS2-10](#)  
378. [CS3-1](#)  
379. [CS3-2](#)  
380. [CS3-3](#)  
381. [CS3-4](#)  
382. [CS3-5](#)  
383. [CS4-1](#)  
384. [CS4-2](#)  
385. [CS4-3](#)  
386. [CS4-4](#)  
387. [CS5-1](#)  
388. [CS5-2](#)  
389. [CS5-3](#)  
390. [CS5-4](#)  
391. [CS5-5](#)  
392. [CS5-6](#)  
393. [CS5-7](#)

- 394. [CS5-8](#)
- 395. [CS5-9](#)
- 396. [CS5-10](#)
- 397. [CS5-11](#)
- 398. [CS5-12](#)
- 399. [CS5-13](#)
- 400. [CS5-14](#)
- 401. [CS6-1](#)
- 402. [CS6-2](#)
- 403. [CS6-3](#)
- 404. [CS6-4](#)
- 405. [CS6-5](#)
- 406. [CS6-6](#)
- 407. [CS6-7](#)
- 408. [CS7-1](#)
- 409. [CS7-2](#)
- 410. [CS7-3](#)
- 411. [CS7-4](#)
- 412. [CS7-5](#)
- 413. [CS7-6](#)
- 414. [CS7-7](#)
- 415. [CS7-8](#)
- 416. [CS7-9](#)
- 417. [CS7-10](#)
- 418. [CS7-11](#)
- 419. [CS7-12](#)

# Long description

From bottom to top, the chapter dependency chart is as follows. Chapter 21 Binary Trees, AVL Trees, and Priority Queues depends on Chapter 20 Stacks and Queues. Chapter 20 depends on Chapter 19 Linked Lists. Chapter 19 and Chapter 18 Collections and the Stream API depend on Chapter 17 Generics. Chapter 13 JavaFX: Advanced Controls and Chapter 14 JavaFX: Graphics, Effects, and Media depend on Chapter 12 JavaFX: GUI Programming and Basic Controls. Chapters 12, 17, and 11 Exceptions and Advanced File I/O depend on Chapter 10 Inheritance. Also, Chapter 16 Sorting, Searching, and Algorithm Analysis depends on Chapter 10. Chapter 10 depends on Chapter 8 A Second Look at Classes and Objects. Chapters 9, 8, 22 (Online) Databases, and 15 Recursion depend on Chapters 1 to 7 (Cover in order) Java Fundamentals. A few side notes reveal the following information. Some examples in Chapter 22 use JavaFX, which is introduced in Chapter 12. One example in Chapter 15 uses the JavaFX Circle class, which is introduced in Chapter 14. Some examples in Chapter 19 use recursion, which is introduced in Chapter 15. Chapter 21's Priority Queue section includes a discussion of algorithm analysis, which is introduced in Chapter 16. Some examples in Chapter 21 use recursion, which is introduced in Chapter 16. While the chapter dependencies are indicated with solid arrows the relationships between chapters based on examples used are indicated using dashed arrows. The arrows to and from the respective chapters.

# Long description

The process reads as follows. 1, the programmer uses a text editor to create a java source code file. 2, the programmer runs the compiler, which translates the source code file into a byte code file. 3, the java virtual machine reads and executes each byte code instruction.



# Long description

Within a class's braces, as well as inside the braces of a main method, each line is indented. A long statement can be spread out over several lines. For example, the `println` statement illustrated here, is broken into three lines that each start with extra spaces to indicate that they are continuous.

# Long description

The figure diagrams this statement: `Scanner keyboard = new Scanner (System.in);`

`Scanner keyboard` declares a variable named `keyboard`. The variable can reference an object of the `Scanner` class.

The `=` sign is an operator that assigns the address of the `Scanner` object to the `keyboard` variable.

`new Scanner (System.in);` Creates a `Scanner` object in memory. The object will read input from `System.in`.

# Long description

A message dialog box and an input dialog contain some overlapping and some different features. The message dialog box displays the message, This is a message dialog, and includes an OK button. The input dialog displays the message, This is an input dialog, and includes a text field and OK button.

# Long description

The first dialog box is labeled Input and asks users: What is your first name? The text Joe is written in the text box, and the OK button is selected.

The second dialog box is labeled Input and asks users, What is your last name? The text Clondike is written in the text box, and the OK button is selected.

The third dialog box appears as a message, Hello Joe Clondike. The OK button is selected.

# Long description

The first dialog box is labeled Input and asks users: What is your name? The text Joe Mahoney is written in the text box, and the OK button is selected. The second dialog box is labeled Input and asks users: How many hours did you work this week? The text 40 is written in the text box, and the OK button is selected. The third dialog box is labeled Input and asks users: What is your hourly pay rate? The text 20 is written in the text box, and the OK button is selected. The fourth dialog box is labeled Message and reads: Hello Joe Mahoney. Your gross pay is \$800.00. The OK button is selected.

# Long description

The strings of code read as follows. Line 1, public class Square Area. Line 2, left brace. Line 3, public static void main left parenthesis string left bracket right bracket a r g s right parenthesis. Line 4, left brace. Line 5, double length, width, area semicolon. Step 1, length = 10 semicolon. Step 2, width = 5 semicolon. Step 3: area = length times width semicolon. Step 4, System dot out dot print left parenthesis, the words the area is, in quotation marks, + area right parenthesis semicolon. Two lines each have one right brace.

# Long description

A prompt asks a, yes forward slash no, question or a, true forward slash false, condition. The prompt asks if it is cold outside. If the answer is yes and the condition is true, the action triggers the message, wear a coat. If the answer is no, the program follows another path and skips the action.

# Long description

The action starts with the question, is it cold outside? is yes, the actions to wear a coat, wear a hat, and wear gloves are performed sequentially. If the answer to the question, Is it cold outside? If no, the actions are skipped.



# Long description

An if statement is missing its braces. The code reads as follows. If left parenthesis sales are greater than 50,000 right parenthesis. Bonus = 500.0 semicolon. Only this statement is conditionally executed. Two statements are always executed and read as follows, commission Rate = 0. 12 semicolon. Days Off + = 1 semicolon.

# Long description

The logic of the if dash else statement. The value is number 2 == 0. When a user enters 0, the statements controlled by the if clause display an error message. If the user does not enter 0, the else clause will display the result of number 1 divided by number 2.

# Long description

In this situation, to qualify for a low interest rate on a loan, a bank customer must fulfill two conditions. First, the customer's salary must be at least \$50,000. When this first condition is false, the if statement triggers the display of a message that reads, you must earn at least \$50,000 per year to qualify, in quotation marks. When the first condition is met, the second condition is represented in another if statement nestled inside the first if statement. The condition reads years On Job is greater than or equal to 2. If the condition is true, the if statement triggers the display of a message that reads, you qualify for the loan, in quotation marks. If the condition is false, the if statement triggers the display of a message that reads, you must have been on your current job for at least two years to qualify, in quotation marks.

# Long description

The array reads as follows, the word, double, left bracket right bracket, left bracket right bracket, scores = new double left bracket 3 right bracket, left bracket 4 right bracket semicolon. Two sets of brackets indicate a two-dimensional array. The number, left bracket 3 right bracket, refers to the number of rows, and the number, left bracket 4 right bracket, refers to the number of columns.

