Java Just in Time: Collected concepts after chapter 21

John Latham, School of Computer Science, Manchester University, UK.

April 15, 2011

DRAFT April 15, 2011

Contents

1 Computer basics **21000**
1.1 Computer basics: hardware (page 3) . . . 21000
1.2 Computer basics: hardware: processor (page 3) . . . 21000
1.3 Computer basics: hardware: memory (page 3) . . . 21000
1.4 Computer basics: hardware: persistent storage (page 3) . . . 21001
1.5 Computer basics: hardware: input and output devices (page 3) . . . 21001
1.6 Computer basics: software (page 3) . . . 21001
1.7 Computer basics: software: machine code (page 3) . . . 21001
1.8 Computer basics: software: operating system (page 4) . . . 21001
1.9 Computer basics: software: application program (page 4) . . . 21002
1.10 Computer basics: data (page 3) . . . 21002
1.11 Computer basics: data: files (page 5) . . . 21002
1.12 Computer basics: data: files: text files (page 5) . . . 21002
1.13 Computer basics: data: files: binary files (page 5) . . . 21003

2 Java tools **21003**
2.1 Java tools: text editor (page 5) . . . 21003
2.2 Java tools: javac compiler (page 9) . . . 21003
2.3 Java tools: java interpreter (page 9) . . . 21004
2.4 Java tools: javadoc (page 223) . . . 21004
2.5 Java tools: javadoc: throws tag (page 355) . . . 21005

3 Operating environment **21006**
3.1 Operating environment: programs are commands (page 7) . . . 21006
3.2 Operating environment: standard output (page 7) . . . 21006
3.3 Operating environment: command line arguments (page 8) . . . 21006
3.4 Operating environment: standard input (page 187) . . . 21006
3.5 Operating environment: standard error (page 344) . . . 21006

4 Class **21007**

4.1 Class: programs are divided into classes (page 16) 21007
4.2 Class: public class (page 16) 21007
4.3 Class: definition (page 16) 21007
4.4 Class: objects: contain a group of variables (page 158) 21007
4.5 Class: objects: are instances of a class (page 158) 21008
4.6 Class: objects: this reference (page 180) 21008
4.7 Class: objects: may be mutable or immutable (page 193) 21009
4.8 Class: objects: compareTo() (page 222) 21009
4.9 Class: is a type (page 161) 21009
4.10 Class: is a type: and has three components (page 512) 21010
4.11 Class: making instances with new (page 162) 21010
4.12 Class: accessing instance variables (page 164) 21011
4.13 Class: importing classes (page 188) 21011
4.14 Class: stub (page 191) 21012
4.15 Class: extending another class (page 245) 21012
4.16 Class: generic class (page 491) 21013
4.17 Class: generic class: bound type parameter (page 496) 21013
4.18 Class: generic class: bound type parameter: extends some class (page 496)21014
4.19 Class: generic class: bound type parameter: extends some interface (page 526)21014
4.20 Class: generic class: where type parameters cannot be used (page 501)21015
4.21 Class: generic class: used as a raw type (page 502) 21015

5 Method **21016**

5.1 Method (page 118) 21016
5.2 Method: main method: programs contain a main method (page 17) . 21016
5.3 Method: main method: is public (page 17) 21016
5.4 Method: main method: is static (page 17) 21016
5.5 Method: main method: is void (page 17) 21017
5.6 Method: main method: is the program starting point (page 17) 21017
5.7 Method: main method: always has the same heading (page 18) . . . 21017
5.8 Method: private (page 118) 21018
5.9 Method: accepting parameters (page 118) 21018
5.10 Method: accepting parameters: of a class type (page 164) 21019
5.11 Method: accepting parameters: of an array type (page 297) 21019
5.12 Method: calling a method (page 119) 21019
5.13 Method: void methods (page 120) 21020
5.14 Method: returning a value (page 122) 21021
5.15 Method: returning a value: of a class type (page 176) 21021
5.16 Method: returning a value: multiple returns (page 196) 21022
5.17 Method: returning a value: of an array type (page 312) 21023
5.18 Method: changing parameters does not affect arguments (page 124) . 21023
5.19 Method: changing parameters does not affect arguments: but referenced objects can be ch
5.20 Method: constructor methods (page 159) 21024
5.21 Method: constructor methods: more than one (page 203) 21025
5.22 Method: constructor methods: more than one: using this (page 393) . 21025
5.23 Method: constructor methods: default (page 425) 21026
5.24 Method: class versus instance methods (page 166) 21027

DRAFT April 15, 2011

5.25 Method: a method may have no parameters (page 173) 21029
5.26 Method: return with no value (page 206) 21029
5.27 Method: accessor methods (page 207) 21029
5.28 Method: mutator methods (page 207) 21030
5.29 Method: overloaded methods (page 237) 21030
5.30 Method: that throws an exception (page 354) 21030
5.31 Method: that throws an exception: RuntimeException (page 358) . . 21031
5.32 Method: generic methods (page 522) 21032
5.33 Method: generic methods: bound type parameter (page 526) 21033

6 Command line arguments **21034**
6.1 Command line arguments: program arguments are passed to main (page 17)21034
6.2 Command line arguments: program arguments are accessed by index (page 26)21035
6.3 Command line arguments: length of the list (page 79) 21035
6.4 Command line arguments: list index can be a variable (page 79) . . . 21035

7 Type **21035**
7.1 Type (page 36) 21035
7.2 Type: String (page 135) 21036
7.3 Type: String: literal (page 18) 21036
7.4 Type: String: literal: must be ended on the same line (page 21) . . . 21036
7.5 Type: String: literal: escape sequences (page 49) 21036
7.6 Type: String: concatenation (page 26) 21037
7.7 Type: String: conversion: from int (page 38) 21037
7.8 Type: String: conversion: from double (page 55) 21038
7.9 Type: String: conversion: from object (page 177) 21038
7.10 Type: String: conversion: from object: null reference (page 211) . . 21039
7.11 Type: int (page 36) 21040
7.12 Type: double (page 54) 21040
7.13 Type: casting an int to a double (page 79) 21040
7.14 Type: boolean (page 133) 21040
7.15 Type: long (page 145) 21041
7.16 Type: short (page 145) 21041
7.17 Type: byte (page 145) 21041
7.18 Type: char (page 145) 21041
7.19 Type: char: literal (page 145) 21042
7.20 Type: char: literal: escape sequences (page 146) 21042
7.21 Type: char: comparisons (page 238) 21042
7.22 Type: char: casting to and from int (page 238) 21043
7.23 Type: float (page 146) 21043
7.24 Type: primitive versus reference (page 162) 21044
7.25 Type: array type (page 287) 21044
7.26 Type: enum type (page 309) 21044
7.27 Type: enum type: access from another class (page 312) 21045

8 Standard API **21045**
8.1 Standard API: System: out.println() (page 18) 21045

DRAFT April 15, 2011

8.2 Standard API: System: out.println(): with no argument (page 98) . . 21045
8.3 Standard API: System: out.println(): with any argument (page 427) . 21046
8.4 Standard API: System: out.print() (page 98) 21046
8.5 Standard API: System: out.printf() (page 126) 21047
8.6 Standard API: System: out.printf(): zero padding (page 140) 21048
8.7 Standard API: System: out.printf(): string item (page 289) 21049
8.8 Standard API: System: out.printf(): fixed text and many items (page 289)21049
8.9 Standard API: System: out.printf(): left justification (page 300) . . . 21050
8.10 Standard API: System: in (page 187) 21050
8.11 Standard API: System: in: is an InputStream (page 452) 21050
8.12 Standard API: System: getProperty() (page 195) 21050
8.13 Standard API: System: getProperty(): line.separator (page 195) . . . 21051
8.14 Standard API: System: currentTimeMillis() (page 262) 21051
8.15 Standard API: System: err.println() (page 344) 21051
8.16 Standard API: System: out: is an OutputStream (page 468) 21051
8.17 Standard API: System: err: is an OutputStream (page 468) 21051
8.18 Standard API: Integer: parseInt() (page 41) 21052
8.19 Standard API: Integer: as a box for int (page 487) 21052
8.20 Standard API: Integer: as a box for int: autoboxing (page 494) 21052
8.21 Standard API: Integer: as a box for int: works with collections (page 548)21053
8.22 Standard API: Double: parseDouble() (page 54) 21053
8.23 Standard API: Math: pow() (page 73) 21054
8.24 Standard API: Math: abs() (page 87) 21054
8.25 Standard API: Math: PI (page 87) 21054
8.26 Standard API: Math: random() (page 205) 21054
8.27 Standard API: Math: round() (page 289) 21055
8.28 Standard API: Scanner (page 188) 21055
8.29 Standard API: Scanner: for a file (page 306) 21056
8.30 Standard API: String (page 233) 21057
8.31 Standard API: String: some instance methods (page 234) 21057
8.32 Standard API: String: format() (page 301) 21058
8.33 Standard API: String: split() (page 313) 21059
8.34 Standard API: String: implements Comparable (page 520) 21059
8.35 Standard API: Character (page 342) 21060
8.36 Standard API: Object (page 422) 21061
8.37 Standard API: Object: toString() (page 427) 21061
8.38 Standard API: Object: equals() (page 521) 21061
8.39 Standard API: Object: hashCode() (page 548) 21062
8.40 Standard API: Object: hashCode(): making a good definition (page 566)21062
8.41 Standard API: Arrays (page 518) 21062
8.42 Standard API: Arrays: sort() (page 518) 21062
8.43 Standard API: Arrays: copyOf() (page 523) 21063
8.44 Standard API: Comparable interface (page 520) 21064
8.45 Standard API: Comparable interface: compareTo() and equals() (page 522)21064

9 Statement **21065**

9.1 Statement (page 18) 21065

9.2 Statement: simple statements are ended with a semi-colon (page 18) 21065
9.3 Statement: assignment statement (page 37) 21065
9.4 Statement: assignment statement: assigning a literal value (page 37) 21065
9.5 Statement: assignment statement: assigning an expression value (page 38)21065
9.6 Statement: assignment statement: updating a variable (page 70) . . . 21066
9.7 Statement: assignment statement: updating a variable: shorthand operators (page 87)2106
9.8 Statement: assignment statement: is an expression (page 450) 21067
9.9 Statement: if else statement (page 60) 21068
9.10 Statement: if else statement: nested (page 62) 21069
9.11 Statement: if statement (page 64) 21069
9.12 Statement: compound statement (page 66) 21070
9.13 Statement: while loop (page 71) 21071
9.14 Statement: for loop (page 77) 21072
9.15 Statement: for loop: multiple statements in for update (page 136) . . 21073
9.16 Statement: statements can be nested within each other (page 92) . . . 21073
9.17 Statement: switch statement with breaks (page 107) 21074
9.18 Statement: switch statement without breaks (page 110) 21074
9.19 Statement: do while loop (page 112) 21076
9.20 Statement: for-each loop: on arrays (page 293) 21076
9.21 Statement: for-each loop: on collections (page 562) 21078
9.22 Statement: try statement (page 344) 21079
9.23 Statement: try statement: with multiple catch clauses (page 347) . . 21080
9.24 Statement: try statement: with finally (page 451) 21082
9.25 Statement: throw statement (page 350) 21082

10 Error 21083
10.1 Error (page 20) 21083
10.2 Error: syntactic error (page 20) 21083
10.3 Error: semantic error (page 22) 21083
10.4 Error: compile time error (page 22) 21084
10.5 Error: run time error (page 24) 21084
10.6 Error: logical error (page 29) 21084

11 Execution 21085
11.1 Execution: sequential execution (page 23) 21085
11.2 Execution: conditional execution (page 60) 21085
11.3 Execution: repeated execution (page 70) 21085
11.4 Execution: parallel execution – threads (page 253) 21085
11.5 Execution: parallel execution – threads: the GUI event thread (page 254)21086
11.6 Execution: event driven programming (page 254) 21086

12 Code clarity 21087
12.1 Code clarity: layout (page 31) 21087
12.2 Code clarity: layout: indentation (page 32) 21087
12.3 Code clarity: layout: splitting long lines (page 43) 21088
12.4 Code clarity: comments (page 82) 21088
12.5 Code clarity: comments: marking ends of code constructs (page 83) . 21089

DRAFT April 15, 2011

12.6 Code clarity: comments: multi-line comments (page 189) 21089

13 Design **21090**

13.1 Design: hard coding (page 36) 21090
13.2 Design: pseudo code (page 73) 21090
13.3 Design: object oriented design (page 184) 21090
13.4 Design: object oriented design: noun identification (page 185) 21091
13.5 Design: object oriented design: encapsulation (page 187) 21091
13.6 Design: Sorting a list (page 295) 21092
13.7 Design: Sorting a list: bubble sort (page 296) 21092
13.8 Design: Sorting a list: total order (page 516) 21094
13.9 Design: Sorting a list: tree sort (page 554) 21094
13.10 Design: Searching a list: linear search (page 323) 21095
13.11 Design: Searching a list: binary search (page 525) 21095
13.12 Design: UML (page 381) 21096
13.13 Design: UML: class diagram (page 381) 21096
13.14 Design: Storing data (page 547) 21096
13.15 Design: Storing data: hash table (page 547) 21097
13.16 Design: Storing data: ordered binary tree (page 552) 21097
13.17 Design: Storing data: linked list (page 557) 21098

14 Variable **21099**

14.1 Variable (page 36) 21099
14.2 Variable: int variable (page 37) 21100
14.3 Variable: a value can be assigned when a variable is declared (page 42)21100
14.4 Variable: double variable (page 54) 21100
14.5 Variable: can be defined within a compound statement (page 92) . . 21101
14.6 Variable: local variables (page 124) 21101
14.7 Variable: class variables (page 124) 21102
14.8 Variable: a group of variables can be declared together (page 129) . . 21102
14.9 Variable: boolean variable (page 133) 21102
14.10 Variable: char variable (page 145) 21104
14.11 Variable: instance variables (page 159) 21104
14.12 Variable: instance variables: should be private by default (page 175) 21105
14.13 Variable: of a class type (page 161) 21105
14.14 Variable: of a class type: stores a reference to an object (page 162) . 21105
14.15 Variable: of a class type: stores a reference to an object: avoid misunderstanding (page 17
14.16 Variable: of a class type: null reference (page 192) 21108
14.17 Variable: of a class type: holding the same reference as some other variable (page 216)21
14.18 Variable: final variables (page 194) 21111
14.19 Variable: final variables: class constant (page 205) 21111
14.20 Variable: final variables: class constant: a set of choices (page 308) . 21112
14.21 Variable: final variables: class constant: a set of choices: dangerous (page 308)21112
14.22 Variable: of an array type (page 287) 21112
14.23 Variable: initial value (page 453) 21113

15 Expression **21113**

DRAFT April 15, 2011

15.1 Expression: arithmetic (page 38) 21113
15.2 Expression: arithmetic: int division truncates result (page 52) 21114
15.3 Expression: arithmetic: associativity and int division (page 52) . . . 21114
15.4 Expression: arithmetic: double division (page 55) 21114
15.5 Expression: arithmetic: double division: by zero (page 291) 21115
15.6 Expression: arithmetic: remainder operator (page 149) 21115
15.7 Expression: arithmetic: shift operators (page 473) 21115
15.8 Expression: arithmetic: integer bitwise operators (page 474) 21116
15.9 Expression: brackets and precedence (page 45) 21116
15.10 Expression: associativity (page 48) 21117
15.11 Expression: boolean (page 60) 21118
15.12 Expression: boolean: relational operators (page 60) 21119
15.13 Expression: boolean: logical operators (page 128) 21119
15.14 Expression: boolean: logical operators: conditional (page 323) . . . 21121
15.15 Expression: conditional expression (page 94) 21121

16 Package **21122**
16.1 Package (page 187) 21122
16.2 Package: java.util (page 188) 21122
16.3 Package: java.awt and javax.swing (page 245) 21122

17 GUI API **21123**
17.1 GUI API: JFrame (page 245) 21123
17.2 GUI API: JFrame: setTitle() (page 246) 21123
17.3 GUI API: JFrame: getContentPane() (page 246) 21123
17.4 GUI API: JFrame: setDefaultCloseOperation() (page 247) 21123
17.5 GUI API: JFrame: pack() (page 247) 21124
17.6 GUI API: JFrame: setVisible() (page 248) 21124
17.7 GUI API: Container (page 246) 21124
17.8 GUI API: Container: add() (page 246) 21125
17.9 GUI API: Container: add(): adding with a position constraint (page 268)21125
17.10 GUI API: Container: setLayout() (page 250) 21125
17.11 GUI API: JLabel (page 246) 21125
17.12 GUI API: JLabel: setText() (page 258) 21125
17.13 GUI API: LayoutManager (page 249) 21126
17.14 GUI API: LayoutManager: FlowLayout (page 249) 21126
17.15 GUI API: LayoutManager: FlowLayout: alignment (page 278) . . . 21126
17.16 GUI API: LayoutManager: GridLayout (page 251) 21127
17.17 GUI API: LayoutManager: BorderLayout (page 267) 21127
17.18 GUI API: Listeners (page 254) 21128
17.19 GUI API: Listeners: ActionListener interface (page 257) 21130
17.20 GUI API: Listeners: ActionListener interface: actionPerformed() (page 258)21131
17.21 GUI API: JButton (page 256) 21131
17.22 GUI API: JButton: addActionListener() (page 256) 21131
17.23 GUI API: JButton: setEnabled() (page 266) 21131
17.24 GUI API: JButton: setText() (page 267) 21132
17.25 GUI API: ActionEvent (page 258) 21132

DRAFT April 15, 2011

17.26	GUI API: ActionEvent: getSource() (page 280)	21132
17.27	GUI API: JTextField (page 265)	21132
17.28	GUI API: JTextField: getText() (page 265)	21132
17.29	GUI API: JTextField: setText() (page 265)	21132
17.30	GUI API: JTextField: setEnabled() (page 267)	21133
17.31	GUI API: JTextField: initial value (page 274)	21133
17.32	GUI API: JTextArea (page 267)	21133
17.33	GUI API: JTextArea: setText() (page 269)	21133
17.34	GUI API: JTextArea: append() (page 269)	21133
17.35	GUI API: JPanel (page 270)	21134
17.36	GUI API: JScrollPane (page 274)	21134
17.37	GUI API: Color (page 400)	21134
18	**Interface**	**21135**
18.1	Interface (page 257)	21135
18.2	Interface: definition (page 511)	21135
18.3	Interface: is a type (page 512)	21137
18.4	Interface: method implementation (page 513)	21137
18.5	Interface: generic interface (page 520)	21137
18.6	Interface: extending another interface (page 526)	21137
18.7	Interface: a class can implement many interfaces (page 530)	21138
19	**Array**	**21138**
19.1	Array (page 286)	21138
19.2	Array: array creation (page 287)	21139
19.3	Array: array creation: initializer (page 320)	21139
19.4	Array: element access (page 288)	21140
19.5	Array: element access: in two-dimensional arrays (page 330)	21141
19.6	Array: length (page 292)	21141
19.7	Array: empty array (page 292)	21141
19.8	Array: of objects (page 301)	21142
19.9	Array: partially filled array (page 310)	21142
19.10	Array: partially filled array: deleting an element (page 404)	21143
19.11	Array: array extension (page 311)	21143
19.12	Array: shallow copy (page 314)	21144
19.13	Array: array of arrays (page 329)	21144
19.14	Array: array of arrays: two-dimensional arrays (page 330)	21145
20	**Exception**	**21147**
20.1	Exception (page 340)	21147
20.2	Exception: getMessage() (page 345)	21147
20.3	Exception: there are many types of exception (page 347)	21147
20.4	Exception: creating exceptions (page 350)	21148
20.5	Exception: creating exceptions: with a cause (page 357)	21148
20.6	Exception: getCause() (page 366)	21149
20.7	Exception: inheritance hierarchy (page 434)	21149
20.8	Exception: making our own exception classes (page 435)	21151

DRAFT April 15, 2011

21 Inheritance **21152**

21.1 Inheritance (page 373) . . . 21152
21.2 Inheritance: a subclass extends its superclass (page 378) . . . 21153
21.3 Inheritance: invoking the superclass constructor (page 379) . . . 21154
21.4 Inheritance: invoking the superclass constructor: implicitly (page 423) 21154
21.5 Inheritance: overriding a method (page 380) . . . 21155
21.6 Inheritance: overriding a method: @Override annotation (page 430) . 21155
21.7 Inheritance: abstract class (page 385) . . . 21156
21.8 Inheritance: abstract method (page 386) . . . 21156
21.9 Inheritance: polymorphism (page 390) . . . 21157
21.10 Inheritance: polymorphism: dynamic method binding (page 391) . . 21158
21.11 Inheritance: final methods and classes (page 391) . . . 21158
21.12 Inheritance: adding more object state (page 393) . . . 21159
21.13 Inheritance: adding more instance methods (page 395) . . . 21159
21.14 Inheritance: testing for an instance of a class (page 397) . . . 21159
21.15 Inheritance: casting to a subclass (page 397) . . . 21159
21.16 Inheritance: is a versus has a (page 406) . . . 21160
21.17 Inheritance: using an overridden method (page 414) . . . 21160
21.18 Inheritance: constructor chaining (page 423) . . . 21161
21.19 Inheritance: multiple inheritance (page 509) . . . 21162

22 File IO API **21163**

22.1 File IO API: IOException (page 450) . . . 21163
22.2 File IO API: InputStream (page 451) . . . 21164
22.3 File IO API: InputStreamReader (page 456) . . . 21165
22.4 File IO API: BufferedReader (page 459) . . . 21165
22.5 File IO API: FileInputStream (page 462) . . . 21165
22.6 File IO API: FileReader (page 462) . . . 21165
22.7 File IO API: OutputStream (page 462) . . . 21166
22.8 File IO API: OutputStreamWriter (page 462) . . . 21166
22.9 File IO API: FileOutputStream (page 463) . . . 21167
22.10 File IO API: FileWriter (page 463) . . . 21167
22.11 File IO API: PrintWriter (page 463) . . . 21168
22.12 File IO API: PrintWriter: checkError() (page 464) . . . 21168
22.13 File IO API: PrintWriter: versus PrintStream (page 468) . . . 21169
22.14 File IO API: PrintWriter: can also wrap an OutputStream (page 468) 21169
22.15 File IO API: File (page 469) . . . 21170
22.16 File IO API: DataOutputStream (page 479) . . . 21170
22.17 File IO API: DataInputStream (page 479) . . . 21171

23 Collections API **21171**

23.1 Collections API (page 538) . . . 21171
23.2 Collections API: Lists (page 538) . . . 21171
23.3 Collections API: Lists: List interface (page 538) . . . 21172
23.4 Collections API: Lists: List interface: iterator() (page 553) . . . 21172
23.5 Collections API: Lists: List interface: extends Collection (page 556) 21173
23.6 Collections API: Lists: ArrayList (page 539) . . . 21173

23.7 Collections API: Lists: add(index) and remove(index) (page 557) . . 21173
23.8 Collections API: Lists: LinkedList (page 558) 21174
23.9 Collections API: Collections class (page 543) 21174
23.10 Collections API: Sets (page 546) 21175
23.11 Collections API: Sets: Set interface (page 546) 21175
23.12 Collections API: Sets: Set interface: iterator() (page 554) 21176
23.13 Collections API: Sets: Set interface: extends Collection (page 557) . 21176
23.14 Collections API: Sets: HashSet (page 548) 21176
23.15 Collections API: Sets: TreeSet (page 552) 21176
23.16 Collections API: Sets: TreeSet: iterator() (page 554) 21177
23.17 Collections API: Iterator interface (page 553) 21177
23.18 Collections API: Collection interface (page 556) 21178
23.19 Collections API: Collection interface: constructor taking a Collection (page 568)21179
23.20 Collections API: Maps (page 559) 21180
23.21 Collections API: Maps: Map interface (page 560) 21180
23.22 Collections API: Maps: TreeMap (page 560) 21181
23.23 Collections API: Maps: HashMap (page 567) 21181

1 Computer basics

1.1 Computer basics: hardware (page 3)

The physical parts of a computer are known as **hardware**. You can see them, and touch them.