# Long description

The table is labeled, Column 0, Column 1, Column 2, and Column 3, and three rows labeled, Row 0, Row 1, and Row 2. Each field contains subscripts for each element of the scores array. The contents read as follows. Row 0, Column 0. scores left bracket 0 right bracket, left bracket 0 right bracket. Row 0, Column 1. scores left bracket 0 right bracket, left bracket 1 right bracket. Row 0, Column 2. scores left bracket 0 right bracket, left bracket 2 right bracket. Row 0, Column 3. scores left bracket 0 right bracket, left bracket 3 right bracket. Row 1, Column 0. scores left bracket 1 right bracket, left bracket 0 right bracket. Row 1, Column 1. scores left bracket 1 right bracket, left bracket 1 right bracket. Row 1, Column 2. scores left bracket 1 right bracket, left bracket 2 right bracket. Row 1, Column 3. scores left bracket 1 right bracket, left bracket 3 right bracket. Row 2, Column 0. scores left bracket 2 right bracket, left bracket 0 right bracket. Row 2, Column 1. scores left bracket 2 right bracket, left bracket 1 right bracket. Row 2, Column 2. scores left bracket 2 right bracket, left bracket 2 right bracket. Row 2, Column 3. scores left bracket 2 right bracket, left bracket 3 right bracket.

# Long description

The table is labeled, Column 0, Column 1, Column 2, and Column 3, and three rows labeled, Row 0, Row 1, and Row 2. Each field contains subscripts for each element of the scores array. The contents read as follows. Row 0, Column 0. The array reads, sales left bracket 0 right bracket, left bracket 0 right bracket. A note reads, holds data for division 1, quarter 1. Row 0, Column 1. The array reads, sales left bracket 0 right bracket, left bracket 1 right bracket. A note reads, holds data for division 1, quarter 2. Row 0, Column 2. The array reads, sales left bracket 0 right bracket, left bracket 2 right bracket. A note reads, holds data for division 1, quarter 3. Row 0, Column 3. The array reads, sales left bracket 0 right bracket, left bracket 3 right bracket. A note reads, holds data for division 1, quarter 4. Row 1, Column 0. The array reads, sales left bracket 1 right bracket, left bracket 0 right bracket. A note reads, holds data for division 2, quarter 1. Row 1, Column 1. The array reads, sales left bracket 1 right bracket, left bracket 1 right bracket. A note reads, holds data for division 2, quarter 2. Row 1, Column 2. The array reads, sales left bracket 1 right bracket, left bracket 2 right bracket. A note reads, holds data for division 2, quarter 3. Row 1, Column 3. The array reads, sales left bracket 1 right bracket, left bracket 3 right bracket. A note reads, holds data for division 2, quarter 4. Row 2, Column 0. The array reads, sales left bracket 2 right bracket, left bracket 0 right bracket. A note reads, holds data for division 3, quarter 1. Row 2, Column 1. The array reads, sales left bracket 2 right bracket, left bracket 1 right bracket. A note reads, holds data for division 3, quarter 2. Row 2, Column 2. The array reads, sales left bracket 2 right bracket, left bracket 2 right bracket. A note reads, holds data for division 3, quarter 3. Row 2, Column 3. The array reads, sales left bracket 2 right bracket, left bracket 3 right bracket. A note reads, holds data for division 3, quarter 4.

# Long description

The table is labeled, Column 0, Column 1, Column 2, and three rows labeled, Row 0, Row 1, and Row 2. Each box contains a numbers variable and a number. The contents read as follows. Row 0, Column 0. numbers left bracket 0 right bracket, left bracket 0 right bracket. The number, 1. Row 0, Column 1. numbers left bracket 0 right bracket, left bracket 1 right bracket. The number, 2. Row 0, Column 2. numbers left bracket 0 right bracket, left bracket 2 right bracket. The number, 3. Row 1, Column 0. numbers left bracket 1 right bracket, left bracket 0 right bracket. The number, 4. Row 1, Column 1. numbers left bracket 1 right bracket, left bracket 1 right bracket. The number, 5. Row 1, Column 2. numbers left bracket 1 right bracket, left bracket 2 right bracket. The number, 6. Row 2, Column 0. numbers left bracket 2 right bracket, left bracket 0 right bracket. The number, 7. Row 2, Column 1. numbers left bracket 2 right bracket, left bracket 1 right bracket. The number, 8. Row 2, Column 2. numbers left bracket 2 right bracket, left bracket 2 right bracket. The number, 9.

# Long description

A list of three numbers variables labeled, numbers left bracket 0 right bracket, numbers left bracket 1 right bracket, numbers left bracket 2 right bracket. Each variable has a corresponding box with the words, address. Arrows point from each address box to an array of references to arrays in a row of three boxes with numbers variables named as follows. The row corresponding to the variable numbers left bracket 0 right bracket reads as follows. Box 1, numbers left bracket 0 right bracket left bracket 0 right bracket. Box 2, numbers left bracket 0 right bracket left bracket 1 right bracket. Box 3, numbers left bracket 0 right bracket left bracket 2 right bracket.

The row corresponding to the variable numbers left bracket 1 right bracket reads as follows. Box 1, numbers left bracket 1 right bracket left bracket 0 right bracket. Box 2, numbers left bracket 1 right bracket left bracket 1 right bracket. Box 3, numbers left bracket 1 right bracket left bracket 2 right bracket.

The row corresponding to the variable numbers left bracket 2 right bracket reads as follows. Box 1, numbers left bracket 2 right bracket left bracket 0 right bracket. Box 2, numbers left bracket 1 right bracket left bracket 1 right bracket. Box 3, numbers left bracket 2 right bracket left bracket 2 right bracket.



# Long description

The six boxes are labeled as Element 0, Element 1, Element 2, Element 3, Element 4, and Element 5 and each have corresponding values of 1, 2, 7, 8, 9, and 5 respectively. Elements 1 and 2 swapped their values from Figure 7.26.

# Long description

The six boxes are labeled as Element 0, Element 1, Element 2, Element 3, Element 4, and Element 5 and each have corresponding values of 1, 2, 5, 8, 9, and 7 respectively. Elements 2 and 5 swapped their values from Figure 7.27.

# Long description

The six boxes are labeled as Element 0, Element 1, Element 2, Element 3, Element 4, and Element 5 and each have corresponding values of 1, 2, 5, 7, 9, and 8 respectively. Elements 3 and 5 swapped their values from Figure 7.28.

# Long description

The six boxes are labeled as Element 0, Element 1, Element 2, Element 3, Element 4, and Element 5 and each have corresponding values of 1, 2, 5, 7, 8, and 9 respectively. Elements 4 and 5 swapped their values from Figure 7.29.

# Long description

Reading clockwise, the top right box containing the number 2 has 2 arrows pointing right. The center right box containing the number 7 has 1 arrow pointing right. The bottom right box containing the number 6 has 1 arrow pointing right, 1 arrow pointing diagonally down, and 1 arrow pointing down. The bottom center box containing the number 1 has 1 arrow pointing down. The bottom left box containing the number 8 has 1 arrow pointing down

# Long description

Three instances of the Countable class, object 1, object 2, and object 3, each point a shared instance Count field left parenthesis, static, right parenthesis, which holds a value of 3.

# Long description

The first input dialog box prompts the user to, Enter a distance in miles, and the user enters the number 10 and clicks OK. A Message box contains the message, 10.00 miles equals 16.09 kilometers, and clicks OK. The second input dialog box prompts the user to, Enter a distance in kilometers, and the user enters the number 100 and clicks OK. A Message box contains the message, 100.00 kilometers equal 62.15 miles.

# Long description

The display Rectangle executes and both box and r reference the same object, the Rectangle reference variable that contains a length value of 12.0 and a width of 5.0. The method reads as follows. Line 1, public static void display Rectangle left parenthesis, Rectangle r, right parenthesis. Line 2, left brace. Line 3, double backward slash, Display the length and width. Line 4, System dot out dot print 1 n, left parenthesis, the word, Length colon, in quotation marks, + r dot get Length left parenthesis right parenthesis +. Line 5, left parenthesis, the word, Width colon, in quotation marks, + r dot get Width left parenthesis double right parenthesis semicolon. Line 6, right brace.



# Long description

The input dialog box prompts the user to, Enter the account balance. The user enters 3200.00 and clicks OK. The message dialog box contains the message, The account has a balance of \$3200.00.

# Long description

A method reads as follows. Line 1, public static Bank Account get Account left parenthesis right parenthesis. Line 2, left brace. Line 3 is separate from the method and reads, left parenthesis, statement appears here, right parenthesis. Line 4, return new Bank Account left parenthesis, balance, right parenthesis semicolon. Line 5, right brace. Line 4 contains the address that references the account = get Account left parenthesis right parenthesis semicolon. The address contains a Bank Account object with a value of a balance of 3200.00.

# Long description

The top section is titled, Stock. The middle section contains two fields. Field 1, minus symbol colon String. Field 2, minus share Price colon double. The third and bottom section contains four methods. Method 1, + Stock left parenthesis s y m colon String, price colon double, right parenthesis colon, Method 2, + get Symbol left parenthesis right parenthesis colon String. Method 3, + get Share Price left parenthesis right parenthesis colon double. Method 4, + to String left parenthesis right parenthesis colon String.

# Long description

An if statement compares two addresses. In Address 1, the company 1 variable holds the address of a Stock object, which has a symbol of, the letters, X Y Z, in quotation marks, and a share Price of 9.62. In Address 2, the company 2 variable holds the address of a Stock object, which has identical contents as the Stock object that the company 1 variable holds.

# Long description

The top section is titled, Instructor. The middle section contains three fields. Field 1, minus last Name colon String. Field 2, minus first Name colon String. Field 3, minus office Number colon String. The third and bottom section contains four methods. Method 1, + Instructor left parenthesis, l name colon String, f name colon String, office colon String right parenthesis colon Method 2, + Instructor left parenthesis, object 2 colon Instructor, right parenthesis colon Method 3, + set left parenthesis, l name colon String, f name colon String, office colon String, right parenthesis colon void. Method 4, + to String left parenthesis right parenthesis colon String.

# Long description

The top section is titled, Textbook. The middle section contains three fields. Field 1, minus title colon String. Field 2, minus author colon String. Field 3, minus publisher colon String. The third and bottom section contains four methods. Method 1, + Textbook left parenthesis, text Title colon String, a u t h colon String, pub colon String right parenthesis, colon Method 2, + Textbook left parenthesis, object 2 colon Text Book, right parenthesis colon Method 3, + set left parenthesis, text Title colon String, a u t h colon String, pub colon String, right parenthesis colon void. Method 4, + to String left parenthesis right parenthesis colon String.

# Long description

The top section is titled, Course. The middle section contains three fields. Field 1, minus course Name colon String. Field 2, minus instructor colon String. Field 3, minus textbook colon TextBook. The third and bottom section contains five methods. Method 1, + Course left parenthesis, name colon String, i n s t r colon Instructor, text colon Text Book right parenthesis colon Method 2, + get Name left parenthesis, right parenthesis colon String. Method 3, + get Instructor left parenthesis, right parenthesis colon Instructor. Method 4, + get Textbook left parenthesis, right parenthesis colon TextBook. Method 5, + to String left parenthesis, right parenthesis colon String.

# Long description

The Course class is listed first. The top section is titled, Course. The middle section contains three fields. Field 1, minus course Name colon String. Field 2, minus instructor colon String. Field 3, minus textbook colon TextBook. The third and bottom section contains five methods. Method 1, + Course left parenthesis, name colon String, i n s t r colon Instructor, text colon TextBook right parenthesis colon Method 2, + get Name left parenthesis, right parenthesis colon String. Method 3, + get instructor left parenthesis, right parenthesis colon Instructor. Method 4, + get Textbook left parenthesis, right parenthesis colon Textbook. Method 5, + to String left parenthesis, right parenthesis colon String. The Course class is connected to the Textbook and Instructor classes through a bar that links to the Instructor class on the left and the Textbook class on the right. The Instructor class is listed on the bottom left and reads as follows. The top section is titled, Instructor. The middle section contains three fields. Field 1, minus last Name colon String. Field 2, minus first Name colon String. Field 3, minus office Number colon String. The third and bottom section contains four methods. Method 1, + Instructor left parenthesis, l name colon String, f name colon String, office colon String right parenthesis colon Method 2, + Instructor left parenthesis, object 2 colon Instructor, right parenthesis colon Method 3, + set left parenthesis, l name colon String, f name colon String, office colon String, right parenthesis colon void. Method 4, + to String left parenthesis right parenthesis colon String. The TextBook class is listed on the bottom right and reads as follows. The top section is titled, TextBook. The middle section contains three fields. Field 1, minus title colon String. Field 2, minus author colon String. Field 3, minus publisher colon String. The third and bottom section contains four methods. Method 1, + TextBook left parenthesis, text Title colon String, a u t h colon String, pub colon String right parenthesis colon Method 2, + TextBook left parenthesis, object 2 colon TextBook, right parenthesis colon Method 3, + set left parenthesis, text Title colon String, a u t h colon String, pub colon String, right parenthesis colon void. Method 4, + to String left parenthesis right parenthesis colon String.



# Long description

The work Day variable holds the address of the Day dot WEDNESDAY object. Each of the seven objects are of the Day type, which is a specialized class, and read top to bottom as follow. Row 1. Day dot SUNDAY. Row 2. Day dot MONDAY. Row 3. Day dot TUESDAY. Row 4. Day dot WEDNESDAY. Row 5. Day dot THURSDAY. Row 6. Day dot FRIDAY. Row 7. Day dot SATURDAY.

# Long description

The title of the table in the C R C card is Stock Purchase. The left column lists Responsibilities and the right column lists Collaborating classes. The table reads as follows. Row 1. Responsibilities, Know the stock to purchase. Collaborating classes, Stock class. Row 2. Responsibilities, Know the number of shares to purchase. Collaborating classes, None. Row 3. Responsibilities, Calculate the cost of the purchase. Collaborating classes, Stock class.

# Long description

The top section is titled, Room Carpet. The middle section contains two fields. Field 1, minus size colon Room Dimension. Field 2, minus carpet Cost colon double. The third and bottom section contains three methods. Method 1, + Room Carpet left parenthesis dim colon Room Dimension, cost colon double, right parenthesis colon, Method 2, + get Total Cost left parenthesis right parenthesis colon double. Method 3, + to String left parenthesis right parenthesis colon String. The Room Carpet section is connected to the bottom section, Room Dimension, by a vertical line. The top section is titled, Room Dimension. The middle section contains two fields. Field 1, minus length colon double. Field 2, minus width colon double. The third and bottom section contains three methods. Method 1, + Room Dimension left parenthesis len colon double, w colon double, right parenthesis colon, Method 2, + get Area left parenthesis right parenthesis colon double. Method 3, + to String left parenthesis right parenthesis colon String.