1.2 Computer basics: hardware: processor (page 3)

The **central processing unit** (**CPU**) is the part of the **hardware** that actually obeys instructions. It does this dumbly – computers are not inherently intelligent.

1.3 Computer basics: hardware: memory (page 3)

The **computer memory** is part of the computer which is capable of storing and retrieving **data** for short term use. This includes the **machine code** instructions that the **central processing unit** is obeying, and any other data that the computer is currently working with. For example, it is likely that an image from a digital camera is stored in the computer memory while you are editing or displaying it, as are the machine code instructions for the image editing program.

The computer memory requires electrical power in order to remember its data – it is **volatile memory** and will forget its contents when the power is turned off.

DRAFT April 15, 2011

An important feature of computer memory is that its contents can be accessed and changed in any order required. This is known as **random access** and such memory is called **random access memory** or just **RAM**.

1.4 Computer basics: hardware: persistent storage (page 3)

For longer term storage of **data**, computers use **persistent storage** devices such as **hard disc**s and **DVD ROM**s. These are capable of holding much more information than **computer memory**, and are persistent in that they do not need power to remember the information stored on them. However, the time taken to store and retrieve data is *much* longer than for computer memory. Also, these devices cannot as easily be accessed in a random order.

1.5 Computer basics: hardware: input and output devices (page 3)

Some parts of the **hardware** are dedicated to receiving input from or producing output to the outside world. Keyboards and mice are examples of **input device**s. Displays and printers are examples of **output device**s.

1.6 Computer basics: software (page 3)

One part of a computer you cannot see is its **software**. This is stored on **computer media**, such as **DVD ROM**s, and ultimately inside the computer, as lots of numbers. It is the instructions that the computer will obey. The closest you get to seeing it might be if you look at the silver surface of a DVD ROM with a powerful magnifying glass!

1.7 Computer basics: software: machine code (page 3)

The instructions that the **central processing unit** obeys are expressed in a language known as **machine code**. This is a very **low level language**, meaning that each instruction gets the computer to do only a very simple thing, such as the **addition** of two numbers, or sending a **byte** to a printer.

1.8 Computer basics: software: operating system (page 4)

A collection of **software** which is dedicated to making the computer generally usable, rather than being able to solve a *particular* task, is known as an **operating system**. The most popular examples for modern personal computers are Microsoft Windows, Mac OS X and Linux. The

DRAFT April 15, 2011

latter two are implementations of Unix, which was first conceived in the early 1970s. The fact it is still in widespread use today, especially by computer professionals, is proof that it is a thoroughly stable and well **design**ed and integrated platform for the expert (or budding expert) computer scientist.

1.9 Computer basics: software: application program (page 4)

A piece of **software** which is dedicated to solving a particular task, or application, is known as an **application program**. For example, an image editing program.

1.10 Computer basics: data (page 3)

Another part of the computer that you cannot see is its **data**. Like **software** it is stored as lots of numbers. Computers are processing and producing data all the time. For example, an image from a digital camera is data. You can only see the picture when you display it using some image displaying or editing software, but even this isn't showing you the actual data that makes up the picture. The names and addresses of your friends is another example of data.

1.11 Computer basics: data: files (page 5)

When **data** is stored in **persistent storage**, such as on a **hard disc**, it is organized into chunks of related information known as **file**s. Files have names and can be accessed by the computer through the **operating system**. For example, the image from a digital camera would probably be stored in a jpeg file, which is a particular type of image file, and the name of this file would probably end in `.jpg` or `.jpeg`.

1.12 Computer basics: data: files: text files (page 5)

A **text file** is a type of **file** that contains **data** stored directly as **character**s in a human readable form. This means if you were to send the raw contents directly to the printer, you would (for most printers) be immediately able to read it. Examples of text files include `README.txt` that sometimes comes with **software** you are installing, or source text for a document to be processed by the LaTeX[6] document processing system, such as the ones used to produce this book (prior to publication). As you will see shortly, a more interesting example for you, is computer program **source code** files.

DRAFT April 15, 2011

1.13 Computer basics: data: files: binary files (page 5)

A **binary file** is another kind of **file** in which **data** is stored as **binary** (base 2) numbers, and so is not human readable. For example, the image from a digital camera is probably stored as a jpeg file, and if you were to look directly at its contents, rather than use some **application program** to display it, you would see what appears to be nonsense! An interesting example of a binary file is the **machine code** instructions of a program.

2 Java tools

2.1 Java tools: text editor (page 5)

A **text editor** is a program that allows the user to type and edit **text file**s. You may well have used `notepad` under Microsoft Windows; that is a text editor. More likely you have used `Microsoft Word`. If you have, you should note that it is not a text editor, it is a **word processor**. Although you can save your documents as text files, it is more common to save them as `.doc` **file**s, which is actually a **binary file** format. `Microsoft Word` is not a good tool to use for creating program **source code** files.

If you are using an **integrated development environment** to support your programming, then the text editor will be built in to it. If not, there are a plethora of text editors available which are suited to Java programming.

2.2 Java tools: javac compiler (page 9)

The Java **compiler** is called `javac`. Java program source is saved by the programmer in a **text file** that has the suffix `.java`. For example, the text file `HelloWorld.java` might contain the source text of a program that prints `Hello world!` on the **standard output**. This text file can then be **compile**d by the Java compiler, by giving its name as a **command line argument**. Thus the command

```
javac HelloWorld.java
```

will produce the **byte code** version of it in the **file** `HelloWorld.class`. Like **machine code** files, byte code is stored in **binary file**s as numbers, and so is not human readable.

DRAFT April 15, 2011

2.3 Java tools: java interpreter (page 9)

When the end user wants to run a Java program, he or she invokes the `java` **interpreter** with the name of the program as its **command line argument**. The program must, of course, have been **compile**d first! For example, to run the `HelloWorld` program we would issue the following command.

```
java HelloWorld
```

This makes the **central processing unit** run the interpreter or **virtual machine** `java`, which itself then **execute**s the program named as its first argument. Notice that the suffix `.java` is needed when compiling the program, but no suffix is used when **run**ning it. In our example here, the virtual machine finds the **byte code** for the program in the **file** `HelloWorld.class` which must have been previously produced by the **compiler**.

2.4 Java tools: javadoc (page 223)

A **class** which is intended to be reusable in many programs should have user documentation to enable another programmer to use it without having to look at the implementation code. In Java this is achieved by the implementer of the class writing **doc comment**s in the code, and then processing them with the `javadoc` program. This tool produces a web page which describes the class from the information in the doc comments and from the structure of the class itself, and this page is linked to the pages for other classes as appropriate. For example, the heading of each **public method** is documented on the web page, with the description of the method being taken by `javadoc` from the doc comment which the implementer supplied for the method.

The resulting user documentation produced by `javadoc` can be placed anywhere we wish – on a web server for example. Meanwhile the *source* of that documentation is kept with the **source code** for the class, indeed it is inside the same **file**. This excellent idea makes it easy for the programmer to maintain information on how to use the class as he or she alters the code, but without restricting where the final documentation can be put.

A doc comment starts with the symbol `/**` and ends with `*/`. These are written in certain places as follows.

- A comment before the start of the class (after any **import statement**s) describing its purpose.
- A comment before each public **variable** describing the meaning of that variable.
- A comment before each public method describing what it does, its **method parameter**s and **return** value.

DRAFT, April 15, 2011

- Optionally, a comment before each **private** variable and method. This is less useful than documentation for public items as normal users of the class do not have access to the private ones. So, many programmers do not write doc comments for these (although of course they do write ordinary **comment**s!). On the other hand, some take the view that anybody who needs to *maintain* the class is, in effect, a user of both the public *and* private parts, and so user documentation of the whole class is of benefit.

The implementer writes user documentation text as appropriate inside the doc comments. The emphasis is on how to use the features, not on how they are implemented. He or she also includes various **doc comment tag**s to help the `javadoc` program process the text properly. Here are some of the most commonly used tags.

Tag	**Meaning**	**Where used**
`@author` author name(s)	State the author of the code.	Before the class starts.
`@param` parameter description	Describe a method parameter.	Before a method.
`@return` description	Describe a method result.	Before a method.

Most doc comments use more than one line, and it is conventional (but not essential) to start continuation lines with an asterisk (`*`) neatly lined up with the first asterisk in the opening comment symbol. The first sentence should be a summary of the whole thing being documented – these are copied to a summary area of the final documentation.

For a doc comment tag to be recognized by `javadoc`, it must be the first word on a line of the comment, preceded only by **white space**, or an asterisk.

Doc comments are sometimes (but wrongly) called **javadoc comment**s.

2.5 Java tools: javadoc: throws tag (page 355)

There is another **doc comment tag** which is used to describe the **exception**s that a **method throws**.

Tag	**Meaning**	**Where used**
`@throws` exception name and description	Describes the circumstances leading to an exception.	Before a method.

DRAFT April 15, 2011

3 Operating environment

3.1 Operating environment: programs are commands (page 7)

When a program is **execute**d, the name of it is passed to the **operating system** which finds and loads the **file** of that name, and then starts the program. This might be hidden from you if you are used to starting programs from a menu or browser interface, but it happens nevertheless.

3.2 Operating environment: standard output (page 7)

When programs **execute**, they have something called the **standard output** in which they can produce text results. If they are **run** from some kind of **command line interface**, such as a Unix **shell** or a Microsoft Windows **Command Prompt**, then this output appears in that interface while the program is running. (If they are invoked through some **integrated development environment**, browser, or menu, then this output might get displayed in some pop-up box, or special console window.)

3.3 Operating environment: command line arguments (page 8)

Programs can be, and often are, given **command line argument**s to vary their behaviour.

3.4 Operating environment: standard input (page 187)

In addition to **standard output**, when programs **execute** they also have a **standard input** which allows text **data** to be entered into the program as it runs. If they are **run** from some kind of **command line interface**, such as a Unix **shell** or a Microsoft Windows **Command Prompt**, then this input is typically typed on the keyboard by the end user.

3.5 Operating environment: standard error (page 344)

When programs **execute**, in addition to **standard output** and **standard input**, they also have another facility called **standard error**. This is intended to be used for output about errors and exceptional circumstances, rather than program results. In some **operating environment**s there might be no difference between these two in practice, but their separation at the program level enables them to be handled differently where that is permitted. For example, on Unix systems, the end user can redirect the standard output into a **file**, whilst leaving the standard error to appear on the screen, or vice versa, etc. as desired. Nowadays, this is also true of Microsoft Windows.

DRAFT April 15, 2011

4 Class

4.1 Class: programs are divided into classes (page 16)

In Java, the source text for a program is separated into pieces called **class**es. The source text for each class is (usually) stored in a separate **file**. Classes have a name, and if the name is `HelloWorld` then the text for the class is saved by the programmer in the **text file** `HelloWorld.java`.

One reason for dividing programs into pieces is to make them easier to manage – programs to perform complex tasks typically contain thousands of lines of text. Another reason is to make it easier to share the pieces between more than one program – such **software reuse** is beneficial to programmer productivity.

Every program has at least one class. The name of this class shall reflect the intention of the program. By convention, class names start with an upper case letter.

4.2 Class: public class (page 16)

A **class** can be declared as being **public**, which means it can be accessed from anywhere in the running Java environment; in particular the **virtual machine** itself can access it. The source text for a public class definition starts with the **reserved word** **`public`**. A reserved word is one which is part of the Java language, rather than a word chosen by the programmer for use as, say, the name of a program.

4.3 Class: definition (page 16)

After stating whether it has **public** access, a **class** next has the **reserved word** **`class`**, then its name, then a left brace ({), its body of text and finally a closing right brace (}).

```
public class MyFabulousProgram
{
  ... Lots of stuff here.
}
```

4.4 Class: objects: contain a group of variables (page 158)

We can group a collection of **variable**s into one entity by creating an **object**. For example, we might wish to represent a point in two dimensional space using an *x* and a *y* value to make up

DRAFT April 15, 2011

a coordinate. We would probably wish to combine our `x` and `y` variables into a single object, a `Point`.

4.5 Class: objects: are instances of a class (page 158)

Before we can make **object**s, we need to tell Java how the objects are to be **construct**ed. For example, to make a `Point` object, we would need to tell Java that there are to be a pair of **variable**s inside it, called `x` and `y`, and tell it what **type**s these variables have, and how they get their values. We achieve this by writing a **class** which will act as a template for the creation of objects. We need to write such a template class for each kind of object we wish to have. For example, we would write a `Point` class describing how to make `Point` objects. If, on the other hand, we wanted to group together a load of variables describing attributes of wardrobes, so we could make objects each of which represents a single wardrobe, then we would probably call that class `Wardrobe`. Java lets us choose any name that we feel is appropriate, except **reserved word**s (although by convention we always start the name with a capital letter).

Once we have described the template, we can get Java to make objects of that class at **run time**. We say that these objects are **instance**s of the class. So, for example, particular `Point` objects would all be instances of the `Point` class. We can create as many different `Point` objects as we wish, each containing its own `x` and `y` variables, all from the one template, the `Point` class.

4.6 Class: objects: this reference (page 180)

Sometimes, in **constructor method**s or in **instance method**s of a **class** we wish to refer to the **object** that the constructor is creating, or to which the instance method belongs. For this purpose, whenever the **reserved word** **`this`** is used in or as an **expression** it means a **reference** to the object that is being created by the constructor or that owns the instance method, etc.. We can only use the **this reference** in places where it makes sense, such as constructor methods, instance methods and **instance variable** initializations. So, **`this`** (when used in this way) behaves somewhat like an extra instance variable in each object, automatically set up to contain a reference to that object.

For example, in a `Point` class we may wish to have an instance method that yields a point which is half way between the origin and **`this`** point.

```
public Point halfThisPoint()
{
  return halfWayPoint(new Point(0, 0));
} // halfThisPoint
```

An alternative implementation would be as follows.

DRAFT April 15, 2011

```
public Point halfThisPoint()
{
  return new Point(0, 0).halfWayPoint(this);
} // halfThisPoint
```

4.7 Class: objects: may be mutable or immutable (page 193)

Sometimes when we **design** a **class** we desire that the **instance**s of it are **immutable object**s. This means that once such an **object** has been **construct**ed, its **object state** cannot be changed. That is, there is no way for the values of the **instance variable**s to be altered after the object is constructed.

By contrast, objects which can be altered are known as **mutable object**s.

4.8 Class: objects: compareTo() (page 222)

It is quite common to require the ability to compare an **object** with another from the same **class**, based on some **total order**, that is, a notion of **less than**, **greater than** and **equivalence**. A Java convention for this is to have an **instance method** called `compareTo` which takes a (**reference** to) another object as its **method parameter**, and **return**s an `int`. A result of `0` indicates the two objects are **equivalent**, a negative value indicates this object is less than the other, and a positive value indicates this object is greater than the other.

```
Date husbandsBirthday = ...
Date wifesBirthday = ...

if (husbandsBirthday.compareTo(wifesBirthday) > 0)
  System.out.println("The husband is older than the wife");
else if (husbandsBirthday.compareTo(wifesBirthday) == 0)
  System.out.println("The husband is the same age as the wife");
else
  System.out.println("The husband is younger than the wife");
```

4.9 Class: is a type (page 161)

A **type** is essentially a **set** of values. The `int` type is all the whole numbers that can be represented using 32 **binary digit**s, the `double` type is all the **real** numbers that can be represented using the **double precision** technique and the `boolean` type contains the values `true` and `false`. A **class** can be used as a template for creating **object**s, and so is regarded in Java as a type: the set of all objects that can be created which are **instance**s of that class. For example, a `Point` class is a type which is the set of all `Point` objects that can be created.

DRAFT April 15, 2017

4.10 Class: is a type: and has three components (page 512)

Whilst a **type** is essentially a **set** of values, it also has two other components. These are the operations which can be performed on those values, and the **operation interface** to those operations. For example, the type `int` is a collection of numbers, with operations such as **addition** and **multiplication**, and each operation has an **operator** as its operation interface, such as `+` and `*`.

The distinction between operation and operation interface is subtle, and may even seem pedantic, but nevertheless, they are not the same thing. For example, one could imagine the **design**ers of Java one day permitting a proper multiplication symbol ($\times$) to be used as an alternative to the `*` operator, without altering the meaning of the multiplication operation.

Each **class** is a type, the set of all (**reference**s to) **object**s that can be created which are **instance**s of that class. It has operations, which are the **method implementation**s of the **instance method**s of the class, and each of these operations has an operation interface, which is the **method interface**.

4.11 Class: making instances with new (page 162)

An **instance** of a **class** is created by calling the **constructor method** of the class, using the **reserved word `new`**, and supplying **method argument**s for the **method parameter**s. At **run time** when this code is **execute**d, the Java **virtual machine**, with the help of the constructor method code, creates an **object** which is an instance of the class. Although it is not stated in its heading, a constructor method always **return**s a value, which is a **reference** to the **new**ly created object. This reference can then be stored in a **variable**, if we wish. For example, if we have a `Point` class, then we might have the following code.

```
Point topLeft     = new Point(-20, 40);
Point bottomLeft  = new Point(-20, -40);
Point topRight    = new Point(20, 40);
Point bottomRight = new Point(20, -40);
```

This declares four variables, of **type** `Point` and creates four instances of the class `Point` representing the four corners of a rectangle. The four variables each contain a reference to one of the points. This is illustrated in the following diagram.

DRAFT April 15, 2011

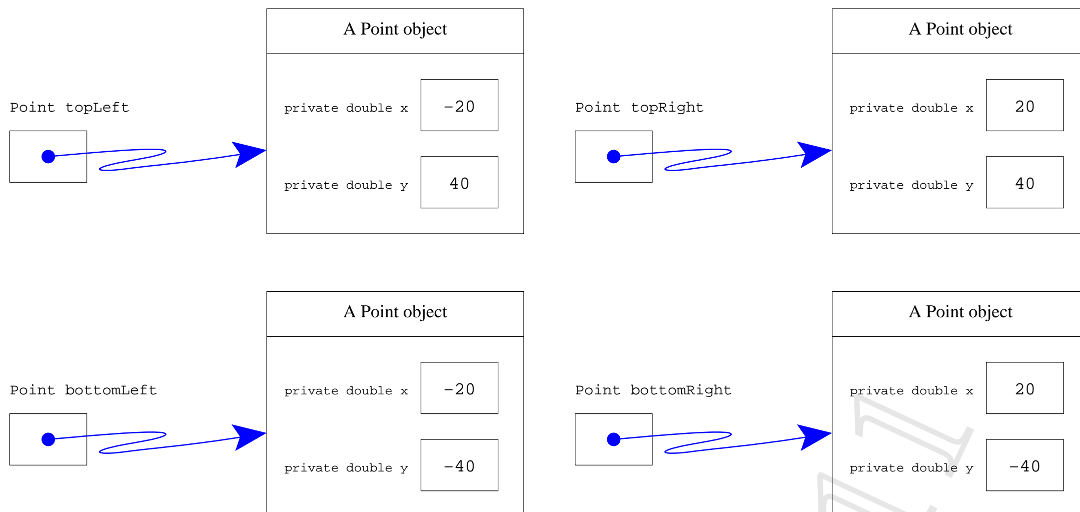

All four `Point` objects each have two **instance variables**, called `x` and `y`.

4.12 Class: accessing instance variables (page 164)

The **instance variables** of an **object** can be accessed by taking a **reference** to the object and appending a dot (`.`) and then the name of the **variable**. For example, if the variable `p1` contains a reference to a `Point` object, and `Point` objects have an instance variable called `x`, then the code `p1.x` is the instance variable `x`, belonging to the `Point` referred to by `p1`.

4.13 Class: importing classes (page 188)

At the start of the source **file** for a Java **class** we can write one or more **import statements**. These start with the **reserved word** **`import`** and then give the **fully qualified name** of a class that lives in some **package** somewhere, followed by a semi-colon(`;`). An **import** for a class permits us to talk about it from then on, by using only its class name, rather than having to always write its fully qualified name. For example, importing `java.util.Scanner` would mean that every time we refer to `Scanner` the Java **compiler** knows we really mean `java.util.Scanner`.

```
import java.util.Scanner;
  ...
  Scanner inputScanner = new Scanner(System.in);
```

If we wish, we can import all the classes in a package using a `*` instead of a class name.

```
import java.util.*;
```

DRAFT April 15, 2011

Many programmers consider this to be lazy, and it is better to import exactly what is needed, if only to help show precisely what is used by the class. There is also the issue of ambiguity: if two different packages have classes with the same name, but this class only needs one of them, then the lazy approach would cause an unnecessary problem.

However, every Java program has an automatic import for every class in the standard **package** `java.lang`, because these classes are used so regularly. That is why we can refer to `java.lang.System` and `java.lang.Integer`, etc. as just `System` and `Integer`, etc.. In other words, every class always implicitly includes the following import statement for convenience.

```
import java.lang.*;
```

4.14 Class: stub (page 191)

During development of a program with several **class**es, we often produce a **stub** for the classes we have not yet implemented. This just contains some or all of the **public** items of the class, with empty, or almost empty, bodies for the **method**s. In other words, it is the bare minimum needed to allow the classes we have so far developed to be **compile**d.

Any **non-void method**s are written with a single **return statement** to yield some temporary value of the right **type**.

These stubs are then developed into the full class code at some later stage.

4.15 Class: extending another class (page 245)

A **class** may be declared to say that it **extend**s another class, using the **reserved word `extends`**. For example, the following says that the class `HelloWorld` extends the class `javax.swing.JFrame`.

```
import javax.swing.JFrame;
public class HelloWorld extends JFrame
```

This means that all **instance**s of `HelloWorld` have the properties that any instance of `JFrame` would have, but also have all the properties that we additionally define in the `HelloWorld` class. It is a way of adding properties to a class without actually changing the class – the new class is an **extension** of the other one.

DRAFT April 15, 2011

4.16 Class: generic class (page 491)

A **generic class** is a **class** which has one or more **type parameter**s written within angled brackets (<>) just after its name in the class heading. When an **instance** of a generic class is made, specific **type**s are supplied as **type argument**s for the type parameters, in a similar way that **method argument**s are supplied for **method parameter**s in a **method call**.

In the following symbolic example, `T1` and `T2` are type parameters.

```
public class MyGenericClass<T1, T2>
{
  ... Typical class stuff here,
  ... but using T1 and T2 as though they are types
  ... (in permitted ways).
  T1 someVariable = ...
  T2 someOtherVariable = ...
  ...
} // class MyGenericClass
```

When we make an instance of `MyGenericClass`, we can supply a specific type for each type parameter, as in the following example.

```
MyGenericClass<String, Date> myVariable = new MyGenericClass<String, Date>();
```

A class is a **type**. However, the intention with a generic class is that we supply specific type arguments for the type parameters before we use it, and in doing so, we identify a **parameterized type**. For example, from the generic class `MyGenericClass` we can have parameterized types such as `MyGenericClass<String, Date>`,
`MyGenericClass<Integer, String>`, etc., including ones involving **array**s, like
`MyGenericClass<String[], Integer>`, and so on.

A parameterized type almost behaves as though we have made a textual copy of the generic class, and replaced each type parameter with its corresponding type argument. But not quite. Instead, due to the way Java actually implements generic classes, there are some restrictions. In particular, type arguments must be **reference type**s, such as classes and arrays. This means they cannot be **primitive type**s.

4.17 Class: generic class: bound type parameter (page 496)

The **type parameter**s of a **generic class** may be **bound type parameter**s, which means we specify certain restrictions for the **type argument**s that can be supplied when a **parameterized type** is identified.

DRAFT April 15, 2011

4.18 Class: generic class: bound type parameter: extends some class (page 496)

One kind of restriction we can specify for a **bound type parameter** is that the type argument must **extend** some known **class**. This is done by following the name of the **type parameter** with the **reserved word `extends`** and then the known class. When a type argument is supplied, the **compiler** checks that it is either the known class, or a **subclass** of it.

For example, in the context of some vehicle simulation program, the following is a class that has a type parameter, `VehicleType`, for which any corresponding **type argument** must be `Vehicle` or a subclass of it.

```
public class ServiceCentre<VehicleType extends Vehicle>
{
  ... Etc., using VehicleType as a type (in permitted ways)
  ... but knowing that it is a Vehicle
  ... and so using some Vehicle methods, etc..

  public void service(VehicleType vehicle)
  {
    if (! vehicle.isRoadworthy())
    {
      ...
    } // if
  } // service

  ...
} // class ServiceCentre
```

This would allow us to make `ServiceCentre` **object**s for particular kinds of `Vehicle`.

```
ServiceCentre<Car> garage = new ServiceCentre<Car>();
Car car = new Car(...);
Lorry lorry = new Lorry(...);
garage.service(car);
garage.service(lorry);
```

The last line above would cause a **compile time error**.

4.19 Class: generic class: bound type parameter: extends some interface (page 526)

The **type parameters** of a **generic class** (or **generic interface**, or **generic method**) may be declared to **extend** some known **type**. The known type may be a **class** or an **interface**. Perhaps

DRAFT April 15, 2011

surprisingly, we use the **reserved word `extends`** even if the known type is an interface. This is in recognition of the idea that an **interface** is a **type** in just the same way that a class is. One type can be an **extension** of another through **inheritance**, either by being a **subclass** of another class, a **subinterface** of another interface, or by being a class that **implement**s an interface.

If the known type is an interface, then when a **type argument** is supplied for the type parameter, the **compiler** checks that it is a class which implements that interface, or is that interface or an interface that extends it.

4.20 Class: generic class: where type parameters cannot be used (page 501)

Each **type parameter** of a **generic class** may be treated as a **type** within the generic class, except for certain restrictions, which fall into two categories.

The first is about the meaning of type parameters. A **type argument** is supplied for each of these to identify a **parameterized type**, which is then ready for **instance**s of it to be made. The type arguments only mean anything in the context of creating instances, and make no sense in the **static context** of the generic class (which is not part of the type). So, we cannot refer to the type parameters in `static` parts, that is, in **class variable** and **class method** declarations.[1]

The second set of restrictions are associated with the way Java implements generic classes. In particular, we cannot create any **instance**s of a type parameter, nor create any **array**s whose **array element**s are of that type. (Essentially, the generic features of a **class** is an entirely **compile time** artifact – to enable the **compiler** to undertake more type checking than it otherwise could. At **run time**, the **virtual machine** has no knowledge of the type parameters, and so cannot *create* instances of the correct type.)

4.21 Class: generic class: used as a raw type (page 502)

A **generic class** is still a **class** and hence a **type**, and actually it can be used directly to make **instance**s of it without supplying **type argument**s. This is due to legacy issues: generic classes were added in Java 5.0, and **type parameter**s were added to many standard **application program interface** (**API**) classes at that time. Obviously there already existed millions of Java programs that use those classes, and it would be unacceptable for them all to suddenly stop working!

Java refers to the type of the generic class without type parameters as the **raw type** for the class. If we use the raw type, then the **compiler** assumes the best known actual type for each of its type parameters, and gives us warnings, about types being `unchecked`. But it goes ahead and makes the **byte code** anyway. This way, programmers are encouraged to use the generic

[1]There is actually a separate mechanism for putting type parameters on class methods.

DRAFT April 15, 2017

classes properly for new code and gradually change legacy code to do so. The best known type assumed by the compiler for a type parameter which **extend**s some concrete type is that concrete type, and for ones that do not it is `java.lang.Object`.

5 Method

5.1 Method (page 118)

A **method** in Java is a section of code, dedicated to performing a particular task. All programs have a **main method** which is the starting point of the program. We can have other methods too, and we can give them any name we like – although we should always choose a name which suits the purpose. By convention, method names start with a lower case letter. For example, `System.out.println()` is a method which prints a line of text. Apart from its slightly strange spelling, the name `println` does reflect the meaning of the method.

5.2 Method: main method: programs contain a main method (page 17)

All Java programs contain a section of code called `main`, and this is where the computer will start to **execute** the program. Such sections of code are called **method**s because they contain instructions on how to do something. The **main method** always starts with the following heading.

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

5.3 Method: main method: is public (page 17)

The **main method** starts with the **reserved word** **`public`**, which means it can be accessed from anywhere in the running Java environment. This is necessary – the program could not be **run** by the **virtual machine** if the starting point was not accessible to it.

```
public
```

5.4 Method: main method: is static (page 17)

The **main method** of the program has the **reserved word** **`static`** which means it is allowed to be used in the **static context**. A context relates to the use of **computer memory** during

DRAFT April 15, 2011

the **run**ning of the program. When the **virtual machine** loads a program, it creates the static context for it, allocating computer memory to store the program and its **data**, etc.. A **dynamic context** is a certain kind of allocation of memory which is made later, during the running of the program. The program would not be able to start if the main method was not allowed to run in the static context.

```
public static
```

5.5 Method: main method: is void (page 17)

In general, a **method** (section of code) might calculate some kind of **function** or formula, and **return** the answer as a result. For example, the result might be a number. If a method returns a result then this must be stated in its heading. If it does not, then we write the **reserved word** **`void`**, which literally means (among other definitions) 'without contents'. The **main method** does not return a value.

```
public static void
```

5.6 Method: main method: is the program starting point (page 17)

The starting part, or **main method**, of the program is always called `main`, because it is the main part of the program.

```
public static void main
```

5.7 Method: main method: always has the same heading (page 18)

The **main method** of a Java program must always have a heading like this.

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

This is true even if we do not intend to use any **command line argument**s. So a typical single **class** program might look like the following.

```
public class MyFabulousProgram
{
  public static void main(String[] args)
```

DRAFT April 15, 2011

```
  {
    ... Stuff here to perform the task.
  }
}
```

5.8 Method: private (page 118)

A **method** should be declared with a **private** visibility **modifier** if it is not intended to be usable from outside the **class** it is defined in. This is done by writing the **reserved word** `private` instead of `public` in the heading.

5.9 Method: accepting parameters (page 118)

A **method** may be given **method parameter**s which enable it to vary its effect based on their values. This is similar to a program being given **command line argument**s, indeed the arguments given to a program are passed as parameters to the **main method**.

Parameters are declared in the heading of the method. For example, main methods have the following heading.

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

The text inside the brackets is the declaration of the parameters. A method can have any number of parameters, including zero. If there is more than one, they are separated by commas (`,`). Each parameter consists of a **type** and a name. For example, the following method is given two parameters, a `double` and an `int`.

```
private static void printHeightPerYear(double height, int age)
{
  System.out.println("At age " + age + ", height per year ratio is "
                     + height / age);
} // printHeightPerYear
```

You should think of parameters as being like **variable**s defined inside the method, except that they are given initial values before the method body is **execute**d. For example, the single parameter to the main method is a variable which is given a **list** of strings before the method begins execution, these strings being the command line arguments supplied to the program.

The names of the parameters are not important to Java – as long as they all have different names! The names only mean something to the human reader, which is of course important. The above method could easily have been written as follows.

DRAFT April 15, 2011

```
private static void printHeightPerYear(double howTall, int howOld)
{
  System.out.println("At age " + howOld + ", height per year ratio is "
                     + howTall / howOld);
} // printHeightPerYear
```

You might think the first version is subjectively nicer than the second, but clearly both are better than this next one!

```
private static void printHeightPerYear(double d, int i)
{
  System.out.println("At age " + i + ", height per year ratio is "
                     + d / i);
} // printHeightPerYear
```

And that is only marginally better than calling the parameters, say `x` and `y`. However, Java does not care – it is not clever enough to be able to, as it can have no understanding of the problem being solved by the code.

5.10 Method: accepting parameters: of a class type (page 164)

The **method parameter**s of a **method** can be of any **type**, including **class**es. A parameter which is of a class type must be given a **method argument** value of that type when the method is invoked, for example a **reference** to an **object** which is an **instance** of the class named as the parameter type.

5.11 Method: accepting parameters: of an array type (page 297)

The **method parameter**s of a **method** can be of any **type**, including **array**s. A parameter which is of an **array type** must be given a **method argument** value of that type when the method is invoked. This value will of course be a **reference** to an array which has **array element**s of the **array base type**, or the **null reference**.

The most obvious example of this is the `String[]` **command line argument** array, which is passed to the **main method** by the Java **virtual machine**.

5.12 Method: calling a method (page 119)

The body of a **method** is **execute**d when some other code refers to it using a **method call**. For example, the program calls a method named `println` when it executes `System.out.println("Hello`

DRAFT April 15, 2011

world!"). For another example, if we have a method, named printHeightPerYear, which prints out a height to age ratio when it is given a height (in metres) and an age, then we could make it print the ratio between the height 1.6 and the age 14 using the following method call.

```
printHeightPerYear(1.6, 14);
```

When we call a method we supply a **method argument** for each **method parameter**, separating them by commas (,). These argument values are copied into the corresponding parameters of the method – the first argument goes into the first parameter, the second into the second, and so on.

The arguments passed to a method may be the current values of **variable**s. For example, the above code could have been written as follows.

```
double personHeight = 1.6;
int personAge = 14;

printHeightPerYear(personHeight, personAge);
```

As you may expect, the arguments to a method are actually **expression**s rather than just **literal value**s or variables. These expressions are **evaluate**d at the time the method is called. So we might have the following.

```
double growthLastYear = 0.02;

printHeightPerYear(personHeight - growthLastYear, personAge - 1);
```

5.13 Method: void methods (page 120)

Often, a **method** might calculate some kind of **function** or formula, perhaps based on its **method parameter**s, and **return** the answer as a result. The result might be an **int** or a **double** or some other **type**. If a method returns a result then the **return type** of the result must be stated in its heading. If it does not, then we write the word **void** instead, which literally means (among other definitions) 'without contents'. For example, the **main method** of a program does not return a result – it is always a **void method**.

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

DRAFT April 15, 2011

5.14 Method: returning a value (page 122)

A **method** may **return** a result back to the code that called it. If this is so, we declare the **return type** of the result in the method heading, in place of the **reserved word** **`void`**. Such methods are often called **non-void methods**. For example, the following method takes a Celsius temperature, and returns the corresponding Fahrenheit value.

```
private static double celsiusToFahrenheit(double celsiusValue)
{
  double fahrenheitValue = celsiusValue * 9 / 5 + 32;
  return fahrenheitValue;
} // celsiusToFahrenheit
```

The method is declared with a return type of **`double`**, by writing that **type** name before the method name.

The **return statement** is how we specify what value is to be returned as the result of the method. The **statement** causes the execution of the method to end, and control to transfer back to the code that called the method.

The result of a non-void method can be used in an **expression**. For example, the method above might be used as follows.

```
double celsiusValue = Double.parseDouble(args[0]);
System.out.println("The Fahrenheit value of "
                   + celsiusValue + " Celsius is "
                   + celsiusToFahrenheit(celsiusValue) + ".");
```

The return statement takes any expression after the reserved word **`return`**. So our method above could be implemented using just one statement.

```
private static double celsiusToFahrenheit(double celsiusValue)
{
  return celsiusValue * 9 / 5 + 32;
} // celsiusToFahrenheit
```

5.15 Method: returning a value: of a class type (page 176)

A **method** may **return** a result back to the code that called it, and this may be of any **type**, including a **class**. In such cases, the value returned will typically be a **reference** to an **object** which is an **instance** of the class named as the **return type**.

DRAFT April 15, 2011

For example, in a `Point` class with **instance variable**s `x` and `y`, we might have an **instance method** to return a `Point` which is half way along a straight line between this `Point` and a given other `Point`.

```
public Point halfWayPoint(Point other)
{
  double newX = (x + other.x) / 2;
  double newY = (y + other.y) / 2;
  return new Point(newX, newY);
} // halfWayPoint
```

The method creates a **new object** and then returns a reference to it. This might be used as follows.

```
Point p1 = new Point(3, 4);
Point p2 = new Point(45, 60);

Point halfWayBetweenP1AndP2 = p1.halfWayPoint(p2);
```

The reference to the new `Point` returned by the instance method, is stored in the **variable** `halfWayBetweenP1AndP2`. It would, of course, be the point $(24, 32)$. This is illustrated in the following diagram.

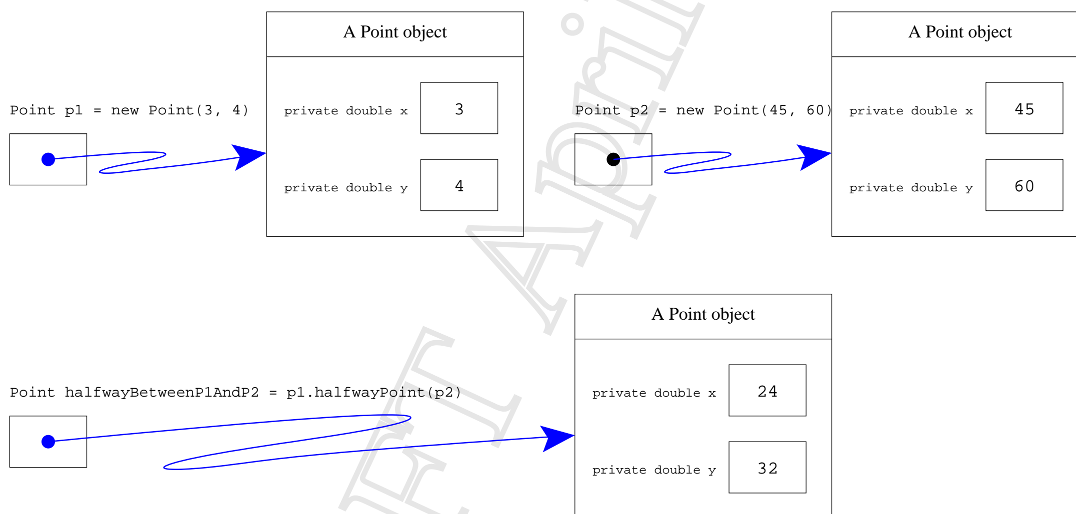

5.16 Method: returning a value: multiple returns (page 196)

The **return statement** is how we specify what value is to be **return**ed as the result of a **non-void method**. The **statement** causes the execution to end, and control to transfer back to the

DRAFT April 15, 2011

code that called the **method**. Typically, this is written as the last statement in the method, but we can actually write one or more anywhere in the method.

The Java **compiler** checks to make sure that we have been sensible, and that:

- There is no path through the method that does not end with a return statement.
- There is no code in the method that can never be reached due to an earlier occurring return statement.

5.17 Method: returning a value: of an array type (page 312)

A **method** may **return** a result back to the code that called it. This result may be of any **type**, including an **array type**. This value will of course be a **reference** to an **array** which contains **array element**s of the appropriate type as stated in the **return type** (or the **null reference**).

5.18 Method: changing parameters does not affect arguments (page 124)

We can think of **method parameter**s as being like **variable**s defined inside the **method**, but which are given their initial value by the code that calls the method. This means the method can change the values of the parameters, like it can for any other variable defined in it. Such changes have no effect on the environment of the code that called the method, regardless of where the **method argument** values came from. An argument value, be it a literal constant, taken straight from a variable, or the result of some more complex **expression**, is simply copied into the corresponding parameter at the time the method is called. This is known as **call by value**.

5.19 Method: changing parameters does not affect arguments: but referenced objects can be changed (page 208)

All **method parameter**s obtain their values from the corresponding **method argument** using the **call by value** principle. This means a **method** cannot have any effect on the calling environment via its method parameters if they are of a **primitive type**.

However, if a method parameter is of a **reference type** then there is nothing to stop the code in the method following the **reference** supplied as the argument, and altering the state of the **object** it refers to (if it is a **mutable object**). Indeed, such behaviour is often exactly what we want.

In the abstract example below, assume that `changeState()` is an **instance method** in the **class** `SomeClass` which alters the values of some of the **instance variable**s.