# Long description

The input dialog box prompts, Enter any single character. The user enters the letter, lowercase a, and clicks OK. A message box is generated with the message, That is a letter. A second message box is generated with the message, That is a lowercase letter.

# Long description

The input dialog box prompts, Enter any single character. The user enters 4 and clicks OK, which generates a message box with the message, That is a digit.

# Long description

The input dialog box prompts, Enter a customer number in the form of L L N N N N, left parenthesis L L = letters and N N N N = numbers, right parenthesis. The user enters R Q S 4 5 6, and clicks OK. A message box with the message, That's a valid customer number, is generated.

# Long description

The input dialog box prompts, Enter a customer number in the form of L L N N N N, left parenthesis L L = letters and N N N N = numbers, right parenthesis. The user enters A X 4 6 7 T p, and clicks OK. A message box with the message, That is not the proper format of a customer number. Here is an example: A B C 1 2 3 4.

# Long description

The table reads as follows. Row 1. A, 87. B, 79. C, 91. D, 82. E, 94. Row 2. A, 72. B, 79. C, 81. D, 74, E, 88. Row 3. A, 94. B, 92. C, 81. D, 89. E, 96. Row 4. A, 77. B, 56. C, 67. D, 81. E, 79. Row 5. A, 79. B, 82. C, 85. D, 81. E, 90.



# Long description

In addition to the common insect characteristics, the bumblebee has its own unique characteristic such as the ability to sting. In addition to the common insect characteristics, the grasshoppers has its own unique characteristics such as the ability to jump.

# Long description

The top section is titled, Rectangle. The middle section contains minus score colon double. The bottom section contains three methods as follows. Plus Set Score left parenthesis, s colon double, right parenthesis, colon void. Plus Get Score left parenthesis right parenthesis colon double. Plus Get Grade left parenthesis right parenthesis colon char.

# Long description

The top section is titled, Final Exam. The middle section contains three variables as follows. Minus num Questions colon i n t. Minus points Each colon double. Minus num Missed colon i n t. The bottom section contains three methods as follows. Plus Final Exam left parenthesis, questions colon int, missed colon int, right parenthesis. Plus get Points Each left parenthesis right parenthesis colon double. Plus get Num Missed left parenthesis, right parenthesis, colon int.

# Long description

An input dialog box prompts, How many questions are on the final exam? The user enters the number 20 and clicks OK. Another input dialog box prompts, How many questions did the student miss? The user enters the number 3 and clicks OK. A message box is generated that contains the message, Each question counts 5.0 points. The exam score is 85.0. The exam grade is B.

# Long description

The top section is titled, Graded Activity. The middle section contains the variable minus score colon double. The bottom section contains three methods as follows. Plus Set Score left parenthesis, s colon double right parenthesis colon void. Plus get Score left parenthesis right parenthesis colon double. Plus get Grade left parenthesis, right parenthesis, colon char. The Final Exam class points to the Graded Activity class. The top section is titled, Final Exam. The middle section contains three variables as follows. Minus num Questions colon i n t. Minus points Each colon double. Minus num Missed colon int. The bottom section contains three methods as follows. Plus Final Exam left parenthesis, questions colon i n t, missed colon i n t, right parenthesis. Plus get Points Each left parenthesis right parenthesis colon double. Plus get Num Missed left parenthesis, right parenthesis, colon i n t.

# Long description

The top section is titled, Rectangle. The middle section contains the variables minus length colon double and minus width colon double. The bottom section contains six methods as follows. Plus Rectangle left parenthesis, len colon double, w colon double, right parenthesis. Plus set Length left parenthesis, len colon double, right parenthesis colon void. Plus set Width left parenthesis w colon double right parenthesis colon void. Plus get Length left parenthesis right parenthesis colon double. Plus get Width left parenthesis right parenthesis colon double. Plus get Area left parenthesis right parenthesis colon double.

# Long description

The top section is titled, Cube. The middle section contains the variable minus height colon double. The bottom section contains four methods as follows. Plus Cube left parenthesis, len colon double, w colon double, h colon double, right parenthesis. Plus get Height left parenthesis right parenthesis colon double. Plus get Surface Area left parenthesis right parenthesis colon double. Plus get Volume left parenthesis right parenthesis colon double. The Cube class points to the Rectangle class. The top section is titled, Rectangle. The middle section contains the variables minus length colon double and minus width colon double. The bottom section contains six methods as follows. Plus Rectangle left parenthesis, len colon double, w colon double, right parenthesis. Plus set Length left parenthesis, len colon double, right parenthesis colon void. Plus set Width left parenthesis w colon double right parenthesis colon void. Plus get Length left parenthesis right parenthesis colon double. Plus get Width left parenthesis right parenthesis colon double. Plus get Area left parenthesis right parenthesis colon double.

# Long description

The bottom class is titled Curved Activity in the top box. The second box contains the variables minus raw Score colon double and minus percentage colon double. The third box contains four methods that read as follows. Plus Curved Activity left parenthesis percent colon double right parenthesis. Plus set Score left parenthesis s colon double right parenthesis colon void. Plus get Raw Score left parenthesis right parenthesis colon double. Plus get Percentage left parenthesis right parenthesis colon double. The Curved Activity class points up to the Graded Activity class. The top section of the box is titled Graded Activity. The second section contains the variable minus score colon double. The third section consists of three methods that read as follows. Plus set Score left parenthesis s colon double right parenthesis colon void. Plus get Score left parenthesis right parenthesis colon double. Plus get Grade left parenthesis right parenthesis colon char.



# Long description

An interaction with the Protected Demo dot java program begins with an input dialog box that prompts, How many questions are on the final exam? The user enters the number 40 and hits OK. A second input dialog box prompts, How many questions did the student miss? The user enters the number 5 and clicks OK. A message box is generated that contains the message, Each question counts 2.5 points. The exam score is 88.0. The exam grade is a B.

# Long description

The top section is titled, Graded Activity 2. The middle section contains the variable pound sign score colon double. The bottom section contains three methods as follows. Plus Set Score left parenthesis, s colon double right parenthesis colon void. Plus get Score left parenthesis right parenthesis colon double. Plus get Grade left parenthesis, right parenthesis, colon char.

# Long description

Each rectangular box contains a title, such as, Class A, Class B, and Class C, and an empty rectangular box below the title section. The boxes are positioned from bottom to top as follows, Class C, Class B, and Class A. An arrow points from Class C to Class B and from Class B to Class A.

# Long description

The top section is titled, Graded Activity. The middle section contains the variable minus score colon double. The bottom section contains three methods as follows. Plus set Score left parenthesis, s colon double right parenthesis colon void. Plus get Score left parenthesis right parenthesis colon double. Plus get Grade left parenthesis, right parenthesis, colon char. The Pass Fail Activity class points to the Graded Activity. The top section is titled, Pass Fail Activity. The middle section contains the variable minus min Passing Score colon double. The bottom section contains two methods as follows. Plus Pass Fail Activity, left parenthesis m p s colon double, right parenthesis. Plus Get Grade, left parenthesis right parenthesis, colon char. The Pass Fail Exam class points to the Pass fail Activity class. The top section is titled, Pass Fail Exam. The middle section contains three variables that read as follows. Minus num Question colon int. Minus points Each colon double. Minus num Missed colon i n t. The bottom section contains three methods as follows. Plus Pass Fail Exam, left parenthesis, questions colon i n t, missed colon i n t, min Passing colon double, right parenthesis. Plus get Points Each left parenthesis right parenthesis colon double. Plus get Num Missed left parenthesis right parenthesis colon i n t.