DRAFT April 15, 2011

```
public static void changeSomething(SomeClass object, SomeType value)
{
  object.changeState(value); // This really changes the object referred to.
  object = null;             // This has no effect outside of this method.
  ...
} // changeSomething
  ...
  SomeClass variable = new SomeClass();
  changeSomething(variable, someValueOfSomeType);
```

At the end of the above code, the change caused by the first line of the method has had an impact outside of the method, whereas the second line has had no such effect.

5.20 Method: constructor methods (page 159)

A **class** which is to be used as a template for making **object**s should be given a **constructor method**. This is a special kind of **method** which contains instructions for the **construct**ion of objects that are **instance**s of the class. A constructor method always has the same name as the class it is defined in. It is usually declared as being **public**, but we do not specify a **return type** or write the **reserved word** **`void`**. Constructor methods can have **method parameter**s, and typically these are the initial values for some or all of the **instance variable**s.

For example, the following might be a constructor method for a `Point` class, which has two instance variables, `x` and `y`.

```
public Point(double requiredX, double requiredY)
{
  x = requiredX;
  y = requiredY;
} // Point
```

This says that in order to construct an object which is an instance of the class `Point`, we need to supply two **`double`** values, the first will be placed in the `x` instance variable, and the second in the `y` instance variable. Constructor methods are called in a similar way to any other **method**, except that we precede the **method call** with the **reserved word** **`new`**. For example, the following code would create a **new** object, which is an instance of the class `Point`, and in which the instance variables `x` and `y` have the values `7.4` and `-19.9` respectively.

```
new Point(7.4, -19.9);
```

We can create as many `Point` objects as we wish, each of them having their own pair of instance variables, and so having possibly different values for `x` and `y`. These next four `Point` objects are the coordinates of a rectangle which is centred around the origin of a graph, point (0, 0).

DRAFT April 15, 2011

```
new Point(-20, 40);
new Point(-20, -40);
new Point(20, 40);
new Point(20, -40);
```

This is illustrated in the following diagram.

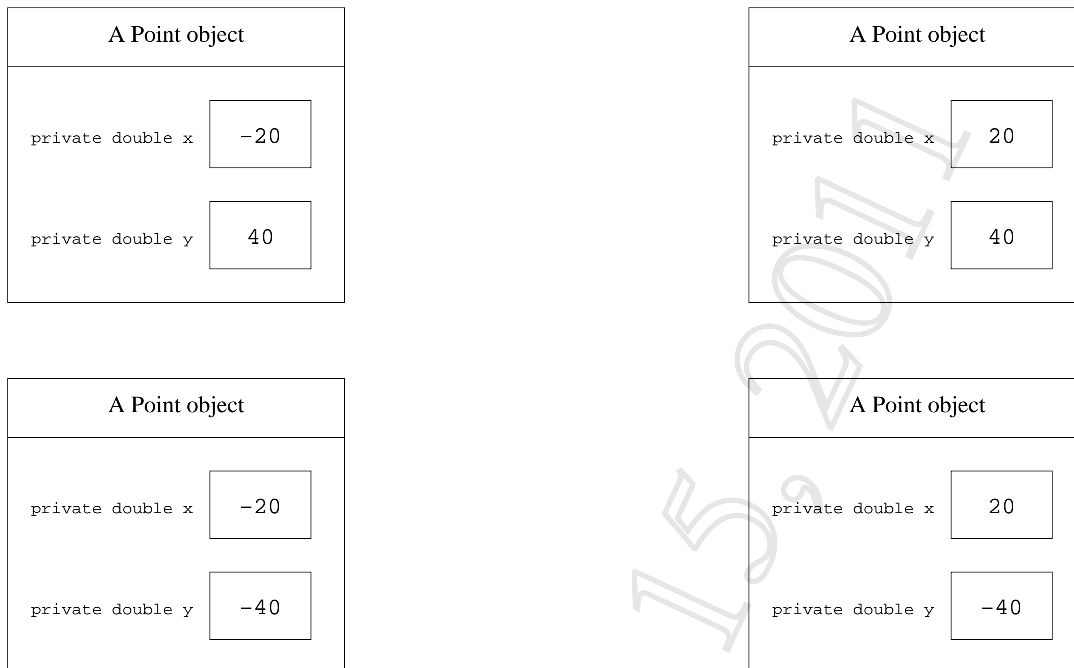

All four `Point` objects each have two instance variables, called `x` and `y`.

5.21 Method: constructor methods: more than one (page 203)

A **class** can have more than one **constructor method** as long as the number, order and/or **type**s of the **method parameter**s are different. This distinction is necessary so that the **compiler** can tell which constructor should be used when an **object** is being created.

5.22 Method: constructor methods: more than one: using this (page 393)

Typically, the **method parameter**s to **constructor method**s are values for **instance variable**s, and in **class**es where there are several instance variables it can be convenient to have multiple constructor methods, some of which assume sensible default values for some instance variables.

For example, in a `Point` class, it might quite reasonably be decided that for convenience, we can easily obtain a representation of the origin by **construct**ing a `Point` using no **method argument**s.

DRAFT April 15, 2011

```
public class Point
{
  private double x, y;

  public Point(double requiredX, double requiredY)
  {
    x = requiredX;
    y = requiredY;
  } // Point

  public Point()
  {
    x = 0;
    y = 0;
  } // Point

  ...
} // class Point
```

In effect, the second constructor method above is rather like a wrapper around the first one, and we can make this relationship explicit by actually calling the first constructor method from the second. We do this using the **reserved word `this`**, and passing the desired parameters in brackets. So, another way of writing the second constructor above is as follows.

```
  public Point()
  {
    this(0, 0);
  } // Point
```

Such an **alternative constructor call** must be the first **statement** in the body of the constructor method, and, of course, the class must have another constructor method which matches the supplied arguments.

5.23 Method: constructor methods: default (page 425)

If we write a **class** and do not include a **constructor method** in it, then Java implicitly treats it as though we have defined a **public** empty one, which takes no **method arguments**. For example, for a class called `FabulousThing`, this would be as follows.

```
public FabulousThing()
{
} // FabulousThing
```

DRAFT April 15, 2011

This is called a **default constructor** and is of course the same as one which simply invokes the constructor method of the **superclass**.

```
public FabulousThing()
{
  super();
} // FabulousThing
```

The default constructor is only assumed for classes that do not explicitly define a constructor method, which means that not every class actually has a constructor method which takes no arguments. For example, the class VeryFabulousThing, partially defined below, does not have such a constructor method.

```
public class VeryFabulousThing
{
  ... Some code, but no more constructor methods.
  public VeryFabulousThing(String name)
  {
    ...
  } // VeryFabulousThing
  ... Some code, but no more constructor methods.
} // class VeryFabulousThing
```

As a result, the following is illegal.

```
public class TheMostFabulousThingInTheUniverse extends VeryFabulousThing
{
  ... Code here, but no constructor method.
} // class TheMostFabulousThingInTheUniverse
```

This is because the class TheMostFabulousThingInTheUniverse cannot have a default constructor because its superclass does not have a constructor method that takes no arguments.

In practice, default constructors are not often what we want anyway. This author recommends that you *always* explicitly write at least one constructor method for every class which you intend there to be **instance**s of, even when that constructor method is empty. This shows to anybody reading your code that it is deliberately empty, rather than has been omitted by mistake.

5.24 Method: class versus instance methods (page 166)

When we define a **method**, we can write the **reserved word** **`static`** in its heading, meaning that it can be **execute**d in the **static context**, that is, it can be used as soon as the **class** is

DRAFT April 15, 2014

loaded into the **virtual machine**. These are known as **class method**s, because they belong to the class. By contrast, if we omit the `static` **modifier** then the method is an **instance method**. This means it can only be run in a **dynamic context**, attached to a particular **instance** of the class.

This parallels the distinction between **class variable**s and **instance variable**s. There is one copy of a class variable, created when the class is loaded. There is one copy of an instance variable for every instance, created when the instance is created.

We can think of methods in the same way: class methods belong to the class they are defined in, and there is one copy of their code at **run time**, ready for use immediately. Instance methods belong to an instance, and there are as many copies of the code at run time as there are instances. Of course, the virtual machine does not really make copies of the code of instance methods, but it *behaves* as though it does, in the sense that when an instance method is executed, it runs in the context of the instance that it belongs to.

For example, suppose we have a `Point` class with instance variables `x` and `y`. We might wish to have an instance method which takes no **method parameter**s, but **return**s the distance of a point from the origin. Pythagoras[18] tells us that this is $\sqrt{x^2+y^2}$. (We can use the `sqrt()` method from the `Math` class.)

```
public double distanceFromOrigin()
{
  return Math.sqrt(x * x + y * y);
} // distanceFromOrigin
```

A class method can be accessed by taking the name of the class, and appending a dot (`.`) and then the name of the method. `Math.sqrt` is a handy example right now.

An instance method belonging to an **object** can be accessed by taking a **reference** to the *object* and appending a dot (`.`) and then the name of the method. For example, if the **variable** `p1` contains a reference to a `Point` object, then the code `p1.distanceFromOrigin()` invokes the instance method `distanceFromOrigin()`, belonging to the `Point` referred to by `p1`.

The following code would print the numbers `5` and `75`.

```
Point p1 = new Point(3, 4);
Point p2 = new Point(45, 60);

System.out.println(p1.distanceFromOrigin());
System.out.println(p2.distanceFromOrigin());
```

When the method is called via `p1` it uses the instance variables of the object referred to by `p1`, that is the values `3` and `4` respectively. When the method is called via `p2` it uses the values `45` and `60` instead.

DRAFT April 15, 2011

For another example, we may wish to have a method which determines the distance between a point and a given other point.

```
public double distanceFromPoint(Point other)
{
  double xDistance = x - other.x;
  double yDistance = y - other.y;

  return Math.sqrt(xDistance * xDistance + yDistance * yDistance);
} // distanceFromPoint
```

The following code would print the number `70.0`, twice.

```
System.out.println(p1.distanceFromPoint(p2));
System.out.println(p2.distanceFromPoint(p1));
```

5.25 Method: a method may have no parameters (page 173)

The list of **method parameter**s given to a **method** may be empty. This is typical for methods which always have the same effect or **return** the same result, or their result depends on the value of **instance variable**s rather than some values in the context where the method is called.

5.26 Method: return with no value (page 206)

A **void method** may contain **return statement**s which do not have an associated **return** value – just the **reserved word `return`**. These cause the execution of the **method** to end, and control to transfer back to the code that called the method. Every void method behaves as though it has an implicit return statement at the end, unless it has one explicitly written.

The use of return statements throughout the body of a method permits us to design them using a **single entry, multiple exit** principle: every call of the method starts at the beginning, but depending on **condition**s the execution may exit at various points.

5.27 Method: accessor methods (page 207)

A **public instance method** whose job it is to reveal all or some part of the **object state**, without changing it, is known as an **accessor method**. Perhaps the most obvious example of this is an instance method called `getSomeVariable`, where `someVariable` is the name of an **instance variable**. However, a well **design**ed **class** with good **encapsulation** does not systematically

DRAFT April 15, 2011

reveal to its user what its instance variables are. Hence the more general idea of an accessor method: it exposes the value of some *feature*, which might or might not be directly implemented as an instance variable.

5.28 Method: mutator methods (page 207)

A **public instance method** whose job it is to set or update all or some part of the **object state** is known as a **mutator method**. Perhaps the most obvious example of this is an instance method called `setSomeVariable`, where `someVariable` is the name of an **instance variable**. However, the more general idea of a mutator method is that it changes the value of some feature, which might or might not be directly implemented as an instance variable.

Obviously, only **mutable object**s have mutator methods.

5.29 Method: overloaded methods (page 237)

The **method signature** of a method is its name and list of **type**s of its **method parameter**s. Java permits us to have **overloaded method**s, that is, more than one **method** with the same name within one **class**, as long as they have different signatures. E.g. they may have a different number of parameters, different types, the same types but in a different order, etc.. If two methods had the same signature then the **compiler** could never know which one was intended by a **method call** with **method argument**s matching both of them.

For example, the method `System.out.println()` can be used with no arguments, with a single `String` as an argument, or with an argument of some other type, such as **`int`** or any **object**. These are in fact different methods with the same name!

5.30 Method: that throws an exception (page 354)

A **method** has a body of code which is **execute**d when a **method call** invokes it. If it is possible for that code to cause an **exception** to be **throw**n, either directly or indirectly, which is not caught by it, then the method must have a **throws clause** stating this in its heading. We do this by writing the **reserved word** **`throws`** followed by the kind(s) of exception, after the **method parameter** list. For example, the `charAt()` **instance method** of the `java.lang.String` **class** **throws** an exception if the given **string index** is not in range.

```
public char charAt(int index) throws IndexOutOfBoundsException
{
  ...
} // charAt
```

DRAFT April 15, 2011

As another example, suppose in some program we have a **class** which provides **mutable objects** representing customer details. An **instance** of the class is allowed to have the customer name changed, but the new name is not allowed to be empty.

```
public class Customer
{
  private String familyName, firstNames;
  ...
  public void setName(String requiredFamilyName, String requiredFirstNames)
              throws IllegalArgumentException
  {
    if (requiredFamilyName == null || requiredFirstNames == null
        || requiredFamilyName.equals("") || requiredFirstNames.equals(""))
      throw new IllegalArgumentException("Name cannot be null or empty");

    familyName = requiredFamilyName;
    firstNames = requiredFirstNames;
  } // setName
  ...
} // class Customer
```

5.31 Method: that throws an exception: RuntimeException (page 358)

Generally, every **exception** that *possibly* can be **thrown** by a **method**, either directly by a **throw statement** or indirectly via another method, etc., must either be caught by the method, or it must say in its **throws clause** that it **throws** the exception. However, Java relaxes this rule for certain kinds of exception known as `RuntimeException`. These represent common erroneous situations which are usually avoidable and for which we typically write code to ensure they do not happen. The `java.lang.RuntimeException` **class** is a kind of `Exception`, and examples of more specific classes which are kinds of `RuntimeException` include `ArrayIndexOutOfBoundsException`, `IllegalArgumentException`, `NumberFormatException`, `ArithmeticException` and `NullPointerExcept` (all from the `java.lang` **package**).

It would be a major inconvenience to *have* to always declare that these common cases might happen, or to explicitly **catch** them, in situations where we know they will not be **thrown** due to the way we have written the code. So, for these kinds of exception, Java leaves it as an option for us to declare whether they might be thrown by a method. For example, in the following method there is an **array reference**, and also an (implicit) **array element** access. These could in principle result in a `NullPointerException` and an `ArrayIndexOutOfBoundsException` respectively. The Java **compiler** is not clever enough to be able to reason whether such an exception can actually occur, whereas we know they cannot because of the way our code works.

```
private int sum(int[] array)
{
```

DRAFT April 15, 2011

```
    if (array == null)
      return 0;

    int sum = 0;
    for (int element : array)
      sum += element;
    return sum;
  } // sum
```

On the other hand, the following method *can* cause some kinds of `RuntimeException` – if given a **null reference** or an **empty array**. Java still cannot know this without us declaring it in the heading.

```
  private double mean(int[] array)
                 throws NullPointerException, ArrayIndexOutOfBoundsException
  {
    int sum = array[0];
    for (int index = 1; index <= array.length; index++)
      sum += array[index];
    return sum / array.length;
  } // sum
```

For code which is intended for **software reuse**, it is a good idea for us to be disciplined about this relaxation of the normal rule. If we write a method that can throw some exception which is a `RuntimeException`, because we have not written the code in a way which always avoids the possibility, or indeed we explicitly throw such an exception, then we should still declare it in the method heading, even though we are not forced to.

Exceptions for which we must either have a **catch clause** or list in a **throws clause** are known as **checked exception**s, and those for which the rule is relaxed, that is `RuntimeException` and its specific kinds, are known as **unchecked exception**s.

5.32 Method: generic methods (page 522)

A **generic method** is a **method** which has one or more **type parameter**s written within angled brackets (<>) just before the **return type** in the method heading. These are used in a similar way to type parameters of a **generic class**, but apply only to the method. When we write a **method call** for a generic method, we can supply **type argument**s for the type parameters. Generic methods may be defined inside a generic or non-generic class, and may be **instance method**s or **class method**s. However, they are of most use as **class method**s, because generic features of instance methods are *usually* best achieved via type parameters for the whole **class**.

The following symbolic example is a class method with two type parameters.

DRAFT April 15, 2017

```
public static <T1, T2> void myGenericMethod(T1[] anArray, T2 aValue)
{
  ... Code here that uses T1 and T2 as types.
  ... Some restrictions apply,
  ... such as we cannot make instances of T1, or T2.
} // myGenericMethod
```

This takes (a **reference** to) an **array** of some **type**, `T1[]` and also (a reference to) an **object** of type `T2`. The actual types for `T1` and `T2` can be supplied as type arguments when the method is called. For example, assuming the above is defined in a class called `MyClassWithGenericMethod`, then the following could be a call to it.

```
Date[] aDateArray = ...
String aString = ...

MyClassWithGenericMethod.<Date, String>myGenericMethod(aDateArray, aString);
```

Notice that the type arguments are written, within angled brackets (`<>`), *after* the dot (`.`) separating the class name from the method name – they are not class type parameters, and so are not written after the class name, but instead they occur before the method name. There is also a peculiarity to watch out for. Normally, if we call a class method from within the class where it is defined, we do not need to prepend the class name and a dot(`.`), but if we are going to supply type arguments then we must. And for a generic instance method, we must similarly use the **this reference** and prepend **`this`** and a dot.

However, the good news is that we can *omit* the type arguments completely when we call a generic method, and in nearly all cases the **compiler** is able to work them out from the types of the **method argument**s.

5.33 Method: generic methods: bound type parameter (page 526)

The **type parameter**s of a **generic method** can be **bound type parameter**s as they can be with **generic class**es. For example, here is a **class method** that **return**s the largest element, according to **natural ordering**, of an **array** of items which are `Comparable` with themselves.

```
public class MaxArray
{
  public static <ArrayType extends Comparable<ArrayType>>
                ArrayType getMax(ArrayType[] anArray)
                throws IllegalArgumentException
  {
    try
    {
```

DRAFT April 15, 2011

```
      ArrayType result = anArray[0];
      for (int index = 1; index < anArray.length; index++)
        if (result.compareTo(anArray[index]) < 0)
          result = anArray[index];
      return result;
    } // try
    catch (ArrayIndexOutOfBoundsException e)
    { throw new IllegalArgumentException("Array must be non-empty", e); }
    catch (NullPointerException e)
    { throw new IllegalArgumentException("Array must exist", e); }
  } // getMax

} // class MaxArray
```

This could be used as follows.

```
String[] aStringArray = { "the", "cat", "vaporized", "on", "the", "mat" };
String maxInAStringArray = MaxArray.getMax(aStringArray);
```

The **compiler** was able to figure out the **type argument**, and so our **method call** above is equivalent to the following.

```
String maxInAStringArray = MaxArray.<String>getMax(aStringArray);
```

6 Command line arguments

6.1 Command line arguments: program arguments are passed to main (page 17)

Programs can be given **command line argument**s which typically affect their behaviour. Arguments given to a Java program are strings of text **data**, and there can be any number of them in a **list**. In Java, `String[]` means 'list of strings'. We have to give a name for this list, and usually we call it `args`. The chosen name allows us to refer to the given data from within the program, should we wish to.

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

DRAFT April 15, 2011

6.2 Command line arguments: program arguments are accessed by index (page 26)

The **command line argument**s given to the **main method** are a **list** of strings. These are the **text data string** arguments supplied on the **command line**. The strings are **index**ed by **integer**s (whole numbers) starting from zero. We can access the individual strings by placing the index value in square brackets after the name of the list. So, assuming that we call the list `args`, then `args[0]` is the first command line argument given to the program, if there is one.

6.3 Command line arguments: length of the list (page 79)

The **command line argument**s passed to the **main method** are a **list** of strings. We can find the length of a list by writing a dot followed by the word `length`, after the name of the list. For example, `args.length` yields an **`int`** value which is the number of items in the list `args`.

6.4 Command line arguments: list index can be a variable (page 79)

The **index** used to access the individual items from a **list** of strings does not have to be an **integer literal**, but can be an **int variable** or indeed an **arithmetic expression**. For example, the following code adds together a list of **integer**s given as **command line argument**s.

```
int sumOfArgs = 0;
for (int argIndex = 0; argIndex < args.length; argIndex = argIndex + 1)
  sumOfArgs = sumOfArgs + Integer.parseInt(args[argIndex]);
System.out.println("The sum is " + sumOfArgs);
```

The benefit of being able to use a **variable**, rather than an integer literal is that the access can be done in a **loop** which controls the value of the variable: thus the actual value used as the index is not the same each time.

7 Type

7.1 Type (page 36)

Programs can process various different kinds of **data**, such as numbers, text data, images etc.. The kind of a data item is known as its **type**.

DRAFT April 15, 2011

7.2 Type: String (page 135)

The **type** of **text data string**s, such as **string literal** values and **concatenation**s of such, is called `String` in Java.

7.3 Type: String: literal (page 18)

In Java, we can have a **string literal**, that is a fixed piece of text to be used as **data**, by enclosing it in double quotes. It is called a string literal, because it is a **type** of data which is a string of **character**s, exactly as listed. Such a piece of data might be used as a message to the user.

```
"This is a fixed piece of text data -- a string literal"
```

7.4 Type: String: literal: must be ended on the same line (page 21)

In Java, **string literal**s must be ended on the same line they are started on.