# Long description

The hierarchy is organized as follows. Pass Fail Exam class points to the Pass Fail Activity class. Both are more specialized classes and positioned at the bottom, with Pass Fail Exam pointing upward. The Final Exam class is on the same level as the Pass Fail Activity and points towards the most general class, Graded Activity, positioned at the top.

# Long description

Each rectangular box contains a title and an empty rectangular box below the title section. The boxes are positioned from bottom to top as follows, Pass Fail Exam, Pass Fail Activity, Graded Activity, and Object. An arrow points from Pass Fail Exam to Pass Fail Activity, and from Pass Fail Activity to Graded Activity, and from Graded Activity to Object.

# Long description

The top section is titled, Cube. The middle section contains the variable minus height colon double. The bottom section contains four methods as follows. Plus Cube left parenthesis, len colon double, w colon double, h colon double, right parenthesis. Plus get Height left parenthesis right parenthesis colon double. Plus get Surface Area left parenthesis right parenthesis colon double. Plus get Volume left parenthesis right parenthesis colon double. The Cube class points to the Rectangle class. The top section is titled, Rectangle. The middle section contains the variables minus length colon double and minus width colon double. The bottom section contains six methods as follows. Plus Rectangle left parenthesis, len colon double, w colon double, right parenthesis. Plus set Length left parenthesis, len colon double, right parenthesis colon void. Plus set Width left parenthesis w colon double right parenthesis colon void. Plus get Length left parenthesis right parenthesis colon double. Plus get Width left parenthesis right parenthesis colon double. Plus get Area left parenthesis right parenthesis colon double.

# Long description

The diagram reads as follows. The top section is titled Student. The second section contains three variables that read as follows. Minus name colon String. Minus id Number colon string. Minus year Admitted colon i n t. The bottom section contains three variables that read as follows. Plus Student left parenthesis, n colon String, i d colon String, year colon i n t right parenthesis. Plus to String left parenthesis right parenthesis colon String. Plus get Remaining Hours left parenthesis right parenthesis colon i n t.



# Long description

The Final Exam class contains a top box titled Final Exam 3. The second box contains three variables as follows. Minus num Questions colon i n t. Minus points Each colon double. Minus num Missed colon i n t. The bottom box contains six methods as follows. Plus Final Exam, left parenthesis questions colon i n t, missed colon i n t, right parenthesis colon Plus Get Points Each left parenthesis right parenthesis colon double. Plus get Num Missed left parenthesis right parenthesis colon double. Plus equals, left parenthesis g colon Graded Activity right parenthesis colon Boolean. Plus is Greater left parenthesis g colon Graded Activity right parenthesis colon Boolean. Plus is Less , left parenthesis g colon Graded Activity right parenthesis colon Boolean. The Final Exam class points right to the double left angle bracket interface double right angle bracket Relatable class. The diagram consists of an empty second section and a bottom section that contains three methods that read as follows. Plus equals, left parenthesis g colon Graded Activity right parenthesis colon Boolean. Plus is Greater, left parenthesis g colon Graded Activity right parenthesis colon Boolean. Plus is Less , left parenthesis g colon Graded Activity right parenthesis colon Boolean. The Final Exam 3 class also points up to the Graded Activity class. The top box is titled Graded Activity. The second box contains the variable, minus score colon double. The bottom box contains three methods that read as follows. Plus set Score left parenthesis, s colon double, right parenthesis colon void. Plus get Score left parenthesis right parenthesis colon double. Plus get Grade left parenthesis right parenthesis colon char.

# Long description

The method reads as follows. `interface Calculator { int square = new int Calculator { left parenthesis right parenthesis left brace. public int calculate left parenthesis int number right parenthesis. Left brace. Return number * number; } double right brace, semicolon.` The interface reference variable is identified as the `int Calculator square =`, which creates an instance of an anonymous class that implements `int Calculator`. The method in the anonymous class is identified from `public int calculate left parenthesis int number right parenthesis` through `double right brace semicolon`.

# Long description

The hierarchy reads from the bottom layer to the top as follows. E O F Exception and File Not Found Exception form a level with a programming error represented by three periods on the far right of the bottom row. A line points from the bottom row to the next level and the I O Exception. On the far right of this level, the Runtime Exception has failed to connect because of the programming error. This level points to the Exception class and Error class, which both share a level that points to the higher Throwable class and ultimately the Object class at the top.

# Long description

The interaction begins with an input dialog box that prompts, Enter the name of a file. The user enters, Bad File dot t x t, and clicks OK. A message box is generated containing the message, File not found. A different message box is generated containing the message, Done.

# Long description

The interaction begins with an input dialog box that prompts, Enter the name of a file. The user enters, Good File dot t x t, and clicks OK. A message box is generated containing the message, The file was found. A different message box is generated containing the message, Done.

# Long description

The statement reads as follows. Try. Left brace. File = new File left parenthesis file Name right parenthesis semicolon Input File = new Scanner left parenthesis file right parenthesis semicolon J Option Pane dot show Message Dialog left parenthesis, null, the words, The file was found., in quotation marks, right parenthesis semicolon Right brace. Catch left parenthesis File Not found Exception e right parenthesis. Left brace. J Option Pane dot show Message Dialog left parenthesis null, the words, File not found., in quotation marks, right parenthesis semicolon Right brace. If the statement, input File = new Scanner left parenthesis file right parenthesis, throws an exception, then the following statement is skipped, from J Option Pane through Right brace. If the exception is an object of the File Not Found Exception class, the program jumps to the catch clause.

# Long description

The statement reads as follows. Try. Left brace. File = new File left parenthesis file Name right parenthesis semicolon Input File = new Scanner left parenthesis file right parenthesis semicolon J Option Pane dot show Message Dialog left parenthesis, null, the words, The file was found., in quotation marks, right parenthesis semicolon Right brace. Catch left parenthesis File Not found Exception e right parenthesis. Left brace. J Option Pane dot show Message Dialog left parenthesis null, the words, File not found., in quotation marks, right parenthesis semicolon Right brace. If no exception is thrown in the try block, the program jumps to the statement that immediately follows the try forward slash catch construct, J Option Pane dot show Message Dialog left parenthesis null, the words, File not found, in quotation marks, right parenthesis semicolon.

# Long description

The interaction begins with an input dialog box that prompts, Enter the name of a file. The user enters, Bad File dot t x t, and clicks OK. A message box is generated containing the message, Bad File dot t x t left parenthesis, The system cannot find the file specified, right parenthesis. A different message box is generated containing the message, Done.



# Long description

An input dialog box prompts, ERROR, Sales Data dot t x t does not exist. Enter another file name. The user enters the file name, Sales Info dot t x t and clicks OK. A message box is generated that contains the message, Number of months, 6. Total Sales, \$169,187.71. Average Sales, \$28,197.95.

# Long description

A message box contains the message, Non-numeric data found in the file. The invalid record will be skipped. A second message box is generated that contains the message, Number of months, 6. Total Sales, \$169,187.71. Average Sales, \$28,197.95.

# Long description

The hierarchy reads from the bottom of the top as follows. Number Format Exception. Illegal  
Argument Exception. Runtime Exception. Exception. Throwable. Object.

# Long description

The top section is titled, Bank Account. The middle section contains one field. Minus balance colon double. The bottom section contains 10 methods as follows. Methods 1, 2, and 3 are identified as overloaded constructors and read as follows. Plus Bank Account left parenthesis right parenthesis. Plus Bank Account left parenthesis, start Balance colon double right parenthesis. Plus Bank Account left parenthesis, s t r colon String right parenthesis. Methods 4 and 5 are identified as overloaded deposit methods and read as follows. Plus deposit left parenthesis, amount colon double, right parenthesis colon void. Plus deposit left parenthesis s t r colon String right parenthesis colon void. Methods 6 and 7 are identified as overloaded withdraw methods and read as follows. Plus withdraw left parenthesis, amount colon double, right parenthesis colon void. Plus withdraw left parenthesis, s t r colon String, right parenthesis colon void. Methods 8 and 9 are identified as overloaded set Balance methods and read as follows. Plus set Balance left parenthesis b colon double right parenthesis colon void. Plus set Balance left parenthesis s t r colon String right parenthesis colon void. Method 10 is not identified and reads as follows. Plus get Balance left parenthesis right parenthesis colon double.

# Long description

Items in a sequential access file are accessed one after the other as shown in a rectangle containing 14 boxes where arrows point to the next boxes. Items in a random-access file are accessed in any order as shown in an identical rectangle where arrows point in different directions across long and short spans.

# Long description

Each box is given a byte number, starting with 0, through 15. Character 0 is comprised of bytes 0 and 1. Character 1 is comprised of bytes 2 and 3. Character 2 is comprised of bytes 4 and 5. Character 3 is comprised of bytes 6 and 7. Character 4 is comprised of bytes 8 and 9. Character 5 is comprised of bytes 10 and 11. Character 6 is comprised of bytes 12 and 13. Character 7 is comprised of bytes 14 and 15. The words, and so forth, imply the file continues as such.

# Long description

Buttons for the Internet include, the local intranet, trusted sites, and restricted sites, for the user to click in response to a prompt, Select a zone to view or change security setting

# Long description

The prompts for the user to enter the, Number of Hours Worked, and enter the Hourly Pay Rate, are label controls. The user enters the answers to these prompts in two Text Field controls. The button to Calculate Gross Pay is the button control.



# Long description

The layout containers are arranged as follows. A window titled H Box Demo arranges Button 1, Button 2, Button 3, in a horizontal row. A window titled V Box Demo arranges Button 1, Button 2, and Button 3, vertically. A window titled Grid Pane Demo arranges Button 1, Button 2, and Button 3, in a horizontal row. A second row arranges Button 3, Button 4, and Button 5, in a horizontal row directly below the first row of buttons.

# Long description

The window contains the text, Hello World, in the center of the window, corresponding to the centered message Label control. The window has an approximate width of 300 pixels and height of 100 pixels.

# Long description

The images are arranged horizontally as follows, full moon, ship, mountain. Each image is arranged with 10 pixels of space between the cells and 30 pixels of padding around the grid.

# Long description

The images are arranged vertically as follows, full moon, ship, mountain. Each image is arranged with 10 pixels of space between the cells and 30 pixels of padding around the grid.

# Long description

The first line of text reads as follows, This is label 1, and, This is label 2, separated by a visible grid line, and the second line of which reads as follows, This is label 3, and, This is label 4, separated by a visible grid line.

# Long description

The Event Handler Object reads as follows. Void handle left parenthesis, Action Event event, right parenthesis. Left brace. Double backward slash, Event handling code is here, dot dot dot, Right brace.

# Long description

When the application starts, a box prompts, Click the button to see a message, with a button labeled, Click Me. After the user clicks the Button, a new box appears with the message, Thanks for clicking the button, and a button labeled, Click Me.

# Long description

The top layer of the application is the root node, known as the v box, connects to the second level of the h box, a branch node, which contains the prompt Label and kilo Text Field leaf nodes. The bottom layer contains the calc Button and result Label leaf nodes.



# Long description

The document is located in the center region. The top region contains the file name and different components of the application, such as, Home, and Layout. The left and right regions contain vertical navigation sections. The bottom region contains information such as word count and page count.

# Long description

Three images of fruit, Cherries, Oranges, and Apples, are aligned in three separate boxes in the center region. A user enters how much money he or she is inserting into the machine in a Text Field that reads, Amount Inserted, \$. The user enters an amount and clicks the button labeled, Spin. To the right, two amounts are calculated and displayed. The top reads, Amount Won This Spin, \$0.00, and the bottom reads, Total Amount Won, \$0.00.

# Long description

The Metric Converter application displays a window titled, Metric Converter, that prompts, Enter a distance in kilometers. The user clicks a Radio Button with three options given as follows, Convert to Miles, Convert to Feet, Convert to Inches, and clicks the Button control, labeled, Convert.

# Long description

In the first version, the user enters a distance of 100 in kilometers and converts the distance to 62.14 miles using the, Convert to Miles, radio button. In the second version, the user enters a distance of 100 in kilometers and converts the distance to 328,100.00 feet using the, Convert to Feet, radio button. In the third version, the user enters a distance of 100 in kilometers and converts the distance to 3,937,000.00 inches using the, Convert to Inches, radio button.

# Long description

Initially, the window displays the photograph of a flower. A radio button labeled, Flower, is selected. When the user selects the, Sunset, radio button, the image changes to the picture of a sunset in the right box.

# Long description

Three Check Box control options are labeled, Pepperoni \$3.00, Cheese \$2.00, Anchovies \$1.00. The user clicks a button labeled, Get Total, to calculate the total cost of the selected options.

# Long description

Three Check Box control options are labeled, Pepperoni \$3.00, Cheese \$2.00, Anchovies \$1.00. The user clicks a button labeled, Get Total, to calculate the total cost of the selected options. In the top left window box, no toppings are selected, and the total cost is \$0.00. In the top right window box, the, Pepperoni \$3.00, box is selected, and the total cost of \$3.00 is displayed. In the bottom left window box, the Pepperoni \$3.00, and, Cheese \$2.00, boxes are selected, and the total cost of \$5.00 is displayed. In the bottom right window box, all three boxes, Pepperoni \$3.00, Cheese \$2.00, and Anchovies \$1.00, are selected, and the total cost of \$6.00 is displayed.

# Long description

The list of dog breeds reads as follows, Poodle, Great Dane, Labrador, Terrier, and does not include a scroll bar. The list of cat breeds reads as follows, Siamese, Persian, Bobtail, Burmese, and does include a scroll bar.



# Long description

Each window contains an identical list of names that read as follows, Steve, Bethany, Will, Katrina, the prompt, Select a Name, and a button labeled, Get the Selection. In the second window, the name Bethany is selected in the list, and the name, Bethany, appears where the, Select a Name, prompt is chosen, after the, Get the Selection, button is clicked.

# Long description

A window contains a list of names that read as follows, Steve, Bethany, Will, Katrina, the prompt, Select a Name, and a button labeled, Get the Selection. When the user clicks the button without selected a name, a message reads, Nothing Selected.