7.5 Type: String: literal: escape sequences (page 49)

We can have a **new line character** embedded in a **string literal** by using the **escape sequence** `\n`. For example, the following code will print out three lines on **standard output**.

```
System.out.println("This text\nspans three\nlines.");
```

It will generate the following.

```
This text
spans three
lines.
```

There are other escape sequences we can use, including the following.

DRAFT April 15, 2019

Sequence	Name	Effect
`\b`	Backspace	Moves the cursor back one place, so the next **character** will over-print the previous.
`\t`	Tab (horizontal tab)	Moves the cursor to the next 'tab stop'.
`\n`	New line (line feed)	Moves the cursor to the next line.
`\f`	Form feed	Moves to a new page on many (text) printers.
`\r`	Carriage return	Moves the cursor to the start of the current line, so characters will over-print those already printed.
`\"`	Double quote	Without the backslash escape, this would mark the end of the string literal.
`\'`	Single quote	This is just for consistency – we don't need to escape a single quote in a string literal.
`\\`	Backslash	Well, sometimes you want the backslash character itself.

Note: `System.out.println()` always ends the line with the platform dependent **line separator**, which on Linux is a new line character but on Microsoft Windows is a **carriage return character** followed by a new line character. In practice you may not notice the difference, but the above code is not strictly the same as using three separate `System.out.println()` calls and is not 100% portable.

7.6 Type: String: concatenation (page 26)

The `+` **operator**, when used with two string **operand**s, produces a string which is the **concatenation** of the two strings. For example `"Hello " + "world"` produces a string which is `Hello ` (including the space) concatenated with the string `world`, and so has the same value as `"Hello world"`.

There would not be much point concatenating together two **string literal**s like this, compared with having one string literal which is already the text we want. We would be more likely to use concatenation when at least one of the operands is not a fixed value, i.e. is a **variable** value. For example, `"Hello " + args[0]` produces a string which is `Hello ` (including the space) concatenated with the first **command line argument** given when the program is **run**.

The resulting string can be used anywhere that a single string literal could be used. For example `System.out.println("Hello " + args[0])` would print the resulting string on the **standard output**.

7.7 Type: String: conversion: from int (page 38)

The Java **operator** `+` is used for both **addition** and **concatenation** – it is an **overloaded operator**. If at least one of the **operand**s is a **text data string**, then Java uses concatenation, otherwise it uses addition. When only one of the two operands is a string, and the other is

DRAFT April 15, 2011

some other **type** of **data**, for example an `int`, the Java **compiler** is clever enough to understand the programmer wishes that data to be converted into a string before the concatenation takes place. It is important to note the difference between an **integer** and the decimal digit string we usually use to represent it. For example, the **integer literal** `123` is an `int`, a number; whereas the **string literal** `"123"` is a text data string – a string of 3 separate **character**s.

Suppose the **variable** `noOfPeopleToInviteToTheStreetParty` had the value `51`, then the code

```
System.out.println("Please invite " + noOfPeopleToInviteToTheStreetParty);
```

would print out the following text.

```
Please invite 51
```

The number `51` would be converted to the string `"51"` and then concatenated to the string `"Please invite "` before being processed by `System.out.println()`.

Furthermore, for our convenience, there is a separate version of `System.out.println()` that takes a single `int` rather than a string, and prints its decimal representation. Thus, the code

```
System.out.println(noOfPeopleToInviteToTheStreetParty);
```

has the same effect as the following.

```
System.out.println("" + noOfPeopleToInviteToTheStreetParty);
```

7.8 Type: String: conversion: from double (page 55)

The Java **concatenation operator**, `+`, for joining **text data string**s can also be used to convert a `double` to a string. For example, the **expression** `"" + 123.4` has the value `"123.4"`.

7.9 Type: String: conversion: from object (page 177)

It is quite common for **class**es to have an **instance method** which is **design**ed to produce a `String` representation of an **object**. It is conventional in Java for such **method**s to be called `toString`. For example, a `Point` class with `x` and `y` **instance variable**s might have the following `toString()` method.

DRAFT April 15, 2011

```
public String toString()
{
  return "(" + x + "," + y + ")";
} // toString
```

For convenience, whenever the Java **compiler** finds an **object reference** as an **operand** of the **concatenation operator** it assumes that the object's `toString()` method is to be invoked to produce the required `String`. For example, consider the following code.

```
Point p1 = new Point(10, 40);
System.out.println("The point is " + p1.toString());
```

Thanks to the compiler's convenient implicit assumption about `toString()`, the above code could, and probably would, have been written as follows.

```
Point p1 = new Point(10, 40);
System.out.println("The point is " + p1);
```

For our further convenience, there is a separate version of `System.out.println()` that takes any single object rather than a string, and prints its `toString()`. Thus, the code

```
System.out.println(p1);
```

has the same effect as the following.

```
System.out.println("" + p1);
```

7.10 Type: String: conversion: from object: null reference (page 211)

For convenience, whenever the Java **compiler** finds an **object reference** as an **operand** of the **concatenation operator** it assumes that the object's `toString()` **instance method** is to be invoked to produce the required `String`. However, the reference might be the **null reference** in which case there is no object on which to invoke `toString()`, so instead, the string `"null"` is used.

In fact, assuming `someString` is some `String` and `myVar` is a **variable** of a **reference type**, then the code:

```
someString + myVar
```

DRAFT April 15, 2011

is actually treated as follows.

```
someString + (myVar == null
              ? "null"
              : (myVar.toString() == null ? "null" : myVar.toString()))
```

The same applies to the first operand of string concatenation if that is an object reference.

For this reason, most Java programmers prefer to use `"" + myVar` rather than `myVar.toString()` when they wish to convert the object referenced by `myVar` to a string, because it avoids the possibility of an **exception** if `myVar` contains the null reference.

7.11 Type: int (page 36)

One of the **type**s of **data** we can use in Java is called `int`. A data item which is an `int` is an **integer** (whole number), such as `0`, `-129934` or `982375`, etc..

7.12 Type: double (page 54)

Another of the **type**s of **data** we can use in Java is known as `double`. A data item which is a `double` is a **real** (fractional decimal number), such as `0.0`, `-129.934` or `98.2375`, etc.. The type is called `double` because it uses a means of storing the numbers called **double precision**. On computers, real numbers are only approximated, because they have to be stored in a finite amount of memory space, whereas in mathematics we have the notion of infinite decimals. The double precision storage approach uses twice as much memory per number than the older **single precision** technique, but gives numbers which are much more precise.

7.13 Type: casting an int to a double (page 79)

Sometimes we have an `int` value which we wish to be regarded as a `double`. The process of conversion is known as **cast**ing, and we can achieve it by writing `(double)` in front of the `int`. For example, `(double)5` is the `double` value `5.0`. Of course, we are most likely to use this feature to cast the value of an **int variable**, rather than an **integer literal**.

7.14 Type: boolean (page 133)

There is a **type** in Java called `boolean`, and this is the type of all **condition**s used in **if else statement**s and **loop**s. It is named after the English mathematician, George Boole whose work

DRAFT April 15, 2011

in 1847 established the basis of modern logic[12]. The type contains just two **boolean literal** values called `true` and `false`. For example, `5 <= 5` is a **boolean expression**, which, because it has no **variable**s in it, always has the same value when **evaluate**d. Whereas the **expression** `age1 < age2 || age1 == age2 && height1 <= height2` has a value which depends on the values of the variables in it.

7.15 Type: long (page 145)

The **type** `int` allows for the storage of **integer**s in the range -2^{31} through to $2^{31}-1$. This is because it uses four **byte**s, i.e. 32 **binary digit**s. $2^{31}-1$ is 2147483647. Although this is plenty for most purposes, we sometimes need whole numbers in a bigger range. The type `long` represents **long integer**s and uses eight bytes, i.e. 64 **bit**s. A `long` **variable** can store numbers from -2^{63} through to $2^{63}-1$. The value of $2^{63}-1$ is 9223372036854775807.

A **long literal** is written with an `L` on the end, to distinguish it from an **int literal**, as in `-15L` and `2147483648L`.

7.16 Type: short (page 145)

The **type** `short` represents **short integer**s using two **byte**s, i.e. 16 **binary digit**s. A **short variable** can store numbers from -2^{15} through to $2^{15}-1$. The value of $2^{15}-1$ is 32767. We would typically use this type when we have a huge number of **integer**s, which happen to lie in the restricted range, and we are concerned about the amount of memory (or **file** space) needed to store them.

7.17 Type: byte (page 145)

The **type** `byte` represents **integer**s using just one **byte**, i.e. 8 **binary digit**s. A `byte` **variable** can store numbers from -2^{7} through to $2^{7}-1$. The value of $2^{7}-1$ is 127.

7.18 Type: char (page 145)

Characters in Java are represented by the **type** `char`. A **char variable** can store a single **character** at any time.

7.19 Type: char: literal (page 145)

A **character literal** can be written in our program by enclosing it in single quotes. For example `'J'` is a character literal.

7.20 Type: char: literal: escape sequences (page 146)

When writing a **character literal** we can use the same **escape sequence**s that are available within **string literal**s. These include the following.

```
char backspace = '\b';              char tab = '\t';
char newline = '\n';                char formFeed = '\f';
char carriageReturn = '\r';         char doubleQuote = '\"';
char singleQuote = '\'';            char backslash = '\\';
```

7.21 Type: char: comparisons (page 238)

Values of **type `char`** may be compared using the usual `<`, `<=`, `==`, `!=`, `>=` and `>` **relational operator**s. Characters are stored in the computer using numeric **character code**s – each one has a unique number – and when two **character**s are compared, the result is formed from the same comparison on the two numbers.

Generally speaking we do not need to know the actual numbers used for specific characters. However, there are certain properties that are useful to know, such as that the number for `'A'` is one **less than** that for `'B'`, which is one less than the number used for `'C'`, and so on. In other words, the upper case alphabetic letters have contiguous character codes. The same is true of the lower case alphabet, and also the digit characters `'0'` through to `'9'`. The character codes for the digits are all less than those for the upper case letters, which are all less than those for the lower case letters.

For example, the following **method** checks whether a given character is a lower case alphabetic character.

```
public static boolean isLowerCase(char aChar)
{
  return aChar >= 'a' && aChar <= 'z';
} // isLowerCase
```

A method similar to this is provided in the standard **class** `java.lang.Character`. That one also works for **locale**s (i.e. languages) other than English.

DRAFT April 15, 2011

Another property worth remembering is that, for the English characters, the code for each upper case letter is 32 less than the code for the corresponding lower case letter.

7.22 Type: char: casting to and from int (page 238)

The numeric **character code** used to store a **character** may be obtained by **cast**ing a `char` value to an `int`. We can achieve this by writing `(int)` in front of it. For example, `(int)'A'` is the numeric code used to store a capital A.

We can also convert in the opposite direction, by casting an `int` to a `char`. For example, at the end of the following fragment of code, the **variable** `letterB` will contain an upper case B character.[2]

```
int codeForA = (int)'A';
char letterB = (char) (codeForA + 1);
```

The following **method return**s the upper case equivalent of a given character, if it is a lower case letter, or the original character if not. It assumes availability of the method `isLowerCase()`.

```
public static char toUpperCase(char aChar)
{
  if (isLowerCase(aChar))
    return (char) ((int)aChar - (int)'a' + (int)'A');
  else
    return aChar;
} // toUpperCase
```

A method similar to this is provided in the standard **class** `java.lang.Character`. That one also works for **locale**s (i.e. languages) other than English.

7.23 Type: float (page 146)

The **type** `float` is for **real** (fractional decimal) numbers, using the **floating point representation** with a **single precision** storage. It uses only four **byte**s per number, compared with `double` which employs **double precision** storage and so is far more accurate, but needs eight bytes per number.

A **float literal** is written with an `f` or `F` on the end, as in `0.0F`, `-129.934F` or `98.2375f`.

[2]Actually, the cast in the first line from `char` to `int` would be implicit, but it is good style to write it anyway. In the second line, the cast from `int` to `char` is required.

DRAFT April 15, 2011

7.24 Type: primitive versus reference (page 162)

Each **type** in Java is either a **primitive type** or a **reference type**. Values of primitive types have a size which is known at **compile time**. For example, every **int** value comprises four **byte**s. Types for which the size of an individual value is only known at **run time**, such as **class**es, are known as reference types because the values are always accessed via a **reference**.

7.25 Type: array type (page 287)

Whilst it is true that **array**s in Java are **object**s, they are treated somewhat differently from **instance**s of **class**es. To obtain an **array type**, we do not write a class and then use its name. Instead we simply write the **type** of the **array element**s followed by a left and then a right square bracket ([]). The type of the elements is known as the **array base type**.

For example, int[] is the type of arrays with **int** as the base type, that is ones which contain elements that are **int** values. String[] is the type of arrays which contain elements that are **reference**s to String objects.

7.26 Type: enum type (page 309)

An **enum type** is a feature which arrived in Java 5.0 that allows us to identify a **type** with an enumeration of named values. For example, we might have four possible directions in some game involving movement.

```
private enum Direction { UP, DOWN, LEFT, RIGHT }
```

This behaves rather like we have defined a **class** called Direction, and four **variable**s, each referring to a unique **instance** of Direction. So, for example, we can have the following.

```
private Direction currentDirection = Direction.UP;
private Direction nextDirection = null;
```

If we wanted the type to be available in other classes, then we would declare it as **public**.

Enum types can also be used in **switch statement**s.

```
switch (currentDirection)
{
  case UP:    ...
```

DRAFT April 15, 2011

```
    case DOWN:  ...
    case LEFT:  ...
    case RIGHT: ...
    default:    ...
  } // switch
```

7.27 Type: enum type: access from another class (page 312)

If we declare a **public enum type**, then it can be used in other **class**es. We access it using dots (`.`) rather like we do for other kinds of access from another class.

For example, if the enum type `Direction` is defined in the class `Movement`, then we could refer to it, and one of its values as follows.

```
Movement.Direction requestedDirection = Movement.Direction.UP;
```

8 Standard API

8.1 Standard API: System: out.println() (page 18)

The simplest way to print a message on **standard output** is to use:

```
System.out.println("This text will appear on standard output");
```

`System` is a **class** (that is, a piece of code) that comes with Java as part of its **application program interface** (**API**) – a large number of classes designed to support our Java programs. Inside `System` there is a thing called `out`, and this has a **method** (section of code) called `println`. So overall, this method is called `System.out.println`. The method takes a string of text given to it in its brackets, and displays that text on the standard output of the program.

8.2 Standard API: System: out.println(): with no argument (page 98)

The **class** `System` also contains a version of the `out.println()` **method** which takes no arguments. This outputs nothing except a **new line**. It has the same effect as calling `System.out.println()` with an empty string as its argument, that is

```
System.out.println();
```

DRAFT April 15, 2011

has the same effect as the following.

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

So, for example

```
System.out.print("Hello world!");
System.out.println();
```

would have the same effect as the following.

```
System.out.println("Hello world!");
```

`System.out.println()` with no argument is most useful when we need to end a line which has been generated a piece at a time, or when we want to have a blank line.

8.3 Standard API: System: out.println(): with any argument (page 427)

The **class**
`java.lang.System` has an **overloaded method** version of `out.println()` and `out.print()` for every **primitive type** of **method argument**, as well as `java.lang.Object`. Each treats its argument, `(arg)`, as `("" + arg)`. So, an **`int`** is output in decimal representation, and a non-null **object reference** has its `toString()` **instance method** used, etc..

Also, there is a version of `System.out.println()` and `System.out.print()` that take a **character array**, `char[]`, and print the characters in it.

8.4 Standard API: System: out.print() (page 98)

The **class** `System` contains a **method** `out.print()` which is almost the same as `out.println()`. The only difference is that `out.print()` does not produce a **new line** after printing its output. This means that any output printed after this will appear on the same line. For example

```
System.out.print("Hello");
System.out.print(" ");
System.out.println("world!");
```

would have the same effect as the following.

DRAFT April 15, 2011

```
System.out.println("Hello world!");
```

`System.out.print()` is most useful when the output is being generated a piece at a time, often within a **loop**.

8.5 Standard API: System: out.printf() (page 126)

The **class** `System` contains a **method** `out.printf()`, introduced in Java 5.0, which is similar to `out.print()` except that we can use it to produce formatted output of values.

A simple use of this is to take an **integer** value and have it printed with **space padding** to a given positive integer field width. This means the output contains leading spaces followed by the usual representation of the integer, such that the number of **character**s printed is at least the given field width.

The following code fragment includes an example which prints a string representation of `123`, with leading spaces so that the result has a width of ten characters.

```
System.out.println("1234567890");
System.out.printf("%10d%n", 123);
```

Here is the effect of these two **statement**s.

```
1234567890
       123
```

The first `%` tells `out.printf()` that we wish it to format something, the `10` tells it the minimum total width to produce, and the following letter says what kind of conversion to perform. A `d` tells it to produce the representation of a decimal whole number, which is given after the **format specifier** string, as the second **method argument**. The `%n` tells `out.printf()` to output the platform dependent **line separator**.

The method can be asked to format a floating point value, such as a **`double`**. In such cases we give the minimum total width, a dot (`.`), the number of decimal places, and an `f` conversion. For example,

```
System.out.printf("%1.2f%n", 123.456);
```

needs more than the given minimum width of `1`, and so produces the following.

DRAFT April 15, 2014

```
123.46
```

Whereas, the format specifier in

```
System.out.println("1234567890");
System.out.printf("%10.2f%n", 123.456);
```

prints a total of ten characters for the number, two of which are decimal places.

```
1234567890
    123.46
```

8.6 Standard API: System: out.printf(): zero padding (page 140)

We can ask
`System.out.printf()` for **zero padding** rather than **space padding** of a number by placing a leading zero on the desired minimum width in the **format specifier**.

The following code fragment contains an example which prints a string representation of `123`, with leading zeroes so that the result is ten **character**s long.

```
System.out.println("1234567890");
System.out.printf("%010d%n", 123);
```

Here is the effect.

```
1234567890
0000000123
```

Similarly,

```
System.out.println("1234567890");
System.out.printf("%010.2f%n", 123.456);
```

produces the following.

```
1234567890
0000123.46
```

DRAFT April 15, 2011

8.7 Standard API: System: out.printf(): string item (page 289)

We can ask
`System.out.printf()` to print a `String` item by using `s` as the conversion **character** in the **format specifier**. For example,

```
System.out.println("123456789012345");
System.out.printf("%15s%n", "Hello World");
```

has this effect.

```
123456789012345
    Hello World
```

If the item following the format specifier string is not itself a string, but some other **object** then its `toString()` is used. For example, assuming a `Point` **class** is defined as expected, then the code

```
System.out.println("123456789012345");
System.out.printf("%15s%n", new Point(3,4));
```

produces the following.

```
123456789012345
      (3.0,4.0)
```

8.8 Standard API: System: out.printf(): fixed text and many items (page 289)

We can give `System.out.printf()` a format string with more than one **format specifier** in it, together with more than one value to be printed. What is more, any text in the format string which is not part of a format specifier is simply printed as it appears. Also, if no width is given for a format specifier then its natural width is used.

For example,

```
Point p1 = new Point(3,4);
Point p2 = new Point(45, 60);
System.out.printf("The distance between %s and %s is %1.2f.%n",
                  p1, p2, p1.distanceFromPoint(p2));
```

DRAFT April 15, 2011

produces the following output.

```
The distance between (3.0,4.0) and (45.0,60.0) is 70.00.
```

8.9 Standard API: System: out.printf(): left justification (page 300)

If we wish an item printed by `System.out.printf()` to be left justified, rather than right justified, then we can place a hyphen in front of the width in the **format specifier**.

For example,

```
System.out.println("123456789012345X");
System.out.printf("%-15sX%n", "Hello World");
```

produces the following.

```
123456789012345X
Hello World    X
```

8.10 Standard API: System: in (page 187)

Inside the `System` **class**, in addition to the **class variable** called `out`, there is another called `in`. This contains a **reference** to an **object** which represents the **standard input** of the program.

Perhaps surprisingly, unlike the **standard output**, the standard input in Java is not easy to use as it is, and we typically access it via some other means, such as a `Scanner`.

8.11 Standard API: System: in: is an InputStream (page 452)

The **class variable** called `in`, inside the `java.lang.System` **class** (i.e. `System.in`) holds a **reference** to an **object** which is an **instance** of `java.io.InputStream`. This enables our programs to access the **byte**s of their **standard input**.

8.12 Standard API: System: getProperty() (page 195)

When a program is **run**ning, various **system property** values hold information about such things as the Java version and platform being used, the home directory of the user, etc.. The

DRAFT April 15, 2011

class method `System.getProperty()` takes the name of such a property as its `String` **method parameter** and **return**s the corresponding `String` value.

8.13 Standard API: System: getProperty(): line.separator (page 195)

`System.getProperty()` maps the name `line.separator` onto the **system property** which is the **line separator** for the platform in use.

8.14 Standard API: System: currentTimeMillis() (page 262)

The **class** `java.lang.System` contains a **class method** called `currentTimeMillis` which **return**s the current date and time expressed as the number of milliseconds since midnight, January 1, 1970. This value is a **`long`**.