# Long description

A window contains a list of names that read as follows, Steve, Bethany, Will, Katrina, the prompt, Select a Name, and a button labeled, Get the Selection. When the user selects the name, Bethany, and clicks the button, the name, Bethany, appears.

# Long description

In single selection mode, only one item can be selected at a time, as depicted in the first window when the month of March is selected. In multiple selection mode, multiple consecutive items can be selected, as depicted in the second window when the months of February, March, April, May, are selected. In multiple selection mode, multiple non-consecutive items can be selected, as depicted in the third window when the months of January, February, March and May are selected.

# Long description

On the left, a window displays the application as it initially appears. The top window contains a list from view 1 with the names January, February, March, April, May, and June. A window at the bottom will hold the names of the months that are selected. A button labeled, Get Selections, activates the Button control. On the right, a window displays the application after the user has selected, February, March, April, and June, and clicked the, Get Selections, button. The names of the months are displayed in the bottom box.

# Long description

The box on the left displays the program as it initially appears without any numbers selected. A List View contains the numbers 1 through 5 in multiple selection mode. The prompt, Select Some Numbers, appears above the button labeled, Get the Total. The box on the right displays the window after the user has selected four items and clicked the button, which generates the amount of 9 above the button.

# Long description

The top box window displays three playing card images in a horizontal List View as follows, Ace of Spade, King of Spades, Queen of Spades. A Label control prompts, Select a Card. The bottom box window displays the program after the user has clicked the Ace of Spades, the program displays the name of the card in the Label control.

# Long description

The first image displays the Combo Box's initial appearance as a drop down box. The middle image displays a list of items, in this case first names, that drop down when the user clicks the Combo Box. The selected name, Amanda, is displayed by the Combo Box.



# Long description

The first image displays the Combo Box's initial appearance. The prompt, Select a Country, appears below a drop down box. The middle image displays a list of items, in this case countries, that drop down when the user clicks the Combo Box. The name of the selected country, Scotland, is displayed by the Combo Box.

# Long description

The first image displays the Combo Box's initial appearance. The prompt, Select a Country, appears below a drop down box. The middle image displays a list of items, in this case countries, that drop down when the user clicks the Combo Box. The name of the selected country, Scotland, is displayed by the Combo Box, and the name, Scotland, appears beneath the drop down box.

# Long description

In the Uneditable Combo Box, the user may only click the Combo Box button and select an item from the dropdown list. In the Editable Combo Box, the user can type an input into a field or alternately click the dropdown arrow button to select an item from the dropdown list.

# Long description

The horizontal slider's track is marked at values of 0, 10, 20, 30, 40, and 50, with the slider control in the center. The vertical slider's track is marked at distances of 1, 0.75, 0.5, 0.25, 0, minus 0.25, minus 0.5, minus 0.75, and minus 1, with the slider control in the center.

# Long description

The Slider's range allows the user to adjust the Celsius temperature from 0 through 100 and instantly see the conversion to Fahrenheit calculated in the Fahrenheit field. Since the slider is set to 0, the Fahrenheit field and the Celsius field have values of 32.0 and 0.0, respectively.

# Long description

The window reads as follows. The word, Applet, is located in the top left corner. Two lines cross in the center of the window and extend to the corners in the shape of the letter X. The text, Applet started, is located in the bottom left corner.

# Long description

The window reads as follows. The word, Applet, is located in the top left corner. The top left of the window contains an empty white square box is bordered by a black line. The bottom right of the window contains a gray square box. The text, Applet started, is located in the bottom left corner.

# Long description

The window reads as follows. The word, Applet, is located in the top left corner. The top left of the window contains an empty white square box is bordered by a black line. The bottom right of the window contains a gray square box. The text, Applet started, is located in the bottom left corner.



# Long description

The figure appears as follows. The box is divided into four quadrants by four lines that point out beyond the box's border and are labeled clockwise as follows. 90 degrees, 0 degrees, 270 degrees, 180 degrees. The circle is bordered by a dashed line except in the quadrant from 90 degrees to 0 degrees, where the circle's border is a solid, black line.

# Long description

The window reads as follows. The word, Applet, is located in the top left corner. The top left corner of the window contains the curved line border that makes up the top right section of a circle. In the top right corner of the window, the top section of the circle is filled in and shaded extending from a 90 degree angle to the curved border of the circle. On the bottom left corner of the window, a smaller wedge of the circle's border is shown in a light gray curved line. On the bottom right corner, the angle of the circle on the left is filled in and shaded extending from 45 degree angle. The text, Applet started, is located in the bottom left corner.

# Long description

The eight points read clockwise as follows. First point, 60, 20. Second point, 10, 20. Third point, 140, 60. Fourth point, 140, 100. Fifth point, 100, 140. Sixth point, 60, 140. Seventh point, 20, 100. Eighth point, 20, 60.

# Long description

The window reads as follows. The word, Applet, is located in the top left corner. A gray octagon with the word, S T O P, in white text is displayed in the window. The text, Applet started, is located in the bottom left corner.

# Long description

The window reads as follows. The word, Applet, is located in the top left corner. On the right side of the window, a list of 7 checkbox options read as follows. Line, Rectangle, Filled Rectangle, Oval, Filled Oval, Arc, Filled Arc. The text, Applet started, is located in the bottom left corner.

# Long description

The window reads as follows. The word, Applet, is located in the top left corner. On the right side of the window, a list of 7 checkbox options read as follows. Line, Rectangle, Filled Rectangle, Oval, Filled Oval, Arc, Filled Arc. The window contains a large gray box and an empty white box with a black border behind it and a gray circle in front of it. A long diagonal line stretches through these shapes. In the top right of the window, a line outlines the top left section of a circle. In the bottom right of the window, an empty circle is outlined by a black white line. To the top left of the circle, the shaded top right of a circle is shaded. The text, Applet started, is located in the bottom left corner.

# Long description

The window reads as follows. The word, Applet, is located in the top left corner. Two rows of 3 shaded text fields read as follows. Top row. Pressed, Clicked, Released. Bottom row. Entered, Exited, Dragged. Two boxes read the values, X: 83, and, Y: 126. The text, Applet started, is located in the bottom left corner.

# Long description

The window reads as follows. The word, Applet, is located in the top left corner. A large square box with a black line border is displayed in the center of the window. The text, Applet started, is located in the bottom left corner.



# Long description

The window reads as follows. The word, Applet, is located in the top left corner. A small gray circle is displayed in center of the window. The text, Applet started, is located in the bottom left corner.

# Long description

The window reads as follows. The word, Applet, is located in the top left corner. A line of bold text reads, Audio source: N A S A. To the left are three buttons labeled, Play, Loop, Stop. The text, Applet started, is located in the bottom left corner.

# Long description

The drawing appears as follows. Each house has a rectangular roof and a front door flanked on the left and the right by windows. The house on the left has its door and windows open, displayed as shaded in full. The house on the right has its door and windows closed. The front door has a doorknob. Each window is shown as a box that is evenly divided into two rows of two windows.

# Long description

The illustration reads clockwise as follows. Top pair of eyes, Mouse not in the window, eyes looking ahead. Second pair of eyes, Mouse above the eyes, eyes looking up. A cursor points up and is positioned above the eyes. Third pair of eyes, Mouse to the right of the eyes, eyes looking right. A cursor points up and is positioned to the right of the eyes. Fourth pair of eyes, Mouse to the left of the eyes, eyes looking left. A cursor points up and is positioned to the left of the eyes. Fifth pair of eyes, Mouse below the eyes, eyes looking down. A cursor points up and is positioned down and in front of the eyes.