8.15 Standard API: System: err.println() (page 344)

Inside the `java.lang.System` **class**, in addition to **class variable**s called `out` and `in` there is another called `err`. This contains a **reference** to an **object** which represents the **standard error** of the program. Via this object we have the **method**s `System.err.println()` and `System.err.print()`. These cause their given **method argument**s to be displayed on the standard error.

8.16 Standard API: System: out: is an OutputStream (page 468)

The **class variable** called `out`, inside the `java.lang.System` **class** (i.e. `System.out`) holds a **reference** to an **object** which is an **instance** of `java.io.OutputStream`. This enables our programs to produce **byte**s on their **standard output**.

More precisely, `System.out` is an instance of `java.io.PrintStream`, which is a **subclass** of `OutputStream`. Unlike basic `OutputStream` objects, a `PrintStream` object also has **instance method**s `print()`, `println()` and (since Java 5.0) `printf()`, which take various **method argument**s and write their **character** representations as bytes.

8.17 Standard API: System: err: is an OutputStream (page 468)

The **class variable** called `err`, inside the `java.lang.System` **class** (i.e. `System.err`) holds a **reference** to an **object** which is an **instance** of `java.io.PrintStream`, a **subclass** of

DRAFT April 15, 2011

`java.io.OutputStream`. This enables our programs to produce **byte**s on their **standard error**.

8.18 Standard API: Integer: parseInt() (page 41)

One simple way to turn a **text data string**, say `"123"` into the **integer** (whole number) it represents is to use the following.

```
Integer.parseInt("123");
```

`Integer` is a **class** (that is, a piece of code) that comes with Java. Inside `Integer` there is a **method** (section of code) called `parseInt`. This method takes a text data string given to it in its brackets, converts it into an `int` and **return**s that number. A **run time error** will occur if the given string does not represent an `int` value.

For example

```
int firstArgument;
firstArgument = Integer.parseInt(args[0]);
```

would take the first **command line argument** and, assuming it represents a number (i.e. it is a string of digits with a possible sign in front), would turn it into the number it represents, then store that number in `firstArgument`. If instead the first argument was some other text data string, it would produce a run time error.

8.19 Standard API: Integer: as a box for int (page 487)

In addition to containing **class method**s to manipulate **integer** related values, the standard **class** `java.lang.Integer` can be used to wrap up `int` values as **object**s. One of the **constructor method**s of the class may be given an `int`, and this makes an **instance** of `Integer` wrapping up, or **box**ing, that number. The **instance method** `intValue()` can then later be used to retrieve the boxed number from the object. This effectively allows an `int`, which is a **primitive type**, to be treated as though it is an **object**.

8.20 Standard API: Integer: as a box for int: autoboxing (page 494)

Use of the standard **class** `java.lang.Integer` to wrap up `int` values as **object**s is so common, that since Java 5.0 the **compiler** can make their use implicit by providing **autoboxing** and

DRAFT April 15, 2011

auto-unboxing. Whenever an **`int`** value is given where an `Integer` is required, the **`int`** is automatically **box**ed (wrapped up) into a **new** `Integer` object. And whenever (a **reference** to) an `Integer` is given where an **`int`** is required, the `intValue()` **instance method** is automatically used to unbox the **`int`** value.

For example, here is some code that explicitly wraps up and extracts an **`int`**.

```
Integer anInteger = new Integer(10);
int anInt = anInteger.intValue() + 1;
System.out.println(anInt);
```

The following code would have exactly the same effect – both would print out `11`.

```
Integer anInteger = 10;
int anInt = anInteger + 1;
System.out.println(anInt);
```

Whilst this convenience can often make the **`int`** and `Integer` **type**s work seamlessly together, it is important to remember the difference between them. **`int`** is a **primitive type**, whereas `Integer` is a **reference type**. So, for example, an **array** of ten **`int`** values would take as much memory as ten times the space of one **`int`** value (plus a little). By contrast, an array of ten `Integer` objects would hold ten **reference**s, each referring to an object storing an **`int`** value.

8.21 Standard API: Integer: as a box for int: works with collections (page 548)

The standard **class** `java.lang.Integer` **implement**s `java.lang.Comparable<Integer>`, and provides the instance methods `compareTo()`, and also `equals()` and `hashCode()` in such a way that `Integer` **object**s behave properly as `Comparable`s and in **hash table**s, etc..

8.22 Standard API: Double: parseDouble() (page 54)

One simple way to turn a **text data string**, say `"123.456"` into the **real** (fractional decimal number) it represents is to use the following.

```
Double.parseDouble("123.456");
```

`Double` is a **class** (that is, a piece of code) that comes with Java. Inside `Double` there is a **method** (section of code) called `parseDouble`. This method takes a text data string given to it in its brackets, converts it into an **`double`** and **return**s that number. A **run time error** will occur if the given string does not represent a number. For example

DRAFT April 15, 2011

```
double firstArgument = Double.parseDouble(args[0]);
```

would take the first **command line argument** and, assuming it represents a number, would turn it into the number it represents, then store that number in `firstArgument`. To represent a number, the string must be a sequence of digits, possibly with a decimal point and maybe a negative sign in front. If instead the first argument was some other text data string, it would produce a run time error.

8.23 Standard API: Math: pow() (page 73)

Java does not have an **operator** to compute powers. Instead, there is a standard **class** called `Math` which contains a collection of useful **method**s, including `pow()`. This takes two numbers, separated by a comma, and gives the value of the first number raised to the power of the second.

For example, the **expression** `Math.pow(2, 10)` produces the value of 2^{10} which is `1024`.

8.24 Standard API: Math: abs() (page 87)

Java does not have an **operator** to yield the **absolute value** of a number, that is, its value ignoring its sign. Instead, the standard **class** called `Math` contains a **method**, called `abs`. This method takes a number and gives its absolute value.

For example, the **expression** `Math.abs(-2.7)` produces the value `2.7`, as does the expression `Math.abs(3.4 - 0.7)`.

8.25 Standard API: Math: PI (page 87)

The standard **class** called `Math` contains a constant value called `PI` that is set to the most accurate value of π that can be represented using the **`double`** number **type**. We can refer to this value using `Math.PI`, as in the following example.

```
double circleArea = Math.PI * circleRadius * circleRadius;
```

8.26 Standard API: Math: random() (page 205)

The standard **class** `java.lang.Math` contains a **class method** called `random`. This takes no **method argument**s and **return**s some **`double`** value, r, such that $0.0 \leq r < 1.0$ is true. The value is chosen in a pseudo random fashion, using an **algorithm** which exhibits the characteristics of an approximately uniform distribution of random numbers.

DRAFT April 15, 2011

8.27 Standard API: Math: round() (page 289)

The standard **class** java.lang.Math contains a **class method** called round. This takes a **double** **method argument** and **return**s a **long** value which is the nearest whole number to the given one. If we wish to turn that result into an **int** then we would of course **cast** it, as in the following example.

```
int myPennies = ... Obtain this somehow.
int myNearlyPounds = (int) Math.round(myPennies / 100.0);
```

8.28 Standard API: Scanner (page 188)

Since the advent of Java 5.0 there is a standard **class** called java.util.Scanner which provides some simple features to read input **data**. In particular, it can be used to read System.in by passing that to its **constructor method** as follows.

```
import java.util.Scanner;
  ...
  Scanner inputScanner = new Scanner(System.in);
  ...
```

Each time we want a line of text we invoke the nextLine() **instance method**.

```
String line = inputScanner.nextLine();
...
```

Or maybe we want to read an **integer** using nextInt().

```
int aNumber = inputScanner.nextInt();
// Skip past anything on the same line following the number.
inputScanner.nextLine();
...
```

Essentially, System.in accesses the **standard input** as a stream of **byte**s of data. A Scanner turns these bytes into a stream of **character**s (i.e. **char** values) and offers a variety of instance methods to scan these into whole lines, or various tokens separated by **white space**, such as spaces, tabs and end of lines. Some of these instance methods are listed below.

DRAFT April 15, 2011

Public method interfaces for class Scanner (some of them).

Method	Return	Arguments	Description
nextLine	String		Returns all the text from the current point in the character stream up to the next end of line, as a String.
nextInt	**int**		Skips any spaces, tabs and end of lines and then reads characters which represent an integer, and **return**s that value as an **int**. It does not skip spaces, tabs or end of lines following those characters. The characters must represent an integer, or a **run time error** will occur.
nextBoolean	**boolean**		Similar to nextInt() except for a **boolean** value.
nextByte	**byte**		Similar to nextInt() except for a **byte** value.
nextDouble	**double**		Similar to nextInt() except for a **double** value.
nextFloat	**float**		Similar to nextInt() except for a **float** value.
nextLong	**long**		Similar to nextInt() except for a **long** value.
nextShort	**short**		Similar to nextInt() except for a **short** value.

There are very many more features in this class, including the ability to change what is considered to be characters that separate the various tokens.

8.29 Standard API: Scanner: for a file (page 306)

The standard **class** java.util.Scanner can be used to read the contents of a **file**, such as my-data.txt, as follows.

```
import java.io.File;
import java.util.Scanner;
  ...
  Scanner input = new Scanner(new File("my-data.txt"));
```

java.io.File is a standard class used to represent file names.

Having obtained a Scanner for the file, we can then use its various **instance method**s, such as nextLine(), to read the **data**.

If we desire to read every line of the file, we might also use the hasNextLine() instance method – this **return**s **true** or **false** depending on whether there are more lines in the file.

```
  while (input.hasNextLine())
```

DRAFT April 15, 2021

```
{
  String line = input.nextLine();
  ...
} // while
```

8.30 Standard API: String (page 233)

Strings in Java are **object**s of the standard **class** `java.lang.String`. This class is defined in the same way as any other, but the Java language also knows about **string literal**s and the string **concatenation operator**. So, strings are semi-built-in to Java. All the other built-in types are **primitive type**s, but `String` is a **reference type**.

When we write

```
String name = "Java";
```

we are asking for an object of **type** `String` to be created, containing the text `Java`, and for a **reference** to that object to be placed in the **variable** called `name`. So, even though we do not use the special word **`new`**, whenever we write a string literal in our code, we are asking for a **new** `String` object to be created.

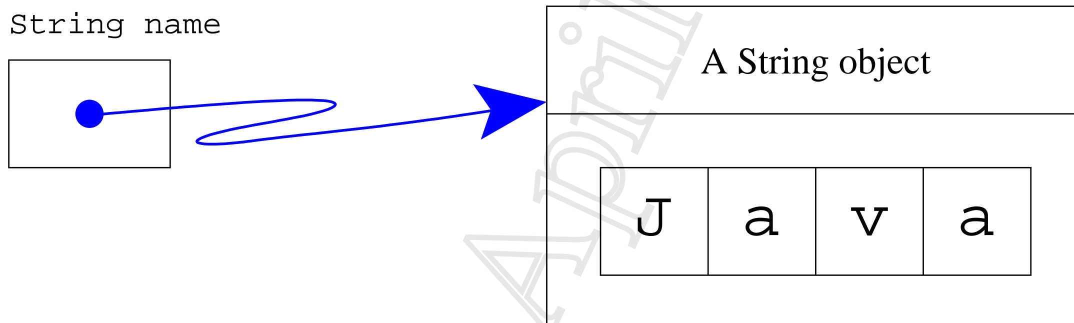

The text of a `String` is stored as a sequence of **character**s, each of these is a member of the **`char`** type. This text cannot be changed: `String`s are **immutable object**s.

8.31 Standard API: String: some instance methods (page 234)

Strings have **instance method**s, some of which are listed below.

DRAFT April 15, 2011

Public method interfaces for class `String` (some of them).

Method	Return	Arguments	Description
charAt	**char**	**int**	This **return**s the **character** at the specified **string index**. The characters are indexed from zero upwards.
compareTo	**int**	String	Compares the text of this with the given other, using **lexicographic order**ing (alphabetic/dictionary order). Returns 0 if they are equal, a negative **int** if this is **less than** the other, a positive **int** otherwise.
endsWith	**boolean**	String	Returns **true** if and only if the text of this string ends with that of the given other.
equals	**boolean**	String	Returns **true** if and only if this string contains the same text as the given other.
indexOf	**int**	String	Returns the index within this string of the first occurrence of the given other string, or -1 if it does not occur.
length	**int**		Returns the length of this string.
startsWith	**boolean**	String	Returns **true** if and only if the text of this string starts with that of the given other.
substring	String	**int**	Returns a **new** string that is a substring of this string. The substring begins with the character at the given index and extends to the end of this string.
substring	String	**int**, **int**	Returns a new string that is a substring of this string. The substring begins at the first given index and extends to the character at the second index minus one.
toLowerCase	String		Returns a new string which is the same as this one except that all upper case letters are replaced with their corresponding lower case letter.
toUpperCase	String		Returns a new string which is the same as this one except that all lower case letters are replaced with their corresponding upper case letter.

8.32 Standard API: String: format() (page 301)

The standard **class** `java.lang.String` has a **class method** to produce formatted `String` representations of values. It is called `format` and was introduced in Java 5.0. It works with a **format specifier** string in precisely the same way as `System.out.printf()` except that the result is **return**ed rather than printed.

For example, the code

DRAFT April 2011

```
System.out.println(String.format("The distance between %s and %s is %1.2f.",
                                 p1, p2, p1.distanceFromPoint(p2)));
```

has precisely the same effect as the following. (Observe the `%n`.)

```
System.out.printf("The distance between %s and %s is %1.2f.%n",
                  p1, p2, p1.distanceFromPoint(p2));
```

8.33 Standard API: String: split() (page 313)

One of the many **instance method**s in the standard **class** `java.lang.String` is called `split`. It **return**s an **array** of `Strings` in which each **array element** is a portion of the `String` to which the instance method belongs. How the string is split into portions depends on the **method argument** given to `split()`. This argument is another `String` containing a **regular expression** describing what separates the portions.

Here are some examples.

String and regular expression	Resulting array
`"The-cat-sat-on-the-mat".split("-")`	`{ "The", "cat", "sat", "on", "the", "mat" }`
`"The--cat--sat--on--the--mat".split("-")`	`{ "The", "", "cat", "", "sat", "", "on", "", "the", "", "mat" }`
`"The--cat--sat--on--the--mat".split("-+")`	`{ "The", "cat", "sat", "on", "the", "mat" }`
`"The-cat--sat---on----the--mat".split("-+")`	`{ "The", "cat", "sat", "on", "the", "mat" }`

In the last two examples, the regular expression `"-+"` means "one or more hyphens".

8.34 Standard API: String: implements Comparable (page 520)

The standard **class** `java.lang.String` **implement**s `java.lang.Comparable`, with an implementation of `compareTo()` which provides a **lexicographic order**ing. This means it orders the strings in dictionary order based on the values of the **character**s in them.

Since Java 5.0, when `Comparable` became a **generic interface**, `String` actually implements `Comparable<String>`.

DRAFT April 15, 2011

```
public final class String implements Comparable<String>
{
  ...
  @Override
  public int compareTo(String other)
  {
    ...
  } // compareTo
  ...
} // class String
```

8.35 Standard API: Character (page 342)

The standard **class** `java.lang.Character` contains many **class methods** to help with manipulation of **characters**, including the following.

Public method interfaces for class `Character` (some of them).			
Method	**Return**	**Arguments**	**Description**
`isWhitespace`	**`boolean`**	**`char`**	Returns **`true`** if the given **`char`** is a **white space** character, (e.g. **space character**, **tab character**, **new line character**), or **`false`** otherwise.
`isDigit`	**`boolean`**	**`char`**	Returns **`true`** if the given **`char`** is a digit (e.g. `'0'`, `'8'`), or **`false`** otherwise.
`isLetter`	**`boolean`**	**`char`**	Returns **`true`** if the given **`char`** is a letter (e.g. `'A'`, `'a'`), or **`false`** otherwise.
`isLetterOrDigit`	**`boolean`**	**`char`**	Returns **`true`** if the given **`char`** is a letter or a digit, or **`false`** otherwise.
`isLowerCase`	**`boolean`**	**`char`**	Returns **`true`** if the given **`char`** is a lower case letter, or **`false`** otherwise.
`isUpperCase`	**`boolean`**	**`char`**	Returns **`true`** if the given **`char`** is an upper case letter, or **`false`** otherwise.
`toLowerCase`	**`char`**	**`char`**	Returns the lower case equivalent of the given **`char`** if it is an upper case letter, or the given **`char`** if it is not.[3]
`toUpperCase`	**`char`**	**`char`**	Returns the upper case equivalent of the given **`char`** if it is a lower case letter, or the given **`char`** if it is not.[1]

[3]For maximum portability of code to different regions of the world, it is better to use the `String` versions of these methods.

DRAFT April 15, 2011

8.36 Standard API: Object (page 422)

All **object**s in Java are also **instance**s of the standard **class** called `java.lang.Object`. Unless a class is explicitly declared to **extend** some other class, then it implicitly extends `Object` directly. This means all classes in Java reside in a single **inheritance hierarchy**, which is a tree structure with the class `Object` at its root. Every class has a **superclass**, except for the class `Object`.

The `Object` class has one **constructor method** and it takes no **method argument**s.

```
public class Object
{
  ...
  public Object()
  {
    ... Code here to actually create an object,
    ... allocating memory for it, etc..
  } // Object
  ...
} // class Object
```

8.37 Standard API: Object: toString() (page 427)

The **class** `java.lang.Object` has a `toString()` **instance method**. This produces a `String` consisting of (a representation of) the **type** of the **object** followed by a `'@'` and a **hexadecimal** (i.e. base 16) number which is (by default) unique to the object. Classes which do not provide their own version **inherit** this default one.

8.38 Standard API: Object: equals() (page 521)

The standard **class** `java.lang.Object` contains an **instance method** `equals()` which is designed to model the notion of **equivalence** between two **object**s. The definition is as follows.

```
public boolean equals(Object other)
{
  return this == other;
} // equals
```

This is **inherit**ed by all other classes, and so by default all **object**s have this *finest* notion of equivalence: two objects are **equivalent** if and only if they are **equal**, i.e. are the same object. This is often too fine, and so, many classes **override** this definition with one which models the appropriate notion of equivalence for that particular class.

DRAFT April 15, 2011

8.39 Standard API: Object: hashCode() (page 548)

Every object has an **instance method** called hashCode, defined in the java.lang.Object **class**, which is designed to help with classes that use a **hash table** to store **object**s, such as java.util.HashSet. The definition in Object is such that distinct objects have a distinct **hash code**, (usually) based on the memory address of the **reference** at **run time**. Classes that **override** equals() should really also override hashCode() with one that **return**s the same hash code for objects that are **equivalent**, rather than distinct, and yet tend to return a different code for those that are not equivalent. This is so that they will work properly if needed to be used as elements of a HashSet, etc..

```
MyClass v1 = new MyClass(...);
MyClass v2 = new MyClass(...);

if (v1.equals(v2) && v1.hashCode() != v2.hashCode())
  System.out.println("Your hash tables will not work!");
else if (! v1.equals(v2) && v1.hashCode() == v2.hashCode())
  System.out.println("Your hash tables may operate slowly.");
```

8.40 Standard API: Object: hashCode(): making a good definition (page 566)

Classes that **override** equals() ought to also override hashCode() with one that **return**s the same value for **equivalent object**s. Thus, the **function** should be based on the same **instance variable**s that are used to define **equivalence** in equals(). A good **hash code** function should tend to give different hash codes for objects that are not equivalent, otherwise **hash table**s that use them will have too many clashes. One way of achieving a good spread of numbers in a hash code, is to turn these instance variables into numbers, if they are not already so (e.g. by using their own hashCode()) and multiply each by a different, arbitrarily chosen, **prime number**, before adding the products together.

8.41 Standard API: Arrays (page 518)

The standard **class** java.util.Arrays provides various **class method**s to perform complex manipulations of **array**s.

8.42 Standard API: Arrays: sort() (page 518)

One of the **class method**s in java.util.Arrays is called sort, and it takes an **array** of Objects which it **sort**s into their **natural ordering**. For this to work without **throw**ing an **ex-**

DRAFT, April 15, 2011

ception, the items in the array must all be of **type** `Comparable` and be **mutually comparable**. The **algorithm** used is called **merge sort**. This is much more efficient than **bubble sort**.

In fact, the **class** also contains several more class methods called `sort`, one for each array of a **primitive type**, such as `int[]`, etc.. There is even a second version for each type which takes three **method parameter**s: an array, and a pair of `int` indices, $from$ and to. These sort all the items in the array which have an **array index** $\geq from$ and $< to$. This enables **partially filled array**s to be sorted by making $from = 0$ and $to =$ the number of **array element**s used.

8.43 Standard API: Arrays: copyOf() (page 523)

The standard **class** `java.util.Arrays` provides, since Java 6.0, another **class method** called `copyOf` which makes a copy of an **array**. It is a **generic method** and so can handle any kind of **reference type** array. The **new** array returned can be bigger or smaller than the original, and the **array element**s will be the same as in the original for the **array index** positions they have in common.

```
public static <T> T[] copyOf(T[] original, int newLength)
{
  T[] result = ... make a new array of length newLength,
               ... where result[i] = original[i]
               ...   for all 0 <= i < min(original.length, newLength)
  return result;
} // copyOf
```

The single **type parameter**, `T`, specifies the **type** of the array elements, and the two **method parameter**s are the original array of type `T[]`, and an `int` required length for the copy. It **return**s a new array of type `T[]`. (The method uses reflection – an advanced topic not covered by this book – to get around the restrictions on the use of type parameters.)

In fact, the class also contains several more class methods called `copyOf`, one for each array of a **primitive type**, such as `int[]`, etc..

These methods are particularly useful for **array extension**.

```
SomeType[] myArray = new SomeType[INITIAL_SIZE];
...
if ... myArray is now full and I need more room
  myArray = Arrays.copyOf(myArray, myArray.length * RESIZE_FACTOR);
...
```

DRAFT April 15, 2011

8.44 Standard API: Comparable interface (page 520)

The standard **interface** `java.lang.Comparable` provides a **type** for **object**s which can be compared with similar items. Having one type for this notion enables general **algorithm**s to be implemented, such as ones for **sort**ing and efficient searching of **array**s. It was introduced in Java 1.2, but at Java 5.0 it became a **generic interface**. It contains just one **instance method** definition.

```
public interface Comparable<T>
{
  int compareTo(T o);
} // Comparable
```

Any non-**abstract class** that **implement**s this interface, must contain a **method implementation** of `compareTo()` which provides a **total order** for its objects. The **type parameter**, `T`, is for the **type** of objects that can be compared and **class**es that (directly) implement `Comparable` are intended to supply their own class name as the **type argument**. For example, if we say that class `SomeClass` implements `Comparable<SomeClass>` we are stating that `SomeClass` provides an **instance method** `compareTo()`, enabling a `SomeClass` object to compare itself with a given other one.

If a class implements `Comparable`, then the order defined by `compareTo()` is known as the **natural ordering** of that class.

8.45 Standard API: Comparable interface: compareTo() and equals() (page 522)

A **class** that **implement**s `java.lang.Comparable` should have a **method implementation** of `compareTo()` which is consistent with `equals()`, wherever this is possible. By consistent, we mean that

```
x.equals(y)
```

always gives the same value as the following.

```
x.compareTo(y) == 0
```

DRAFT April 15, 2011

9 Statement

9.1 Statement (page 18)

A command in a programming language, such as Java, which makes the computer perform a task is known as a **statement**. `System.out.println("I will output whatever I am told to")` is an example of a statement.

9.2 Statement: simple statements are ended with a semi-colon (page 18)

All simple **statement**s in Java must be ended by a semi-colon (`;`). This is a rule of the Java language **syntax**.

9.3 Statement: assignment statement (page 37)

An **assignment statement** is a Java **statement** which is used to give a value to a **variable**, or change its existing value. This is only allowed if the value we are assigning has a **type** which matches the type of the variable.

9.4 Statement: assignment statement: assigning a literal value (page 37)

We can assign a **literal value**, that is a constant, to a **variable** using an **assignment statement** such as the following.

```
noOfPeopleLivingInMyStreet = 47;
```

We use a single **equal sign** (`=`), with the name of the variable to the left of it, and the value we wish it to be given on the right. In the above example, the **integer literal** `47` will be placed into the variable `noOfPeopleLivingInMyStreet`. Assuming the variable was declared as an **int variable** then this assignment would be allowed because 47 is an **`int`**.

9.5 Statement: assignment statement: assigning an expression value (page 38)

More generally than just assigning a **literal value**, we can use an **assignment statement** to assign the value of an **expression** to a **variable**. For example, assuming we have the variable

DRAFT April 15, 2011

```
int noOfPeopleToInviteToTheStreetParty;
```

then the code

```
noOfPeopleToInviteToTheStreetParty = noOfPeopleLivingInMyStreet + 4;
```

when **execute**d, would **evaluate** the expression on the right of the **equal sign** (=) and then place the resulting value in the variable `noOfPeopleToInviteToTheStreetParty`.

9.6 Statement: assignment statement: updating a variable (page 70)

Java **variable**s have a name and a value, and this value can change. For example, the following code is one way of working out the maximum of two numbers.

```
int x;
int y;
int z;
... Code here that gives values to x, y and z.

int maximumOfXYandZ = x;
if (maximumOfXYandZ < y)
  maximumOfXYandZ = y;
if (maximumOfXYandZ < z)
  maximumOfXYandZ = z;
```

See that the variable `maximumOfXYandZ` is given a value which then might get changed, so that after the end of the second **if statement** it holds the correct value.

A very common thing we want the computer to do, typically inside a **loop**, is to perform a **variable update**. This is when a variable has its value changed to a new value which is based on its current one. For example, the code

```
count = count + 1;
```

will add one to the value of the variable `count`. Such examples remind us that an **assignment statement** is *not* a definition of **equality**, despite Java's use of the single **equal sign**!

DRAFT April 15, 2011

9.7 Statement: assignment statement: updating a variable: shorthand operators (page 87)

The need to undertake a **variable update** is so common, that Java provides various **shorthand operator**s for certain types of update.

Here are some of the most commonly used ones.

Operator	Name	Example	Longhand meaning
`++`	postfix increment	`x++`	`x = x + 1`
`--`	postfix decrement	`x--`	`x = x - 1`
`+=`	compound assignment: add to	`x += y`	`x = x + y`
`-=`	compound assignment: subtract from	`x -= y`	`x = x - y`
`*=`	compound assignment: multiply by	`x *= y`	`x = x * y`
`/=`	compound assignment: divide by	`x /= y`	`x = x / y`

The point of these **postfix increment**, **postfix decrement** and **compound assignment** operators is not so much to save typing when a program is being written, but to make the program easier to read. Once you are familiar with them, you will benefit from the shorter and more obvious code.

There is also a historical motivation. In the early days of the programming language C, from which Java inherits much of its **syntax**, these shorthand **operator**s caused the **compiler** to produce more efficient code than their longhand counterparts. The modern Java compiler with the latest optimization technology should remove this concern.

9.8 Statement: assignment statement: is an expression (page 450)

In Java, the **assignment statement** is actually an **expression**. The `=` symbol is an **operator**, which takes a **variable** as its left **operand**, and an expression as its right operand. It evaluates the expression, assigns it to the variable, *and then* yields the value of the expression as its result.

This allows us to write *horrible* code, such as the following.

```
int x = 10, y = 20, z;

int result = (z = x * y) + (y = z * 2);
```

This is an example of the more general idea of **side effect expression**s – expressions that change the value of some variables while they are being **evaluate**d. Generally speaking, side effect expressions are bad idea, as their use leads to code that is difficult to understand and hence maintain – as the above example illustrates!

DRAFT April 15, 2011

However, there are a few appropriate uses of treating assignment statements as expressions. One is when we wish to assign the same value to a number of variables in one go.

```
x = y = z = 10;
```

Unlike most operators, = has **right associativity**, which means the above example is the same as

```
x = (y = (z = 10));
```

and so makes sense. However, situations where we wish to give several variables the same value at once are not actually very common.

9.9 Statement: if else statement (page 60)

The **if else statement** is one way in Java of having **conditional execution**. It essentially consists of three parts: a **condition** or **boolean expression**, a **statement** which will be **execute**d when the condition is `true` (the **true part**), and another statement which will be executed when the condition is `false` (the **false part**). The whole statement starts with the **reserved word** `if`. This is followed by the condition, written in brackets. Next comes the statement for the true part, then the reserved word `else` and finally the statement for the false part.

For example, assuming we have the **variable** `noOfPeopleToInviteToTheStreetParty` containing the number suggested by its name, then the code

```
if (noOfPeopleToInviteToTheStreetParty > 100)
  System.out.println("We will need a big sound system!");
else
  System.out.println("We should be okay with a normal HiFi.");
```

will cause the computer to compare the current value of `noOfPeopleToInviteToTheStreetParty` with the number `100`, and if it is greater then print out the message `We will need a big sound system!` or otherwise print out the message `We should be okay with a normal HiFi.` – it will never print out both messages. Notice the brackets around the condition and the semi-colons at the end of the two statements inside the if else statement. Notice also the way we lay out the code to make it easy to read, splitting the lines at sensible places and adding more **indentation** at the start of the two inner statements.

DRAFT April 15, 2011

9.10 Statement: if else statement: nested (page 62)

The **true part** or **false part** statements inside an **if else statement** may be any valid Java **statement**, including other if else statements. When we place an if else statement inside another, we say they are **nested**.

For example, study the following code.

```
if (noOfPeopleToInviteToTheStreetParty > 300)
  System.out.println("We will need a Mega master 500 Watt amplifier!");
else
  if (noOfPeopleToInviteToTheStreetParty > 100)
    System.out.println("We will need a Maxi Master 150 Watt amplifier!");
  else
    System.out.println("We should be okay with a normal HiFi.");
```

Depending on the value of `noOfPeopleToInviteToTheStreetParty`, this will report one of *three* messages. Notice the way we have laid out the code above – this is following the usual rules that inner statements have more **indentation** than those they are contained in, so the second if else statement has more spaces because it lives inside the first one. However, typically we make an exception to this rule for if else statements nested in the false part of another, and we would actually lay out the code as follows.

```
if (noOfPeopleToInviteToTheStreetParty > 300)
  System.out.println("We will need a Mega master 500 Watt amplifier!");
else if (noOfPeopleToInviteToTheStreetParty > 100)
  System.out.println("We will need a Maxi Master 150 Watt amplifier!");
else
  System.out.println("We should be okay with a normal HiFi.");
```

This layout reflects our *abstract* thinking that the collection of statements is *one* construct offering three choices, even though it is implemented using two if else statements. This idea extends to cases where we want many choices, using many nested if else statements, without the indentation having to increase for each choice.

9.11 Statement: if statement (page 64)

Sometimes we want the computer to **execute** some code depending on a **condition**, but do nothing if the condition is `false`. We could implement this using an **if else statement** with an empty **false part**. For example, consider the following code.

DRAFT April 15, 2011

```
if (noOfPeopleToInviteToTheStreetParty > 500)
  System.out.println("You may need an entertainment license!");
else ;
```

This will print the message if the **variable** has a value **greater than** 500, or otherwise execute the **empty statement** between the **reserved word** `else` and the semi-colon. Such empty statements do nothing, as you would probably expect!

It is quite common to wish nothing to be done when the condition is `false`, and so Java offers us the **if statement**. This is similar to the if else statement, except it simply does not have the word `else`, nor a false part.

```
if (noOfPeopleToInviteToTheStreetParty > 500)
  System.out.println("You may need an entertainment license!");
```

9.12 Statement: compound statement (page 66)

The Java **compound statement** is simply a list of any number of **statement**s between an opening left brace ({) and a closing right brace (}). You could think of the body of a **method**, e.g. `main()`, as being a compound statement if that is helpful. The meaning is straightforward: when the computer **execute**s a compound statement, it merely executes each statement inside it, in turn. More precisely of course, the Java **compiler** turns the **source code** into **byte code** that has this effect when the **virtual machine** executes the **compile**d program.

We can have a compound statement wherever we can have any kind of statement, but it is most useful when combined with statements which have another statement within them, such as **if else statement**s and **if statement**s.

For example, the following code reports three messages when the **variable** has a value **greater than** 500.

```
if (noOfPeopleToInviteToTheStreetParty > 500)
{
  System.out.println("You may need an entertainment license!");
  System.out.println("Also hire some street cleaners for the next day?");
  System.out.println("You should consider a bulk discount on lemonade!");
}
```

When the **condition** of the if statement is `true`, the body of the if statement is executed. This single statement is itself a compound statement, and so the three statements within it are executed. It is for this sort of purpose that the compound statement exists.

DRAFT April 15, 2011

Note how we lay out the compound statement, with the opening brace at the same **indentation** as the if statement, the statements within it having extra indentation, and the closing brace lining up with the opening one.

Less usefully, a compound statement can be empty, as in the following example.

```
if (noOfPeopleToInviteToTheStreetParty > 500)
{
  System.out.println("You may need an entertainment license!");
  System.out.println("Also hire some street cleaners for the next day?");
  System.out.println("You should consider a bulk discount on lemonade!");
}
else {}
```

As you might expect, the meaning of an empty compound statement is the same as the meaning of an **empty statement**!

9.13 Statement: while loop (page 71)

The **while loop** is one way in Java of having **repeated execution**. It essentially consists of two parts: a **condition**, and a **statement** which will be **execute**d repeatedly while the condition is `true`. The whole statement starts with the **reserved word** `while`. This is followed by the condition, written in brackets. Next comes the statement to be repeated, known as the **loop body**.

For example, the following code is a long winded and inefficient way of giving the **variable** `x` the value `21`.

```
int x = 1;
while (x < 20)
  x = x + 2;
```

The variable starts off with the value `1`, and then repeatedly has `2` added to it, until it is no longer **less than** `20`. This is when the **loop** ends, and `x` will have the value `21`.

Notice the brackets around the condition and the semi-colon at the end of the statement inside the loop. Notice also the way we lay out the code to make it easy to read, splitting the lines at sensible places and adding more **indentation** at the start of the inner statement.

Observe the similarity between the while loop and the **if statement** – the *only* difference in **syntax** is the first word. There is a similarity in meaning too: the while loop executes its body zero or *more* times, whereas the if statement executes its body zero or *one* time. However, **if statement**s are *not* loops and you should avoid the common novice phrase "if loop" when referring to them!

DRAFT April 5, 2011

9.14 Statement: for loop (page 77)

Another kind of **loop** in Java is the **for loop**, which is best suited for situations when the number of **iteration**s of the **loop body** is known before the loop starts. We shall describe it using the following simple example.

```
for (int count = 1; count <= 10; count = count + 1)
  System.out.println("Counting " + count);
```

The **statement** starts with the **reserved word** `for`, which is followed by three items in brackets, separated by semi-colons. Then comes the loop body, which is a single statement (often a **compound statement** of course). The first of the three items in brackets is a **for initialization**, which is performed once just before the loop starts. Typically this involves declaring a **variable** and giving an initial value to it, as in the above example `int count = 1`. The second item is the **condition** for continuing the loop – the loop will only **execute** and will continue to execute while that condition is `true`. In the example above the condition is `count <= 10`. Finally, the third item, a **for update**, is a statement which is executed at the *end* of each iteration of the loop, that is *after* the loop body has been executed. This is typically used to change the value of the variable declared in the first item, as in our example `count = count + 1`.

So the overall effect of our simple example is: declare `count` and set its value to `1`, check that it is **less than** `10`, print out `Counting 1`, add one to `count`, check again, print out `Counting 2`, add one to `count`, check again, and so on until the condition is `false` when the value of `count` has reached `11`.

We do not really need the for loop, as the **while loop** is sufficient. For example, the code above could have been written as follows.

```
int count = 1;
while (count <= 10)
{
  System.out.println("Counting " + count);
  count = count + 1;
}
```

However you will see that the for loop version has placed together all the code associated with the control of the loop, making it easier to read, as well as a little shorter.

There is one very subtle difference between the for loop and while loop versions of the example above, concerning the **scope** of the variable `count`, that is the area of code in which the variable can be used. Variables declared in the initialization part of a for loop can only be used in the for loop – they do not exist elsewhere. This is an added benefit of using for loops when appropriate: the variable, which is used solely to control the loop, cannot be accidentally used in the rest of the code.

DRAFT April 15, 2011

9.15 Statement: for loop: multiple statements in for update (page 136)

Java **for loop**s are permitted to have more than one **statement** in their **for update**, that is, the part which is **execute**d after the **loop body**. Rather than always being one statement, this part may be a list of statements with commas (,) between them.

One appropriate use for this feature is to have a for loop that executes twice, once each for the two possible values of a **boolean variable**.

For example, the following code prints out scenarios to help train people to live in the city of Manchester!

```
boolean isRaining = true;
boolean haveUmbrella = true;
for (int countU = 1; countU <= 2; countU++, haveUmbrella = !haveUmbrella)
  for (int countR = 1; countR <= 2; countR++, isRaining = !isRaining)
  {
    System.out.println("It is" + (isRaining ? "" : " not") + " raining.");
    System.out.println
       ("You have " + (haveUmbrella ? "an" : "no") + " umbrella.");
    if (isRaining && !haveUmbrella)
      System.out.println("You get wet!");
    else
      System.out.println("You stay dry.");
    System.out.println();
  } // for
```

9.16 Statement: statements can be nested within each other (page 92)

Statements that control execution flow, such as **loop**s and **if else statement**s have other **statement**s inside them. These inner statements can be any kind of statement, including those that control the flow of execution. This allows quite complex **algorithm**s to be constructed with unlimited nesting of different and same kinds of control statements.

For example, one simple (but inefficient) way to print out the non-negative multiples of x which lie between y (≥ 0) and z inclusive, is as follows.

```
for (int number = 0; number <= z; number += x)
  if (number >= y)
    System.out.println("A multiple of " + x + " between " + y
                       + "and " + z + " is " + number);
```

DRAFT April 15, 2011

9.17 Statement: switch statement with breaks (page 107)

Java provides a **conditional execution statement** which is ideal for situations where there are many choices based on some value, such as a number, being **equal** to specific fixed values for each choice. It is called the **switch statement**. The following example code will applaud the user when they have correctly guessed the winning number of `100`, encourage them when they are one out, or insult them otherwise.

```
int userGuess = Integer.parseInt(args[0]);

switch (userGuess)
{
  case 99: case 101:
    System.out.println("You are close!");
    break;
  case 100:
    System.out.println("Bingo! You win!");
    System.out.println("You have guessed correctly.");
    break;
  default:
    System.out.println("You are pathetic!");
    System.out.println("Have another guess.");
    break;
} // switch
```

The switch statement starts with the **reserved word** **`switch`** followed by a bracketed **expression** of a **type** that has discrete values, such as **`int`** (notably not **`double`**). The body of the statement is enclosed in braces, ({ and }), and consists of a list of entries. Each of these starts with a list of labels, comprising the reserved word **`case`** followed by a value and then a colon (`:`). After the labels we have one or more statements, typically ending with a **break statement**. One (at most) label is allowed to be the reserved word **`default`** followed by a colon – usually written at the end of the list.

When a switch statement is **execute**d, the expression is **evaluate**d and then each label in the body is examined in turn to find one whose value is equal to that of the expression. If such a match is found, the statements associated with that label are executed, down to the special **break statement** which causes the execution of the switch statement to end. If a match is not found, then instead the statements associated with the **`default`** label are executed, or if there is no **`default`** then nothing is done.

9.18 Statement: switch statement without breaks (page 110)

A less common form of the **switch statement** is when we omit the **break statement**s at the end of the list of **statement**s associated with each set of **`case`** labels. This, perhaps surprisingly,

DRAFT April 15, 2011

causes execution to "fall through" to the statements associated with the next set of `case` labels. Most of the time we do *not* want this to happen – so we have to be careful to remember the break statements.

We can also mix the styles – having break statements for some entries, and not for some others. The following code is a bizarre, but interesting way of doing something reasonably simple. It serves as an illustration of the switch statement, and as a puzzle for you. It takes two **integer**s, the second of which is meant to be in the range one to ten, and outputs a result which is some **function** of the two numbers. What is that result?

```
int value = Integer.parseInt(args[0]);
int power = Integer.parseInt(args[1]);

int valueToThePower1 = value;
int valueToThePower2 = valueToThePower1 * valueToThePower1;
int valueToThePower4 = valueToThePower2 * valueToThePower2;
int valueToThePower8 = valueToThePower4 * valueToThePower4;

int result = 1;

switch (power)
{
  case 10: result *= valueToThePower1;
  case 9:  result *= valueToThePower1;
  case 8:  result *= valueToThePower8;
           break;
  case 7:  result *= valueToThePower1;
  case 6:  result *= valueToThePower1;
  case 5:  result *= valueToThePower1;
  case 4:  result *= valueToThePower4;
           break;
  case 3:  result *= valueToThePower1;
  case 2:  result *= valueToThePower2;
           break;
  case 1:  result *= valueToThePower1;
           break;
} // switch

System.out.println(result);
```

If you find the semantics of the switch statement somewhat inelegant, then do not worry – you are not alone! Java inherited it from C, where it was designed more to ease the work of the **compiler** than to be a good construct for the programmer. You will find the switch statement is less commonly used than the **if else statement**, and the majority of times you use it, you will want to have break statements on every set of `case` labels. Unfortunately, due to them being optional, accidentally missing them off does not cause a **compile time error**.

DRAFT April 15, 2011

9.19 Statement: do while loop (page 112)

The **do while loop** is the third way in Java of having **repeated execution**. It is similar to the **while loop** but instead of having the **condition** at the start of the **loop**, it appears at the end. This means the condition is **evaluate**d *after* the **loop body** is **execute**d rather than before. The whole **statement** starts with the **reserved word** `do`. This is followed by the statement to be repeated, then the reserved word `while` and finally the condition, written in brackets.

For example, the following code is a long winded and inefficient way of giving the **variable** `x` the value `21`.

```
int x = 1;
do
  x += 2;
while (x < 20);
```

Observe the semi-colon that is needed after the condition.

Of course, the body of the do while loop might be a **compound statement**, in which case we might lay out the code as follows.

```
int x = 0;
int y = 100;
do
{
  x++;
  y--;
} while (x != y);
```

The above is a long winded and inefficient way of giving both the variables `x` and `y` the value `50`.

Note that, because the condition is evaluated *after* the body is executed, the body is executed at least once. This is in contrast to the while loop, which might have have its body executed zero times.