# Long description

The page contains an index divided into two rows of three topics as follows. First row. Get Started, Tools, Language. Second row. Specifications, Core Libraries, Security.

# Long description

The page is titled, Java Platform Standard Edition and Java Development Kit, Version 9 A P I Specification. The page's three sections are introduced as follows: Java S E, J D K, Java F X.

# Long description

The structure reads as follows. Beneath the C L A S S P A T H, written as, C colon forward slash file, the my Packages folder contains a vehicles folder and an employees folder, which contains an hourly folder and a salaried folder.

# Long description

The structure reads as follows. Beneath the C L A S S P A T H, written as, C colon forward slash file, the my Packages folder contains a vehicles folder and an employees folder, which contains an hourly folder and a salaried folder. The hourly folder also contains the Clerical dot java and Assembly Line java files.



# Long description

The page is arranged as follows. The page is titled, Apache Derby: Downloads and an index of 6 links to the contents of the page, which read as follows. For Java 8 and Higher. For Java 6 and Higher. For Java 1 point 4 and Higher. For Java 1 point 3 and Higher. Deprecated Releases. Change History. The entirety of the For Java 8 and Higher section and For Java 6 and Higher sections are visible, as is the header for the For Java 1 point 4 and Higher section.

# Long description

The two sections of the page are introduced as Distributions and Release Notes for Apache Derby 10 point 14 point 1 point 0. Beneath the Distribution section, the file d b dash derby dash 10 point 14 point 1 point 0 dash bin dot zip is circled with an arrow pointing to it from the left.

# Long description

Each rectangle object variable has two member variables: length and width. The address r1 points to a rectangle object variable whose length assigned 5 and width is assigned 20. The address r2 points to a rectangle object variable whose length assigned 10 and width is assigned 15.

# Long description

The Circle (subclass) is inherited from the Shape (superclass). The UML diagram is as follows. The subclass is at the bottom and the superclass is at the top. The inheritance is indicated using an upward arrow that connects the subclass with the superclass. The Circle object has a double type variable, radius. It has two methods: setRadius(r : double) : void and getRadius() : double. The Shape object has a double type variable, area. It has two methods: setArea(a : double) :void and getArea() : double.

# Long description

The UML diagram for the class Pet is shown. The class has three variables: name, animal, and age. Name and animal are string type. Age is integer type. It has six methods and they are setName(n : String) : void, setAnimal(a : String) : void, setAge(a : int) : void, getName() : String, getAnimal() : String, and getAge() : int.

# Long description

The interaction is displayed as follows. The Input dialog box prompts, Enter the amount of monthly sales. The user enters \$19600 and clicks O K. The next box prompts, Enter the amount of advanced pay. The user enters \$1000 and clicks O K. A Message box with the message, Commission rate is 15 percent. The amount of pay is \$1,940, is generated.

# Long description

The sections read as follows. The top section contains 7 fields as follows. Negative loan Amount colon double. Negative Interest Rate colon double. Negative loan Balance colon double. Negative term colon double. Negative payment colon double. Negative loan Years colon i n t. The bottom section contains 7 fields as follows. Plus Amortization left parenthesis colon double, rate colon double, years colon i n t right parenthesis. Negative c a l c Payment left parenthesis right parenthesis colon void. Plus get Number Of Payments left parenthesis right parenthesis colon i n t. Plus save Report left parenthesis file name colon String right parenthesis colon void. Plus get Loan Amount left parenthesis right parenthesis colon double. Plus get Interest Rate left parenthesis right parenthesis colon double. Plus get Loan Years left parenthesis right parenthesis colon i n t.

# Long description

The program is displayed as follows. The first Input dialog box prompts, Enter the loan amount. The user enters the amount 5000 and clicks O K. The second Input dialog prompts, Enter the annual interest rate. The user enters 0 point 059 and clicks O K. The third Input dialog box prompts, Enter the years of the loan. The user enters the amount 1 and clicks O K, which generates a Message dialog box with the message, Report saved to the file, Loan Amortization dot t x t. The user clicks OK. A final Input dialog box prompts, Would you like to run another report? Enter Y for yes or N for no. The user enters the letter, n, and clicks OK.



# Long description

Two rectangular rows are divided into 5 equal sections and read as follows. Top row. Names left bracket 0 right bracket. Names left bracket 1 right bracket. Names left bracket 2 right bracket. Names left bracket 3 right bracket. Names left bracket 4 right bracket. Bottom row. Addresses left bracket 0 right bracket. Addresses left bracket 1 right bracket. Addresses left bracket 2 right bracket. Addresses left bracket 3 right bracket. Addresses left bracket 4 right bracket. Between the top and bottom rows are 5 labels that read as follows. Person number 1. Person number 2. Person number 3. Person number 4. Person number 5. Each label is assigned between the corresponding names and addresses, such as Person number 1 appears between the cells for names left bracket 0 right bracket on the top row and addresses left bracket 0 right bracket. Arrows point up and down from the top and bottom of each label.

# Long description

The U M L diagram is divided into two sections that read as follows. The top section contains two fields that read as follows. Negative feet colon i n t. Negative inches colon i n t. The bottom section contains 11 fields that read as follows. Plus Feet Inches left parenthesis right parenthesis colon. Plus Feet Inches left parenthesis f colon i n t comma i colon i n t right parenthesis colon. Plus Feet Inches left parenthesis object 2 colon Feet Inches left parenthesis colon. Negative simplify left parenthesis right parenthesis colon void. Plus set Feet left parenthesis f colon i n t right parenthesis colon void. Plus set Inches left parenthesis f colon i n t right parenthesis colon void. Plus get Feet left parenthesis right parenthesis colon i n t. Plus get Inches left parenthesis right parenthesis colon i n t. Plus to String left parenthesis right parenthesis colon i n t. Plus add left parenthesis object 2 colon Feet Inches right parenthesis colon Feet Inches. Plus equals left parenthesis object 2 colon Feet Inches right parenthesis colon boolean.

# Long description

The if statement compares two addresses as follows. The distance 1 variable holds the address of a Feet Inches object, which holds values of 6 Feet and 5 Inches. The distance 2 variable holds the address of a Feet Inches object, which holds values of 6 Feet and 5 Inches.

# Long description

The top section consists of 4 fields that read as follows. Negative first colon string. Negative second colon String. Negative third colon. Negative valid colon boolean. The bottom section consists of 6 fields that read as follows. Plus Serial Numbers left parenthesis s n colon String right parenthesis colon. Plus is Valid left parenthesis right parenthesis colon boolean. Negative validate left parenthesis right parenthesis colon void. Negative is First Group Valid left parenthesis colon boolean. Negative is Second Group Valid colon boolean. Negative is Third Group Valid colon Boolean.

# Long description

Beneath the File menu are 5 options that read as follows. New, Open, Save, Save As, Exit. Beneath the Text menu are two sections. The first section consists of 3 radio buttons that read as follows. Monospaced. Serif. Sans Serif. Currently, the Sans Serif button is selected. A divider separates the bottom section of the Text menu and consists of two checkmark boxes for Italic and Bold.

# Long description

The text is displayed as follows. In the first window, the text is set in Sans Serif font, plain style. In the second window, the text is set in the Monospaced font, bold style. In the third window, the text is set in the Serif font, bold and italic style.