9.20 Statement: for-each loop: on arrays (page 293)

Java 5.0 introduced a new **statement** called the **enhanced for statement**, more commonly known as the **for-each loop**.[4]

[4] The popular name for this loop may seem odd, because the word `each` is not used in it, but the meaning of the statement is similar to a concept in languages such as Perl[17], which does use the phrase `for each`. And we actually say ‘for each’ when we read out the Java statement.

DRAFT April 15, 2011

It is best explained by example. Suppose we have the following.

```
double[] myFingerLengths = new double[10];

... Code here to assign values to the array elements.
```

Then we can find the sum of the **array element**s with the following for-each loop.

```
double myTotalFingerLength = 0;
for (double fingerLength : myFingerLengths)
  myTotalFingerLength += fingerLength;
```

This is saying that we want to **loop** over all the elements in the **array** which is **reference**d by `myFingerLengths`, storing each element in turn in the variable `fingerLength`, and adding it to the value of `myTotalFingerLength`. In other words 'for each `fingerLength` in `myFingerLengths`, add `fingerLength` to `myTotalFingerLength`'.

The above for-each loop is actually a shorthand for the following **for loop**.

```
double myTotalFingerLength = 0;
for (int index = 0; index < myFingerLengths.length; index++)
{
  double fingerLength = myFingerLengths[index];
  myTotalFingerLength += fingerLength;
} // for
```

Here is the general case of the for-each loop when used with arrays, where `anArray` is a variable referring to some array with **array base type** `SomeType` and `elementName` is any suitable **variable** name.

```
for (SomeType elementName : anArray)
  ... Statement using elementName.
```

This general case is simply a shorthand for the following.

```
for (int index = 0; index < anArray.length; index++)
{
  SomeType elementName = anArray[index];
  ... Statement using elementName.
} // for
```

DRAFT April 15, 2011

A for-each loop can and should be used instead of a for loop in places where we wish to loop over all the elements of a single array, and the **array index** is *only* used to access (not change) the elements of that array. In other words, for processing where the element values matter, but their position in the array is not directly used, and there is only one array. So, for example, the following code cannot be replaced with a for-each loop.

```
int weightedSum = 0;
for (int index = 0; index < numbers.length; index++)
  weightedSum += numbers[index] * index;
```

Neither can this.

```
for (int index = 0; index < numbers.length; index++)
  otherNumbers[index] = numbers[index];
```

Finally, a common error (even in some Java text books!) is to think that a for-each loop can be used to *change* the array elements. For example, the following code **compile**s without errors, but it does not do what you might expect!

```
int[] numbers = new int[100];
for (int number : numbers)
  number = 10;
```

The for-each loop above is a shorthand for the following, which you can see achieves nothing.

```
for (int index = 0; index < numbers.length; index++)
{
  int number = numbers[index];
  number = 10;
} // for
```

9.21 Statement: for-each loop: on collections (page 562)

The **enhanced for statement** introduced in Java 5.0 and more commonly known as the **for-each loop**, can be used with **collection**s as well as **array**s. Suppose we have some `Collection` for which we want to process each of the elements using its `Iterator`.

```
Collection<T> c = ...

Iterator<T> i = c.iterator();
while (i.hasNext())
  ... Statement with one use of i.next().
```

DRAFT April 15, 2011

If, as in the above abstract example, we wish to process all the elements of the collection in the same **loop**, then we can use a for-each loop, as follows.

```
Collection<T> c = ...

for (T e : c)
  ... Statement using e.
```

This is a shorthand for precisely the following **for loop**.

```
Collection<T> c = ...

for (Iterator<T> i = c.iterator(); i.hasNext(); )
{
  T e = i.next();
  ... Statement using e.
} // for
```

As with arrays, the for-each loop is only suitable if we want to process all the elements using the one loop.

9.22 Statement: try statement (page 344)

The **try statement** is used to implement **exception catching** in Java. It uses the **reserved words** **`try`** and **`catch`**, as follows.

```
try
{
  ... Code here that might cause an exception to happen.
} // try
catch (Exception exception)
{
  ... Code here to deal with the exception.
} // catch
```

The **statement** consists of two parts, the **try block** and the **catch clause**. When the try statement is **execute**d, the code inside the try block is obeyed as usual. However, if at some point during this execution an **exception** occurs, an **instance** of `java.lang.Exception` is created, and then control immediately transfers to the catch clause. The newly created `Exception` **object** is available to the code in the catch clause, as an **exception parameter**, which is a bit like

DRAFT April 15, 2011

a **method parameter**. For this reason, we must declare a name (and **type**) for the exception in the round brackets following the reserved word `catch`.

For example, the following **method** computes the mean average of an **array** of `int` values, dealing with the possibility of the **reference** being the **null reference** or the array being an **empty array**, by **catch**ing the exception and **return**ing zero instead.

```
private double average(int[] anArray)
{
  try
  {
    int total = anArray[0];
    for (int i = 1; i < anArray.length; i++)
      total += anArray[i];
    return total / (double) anArray.length;
  } // try
  catch (Exception exception)
  {
    // Report the exception and carry on.
    System.err.println(exception);
    return 0;
  } // catch
} // average
```

Note: unlike most Java **statement**s that may contain other statements, the two parts of the try statement must both be **compound statement**s, even if they only contain one statement!

9.23 Statement: try statement: with multiple catch clauses (page 347)

The **try statement** may have more than one **catch clause**, each of which is designed to **catch** a different kind of **exception**. When an exception occurs in the **try block**, the execution control transfers to the first matching catch clause, if there is one, or continues to propagate out of the try statement if there is not.

For example, consider the following **method** which finds the largest of some numbers stored in an **array** of `String` **object**s.

```
private int maximum(String[] anArray)
{
  try
  {
    int maximumSoFar = Integer.parseInt(anArray[0]);
    for (int i = 1; i < anArray.length; i++)
    {
```

DRAFT April 15, 2011

```
        int thisNumber = Integer.parseInt(anArray[i]);
        if (thisNumber > maximumSoFar)
          maximumSoFar = thisNumber;
      } // for
      return maximumSoFar;
    } // try
    catch(NumberFormatException exception)
    {
      System.err.println("Cannot parse item as an int: "
                         + exception.getMessage());
      return 0;
    } // catch
    catch(ArrayIndexOutOfBoundsException exception)
    {
      System.err.println("There is no maximum, as there are no numbers!");
      return 0;
    } // catch
  } // maximum
```

If the array **reference**d by the **method parameter** is an **empty array**, that is, it has no elements, then an `ArrayIndexOutOfBoundsException` object will be created when the code tries to access the first **array element**. This will be caught by the second catch clause. If, on the other hand, one of the strings in the array does not represent an **`int`** then a `NumberFormatException` object will be created inside the `parseInt()` method, and this will be caught by the first catch clause.

However, if the given **method argument** was actually the **null reference**, that is, there is no array at all – not even an empty one, then a `NullPointerException` object is created when the code tries to follow the array reference to access element zero of it.

```
  int maximumSoFar = Integer.parseInt(anArray[0]);
```

The code `anArray[0]` means "follow the reference in the **variable** `anArray` to the array referenced by it, and then get the value stored at **array index** `0` in that array." In this example there is no catch clause matching a `NullPointerException`, so the execution control transfers out of the try statement altogether, and out of the method. If the **method call** was itself inside the following try statement, then the `NullPointerException` would get caught there.

```
  try
  {
    int max = maximum(null);
    ...
  } // try
  catch (NullPointerException exception)
  {
    System.err.println("Silly me!");
  } // catch
```

DRAFT April 15, 2011

9.24 Statement: try statement: with finally (page 451)

The **try statement** may optionally be given a **finally block**, which is a piece of code that will be **execute**d at the end of the whole try statement, regardless of whether the **try block** successfully completes, or if a **catch clause** is executed, or if control is being **throw**n out of the try statement.

The general form of a **try finally statement** is as follows.

```
try
{
  ... Code here that might cause an exception to happen.
} // try
catch (SomeException exception)
{
  ... Code here to deal with SomeException types of exception.
} // catch
catch (AnotherException exception)
{
  ... Code here to deal with AnotherException types of exception.
} // catch
... more catch clauses as required.
finally
{
  ... Code here that will be run, no matter what,
  ... as the last thing the statement does.
} // finally
```

9.25 Statement: throw statement (page 350)

The **throw statement** is used when we wish our code to trigger the **exception** mechanism of Java. It consists of the **reserved word** `throw`, followed by a **reference** to an `Exception` **object**. When the **statement** is **execute**d, the Java **virtual machine** finds the closest **try statement** that is currently being executed, which has a **catch clause** that matches the kind of exception being thrown, and transfers execution control to that catch clause. If there is no matching catch clause to be found, then the exception is reported and the **thread** is terminated.

For example, here we **throw** an **instance** of the general `java.lang.Exception` **class** without a specific message.

```
throw new Exception();
```

This next one has a message.

DRAFT April 15, 2011

```
throw new Exception("This is the message associated with the exception");
```

And finally, this example is throwing an instance of `java.lang.NumberFormatException` with a message.

```
NumberFormatException exception
  = new NumberFormatException("Only digits please");
throw exception;
```

DRAFT April 15, 2011

10 Error

10.1 Error (page 20)

When we write the **source code** for a Java program, it is very easy for us to get something wrong. In particular, there are lots of rules of the language that our program must obey in order for it to be a valid program.

10.2 Error: syntactic error (page 20)

One kind of error we might make in our programs is **syntactic error**s. This is when we break the **syntax** rules of the language. For example, we might miss out a closing bracket, or insert an extra one, etc.. This is rather like missing out a word in a sentence of natural language, making it grammatically incorrect. The sign below, seen strapped to the back of a poodle, contains bad grammar – it has an `is` missing.

My other dog an Alsatian.

Syntactic errors in Java result in the **compiler** giving us an error message. They can possibly confuse the compiler, resulting in it thinking many more things are wrong too!

10.3 Error: semantic error (page 22)

Another kind of error we might make is a **semantic error**, when we obey the rules of the **syntax** but what we have written does not make any sense – it has no semantics (meaning). Another sign on a different poodle might say

My other dog is a Porsche.

which is senseless because a Porsche is a kind of car, not a dog.

10.4 Error: compile time error (page 22)

Java **syntactic error**s and many **semantic error**s can be detected for us by the **compiler** when it processes our program. Errors that the compiler can detect are called **compile time error**s.

10.5 Error: run time error (page 24)

Another kind of error we can get with programs is **run time error**s. These are errors which are detected when the program is **run** rather than when it is **compile**d. In Java this means the errors are detected and reported by the **virtual machine**, `java`.

Java calls run time errors **exception**s. Unfortunately, the error messages produced by `java` can look very cryptic to novice programmers. A typical one might be as follows.

```
Exception in thread "main" java.lang.NoSuchMethodError: main
```

You can get the best clue to what has caused the error by just looking at the words either side of the colon (`:`). In the above example, the message is saying that `java` cannot find the **method** called `main`.

10.6 Error: logical error (page 29)

The most tricky kind of error we can make in our programs is a **logical error**. For these mistakes we do not get an error message from the **compiler**, nor do we get one at **run time** from the **virtual machine**. These are the kind of errors for which the Java program we have written is meaningful as far as Java is concerned, it is just that our program does the wrong thing compared with what we wanted. There is no way the compiler or virtual machine can help us with these kinds of error: they are far, far too stupid to understand the *problem* we were trying to solve with our program.

For this reason, many logical errors, especially very subtle ones, manage to slip through undetected by human program testing, and end up as **bug**s in the final product – we have all heard stories of computer generated demands for unpaid bills with *negative* amounts, etc..

DRAFT April 15, 2011

11 Execution

11.1 Execution: sequential execution (page 23)

Programs generally consist of more than one **statement**, in a list. We usually place these on separate lines to enhance human readability, although Java does not care about that. Statements in such a list are **execute**d sequentially, one after the other. More correctly, the Java **compiler** turns each one into corresponding **byte code**s, and the **virtual machine** executes each collection of byte codes in turn. This is known as **sequential execution**.

11.2 Execution: conditional execution (page 60)

Having a computer always obey a list of instructions in a certain order is not sufficient to solve many problems. We often need the computer to do some things only under certain circumstances, rather than every time the program is **run**. This is known as **conditional execution**, because we get the computer to **execute** certain instructions **condition**ally, based on the values of the **variable**s in the program.

11.3 Execution: repeated execution (page 70)

Having a computer always obey instructions just once within the **run** of a program is not sufficient to solve many problems. We often need the computer to do some things more than once. In general, we might want some instructions to be **execute**d, zero, one or many times. This is known as **repeated execution**, **iteration**, or **loop**ing. The number of times a loop of instructions is executed will depend on some **condition** involving the **variable**s in the program.

11.4 Execution: parallel execution – threads (page 253)

Computers appear to be able to perform more than one task at the same time. For example, we can run several programs at once and they run in parallel. At the **operating system** level, each program runs in a separate **process**, and the computer shares its **central processing unit** time fairly between the current processes.

The Java **virtual machine** has a built-in notion of processes, called **thread**s, which allows for a single program to be doing more than one thing at a time. When a Java program is started, the virtual machine creates one thread, called the **main thread**, which is set off to **run** the body of the **main method**. This **execute**s the **statement**s in the main method, including the statements of any **method call**s it finds. Upon reaching the end of the main method, this thread terminates, which causes the virtual machine to exit if that was the only thread existing at the

DRAFT April 15, 2011

time. If, however there are any other threads which have not yet terminated, then the virtual machine continues to run them. It exits the program only when all the threads have ended.

11.5 Execution: parallel execution – threads: the GUI event thread (page 254)

When we have a program that places a **graphical user interface** (**GUI**) window on the screen, the Java **virtual machine** creates another **thread**, which we shall call the **GUI event thread**. This is created when the first window of the program is shown. As a result of this, the program does *not* end when the **main thread** reaches the end of the **main method** – this is of course what we want for a program with a GUI.

(In reality, the virtual machine creates several GUI event threads, but it suffices to think of there being just the one.)

The GUI event thread spends most of its life asleep – quietly doing nothing. When the end user of the program does something that might be of interest to the program, the **operating system** informs the virtual machine, which in turn wakes up the GUI event thread. Such interesting things include moving the mouse into, out of, or within a window belonging to the program, pressing a mouse key while the mouse is over such a window, typing a keyboard key while a window of the program has keyboard focus, etc.. These things are collectively known as **event**s.

When it is woken up, the GUI event thread looks to see what might have changed as a result of the end user's action. For example, he or she may have pressed a GUI button belonging to the program. For each event which is definitely interesting, the GUI event thread **execute**s some code which is designed to process that event. Then it goes back to sleep again.

11.6 Execution: event driven programming (page 254)

A large part of writing programs with **graphical user interface**s (**GUI**s) is about constructing the code which will process the **event**s associated with the end user's actions. This is known as **event driven programming**. Essentially, the **main method** sets up the GUI of the program via **method call**s, and then it ends. From then on, the code associated with processing GUI events does all the work – when the end user does things which cause such events to happen. That is, the program becomes driven by the events.

12 Code clarity

12.1 Code clarity: layout (page 31)

Java does not care how we lay our code out, as long as we use some **white space** to separate adjacent symbols that would otherwise be treated as one symbol if they were joined. For example `public void` with no space between the words would be treated as the single symbol `publicvoid` and no doubt cause a **compile time error**. So, if we were crazy, we could write all our program **source code** on one line with the minimum amount of space between symbols!

```
public class HelloSolarSystem{public static void main(String[]args){System.out.println("Hello Mercury!");System.out.println("He
```

Oh dear – it ran off the side of the page (and that was with a smaller font too). Let us split it up into separate lines so that it fits on the page.

```
public class HelloSolarSystem{public static void main(String[]args){
System.out.println("Hello Mercury!");System.out.println(
"Hello Venus!");System.out.println("Hello Earth!");System.out.println
("Hello Mars!");System.out.println("Hello Jupiter!");System.out.
println("Hello Saturn!");System.out.println("Hello Uranus!");System.
out.println("Hello Neptune!");System.out.println("Goodbye Pluto!");}}
```

Believe it or not, this program would still **compile** and **run** okay, but hopefully you will agree that it is not very easy for *us* to read. Layout is very important to the human reader, and programmers must take care and pride in laying out their programs as they are written. So we split our program *sensibly*, rather than arbitrarily, into separate lines, and use **indentation** (i.e. spaces at the start of some lines), to maximize the readability of our code.

12.2 Code clarity: layout: indentation (page 32)

A **class** contains structures **nested** within each other. The outer-most structure is the class itself, consisting of its heading and then containing it's body within the braces. The body contains items such as the **main method**. This in turn consists of a heading and a body contained within braces.

The idea of **indentation** is that the more nested a part of the code is, the more space it has at the start of its lines. So the class itself has no spaces, but its body, within the braces, has two or three. Then the body of the main method has two or three more. You should be consistent: always use the same number of spaces per nesting level. It is also a good idea to avoid using **tab character**s as they can often look okay on your screen, but not line up properly when the code is printed.

DRAFT April 15, 2011

In addition, another rule of thumb is that opening braces ({) should have the same amount of indentation as the matching closing brace (}). You will find that principle being used throughout this book. However, some people prefer a style where opening braces are placed at the end of lines, which this author believes is less clear.

```
public class HelloWorld {

  public static void main(String[] args) {
    System.out.println("Hello world!");
  }
}
```

12.3 Code clarity: layout: splitting long lines (page 43)

One of the features of good layout is to keep our **source code** lines from getting too long. Very long lines cause the reader to have to work harder in horizontal eye movement to scan the code. When code with long lines is viewed on the screen, the reader either has to use a horizontal scroll bar to see them, or make the window so wide that other windows cannot be placed next to it. Worst of all, when code with long lines is printed on paper there is a good chance that the long lines will disappear off the edge of the page! At very least, they will be wrapped onto the next line making the code messy and hard to read.

So a good rule of thumb is to keep your source code lines shorter than 80 **character**s long. You can do this simply in most **text editor**s by never making the text window too wide and never using the horizontal scroll bar while writing the code.

When we do have a **statement** that is quite long, we simply split it into separate lines at carefully chosen places. When we choose such places, we bear in mind that most human readers scan down the left hand side of the code lines, rather than read every word. So, if a line is a continuation of a previous line, it is important to make this obvious at the start of it. This means using an appropriate amount of **indentation**, and choosing the split so that the first symbol on the continued line is not one which could normally start a statement.

A little thought at the writing stage quickly leads to a habit of good practise which seriously reduces the effort required to read programs once they are written. Due to **bug** fixing and general maintenance over the lifetime of a real program, the code is read many more times than it is written!

12.4 Code clarity: comments (page 82)

In addition to having careful layout and **indentation** in our programs, we can also enhance human readability by using **comment**s. These are pieces of text which are ignored by the **compiler**, but help describe to the human reader what the program does and how it works.

DRAFT April 15, 2011

For example, every program should have comments at the start saying what it does and briefly how it is used. Also, **variable**s can often benefit from a comment before their declaration explaining what they are used for. As appropriate, there should be comments in the code too, *before* certain parts of it, explaining what these next **statement**s are going to do.

One form of comment in Java starts with the symbol //. The rest of that source line is then the text of the comment. For example

```
// This is a comment, ignored by the compiler.
```

12.5 Code clarity: comments: marking ends of code constructs (page 83)

Another good use of **comment**s is to mark every closing brace (}) with a comment saying what code construct it is ending. The following skeleton example code illustrates this.

```
public class SomeClass
{
  public static void main(String[] args)
  {
    ...
    while (...)
    {
      ...
      ...
      ...
    } // while
    ...
  } // main

} // class SomeClass
```

12.6 Code clarity: comments: multi-line comments (page 189)

Another form of **comment** in Java allows us to have text which spans several lines. These start with the symbol /* and end with the symbol */, which typically will be several lines later in the code. These symbols, and all text between them, is ignored by the **compiler**.

Less usefully, we can have the start and end symbols on the same line, with program code on either side of the comment, if we wish.

DRAFT April 15, 2011

13 Design

13.1 Design: hard coding (page 36)

Programs typically process input **data**, and produce output data. The input data might be given as **command line argument**s, or it might be supplied by the user through some **user interface** such as a **graphical user interface** or **GUI**. It might be obtained from **file**s stored on the computer.

Sometimes input data might be built into the program. Such data is said to be **hard coded**. This can be quite common while we are developing a program and we haven't yet written the code that obtains the data from the appropriate place. In other cases it might be appropriate to have it hard coded in the final version of the program, if such data only rarely changes.

13.2 Design: pseudo code (page 73)

As our programs get a little more complex, it becomes hard to write them straight into the **text editor**. Instead we need to **design** them *before* we implement them.

We do not design programs by starting at the first word and ending at the last, like we do when we implement them. Instead we can start wherever it suits us – typically at the trickiest bit.

Neither do we express our designs in Java – that would be a bad thing to do, as Java forces our mind to be cluttered with trivia which, although essential in the final code, is distracting during the design.

Instead, we express our **algorithm** designs in **pseudo code**, which is a kind of informal programming language that has all unnecessary trivia ignored. So, for example, we do not bother writing the semi-colons at the end of **statement**s, or the brackets round **condition**s etc.. We might not bother writing the **class** heading, nor the **method** heading, if it is obvious to us what we are designing. And so on.

Also, during design in pseudo code, we can vary the level of **abstraction** to suit us – we do not have to be constrained to use only the features that are available in Java.

13.3 Design: object oriented design (page 184)

When we are developing programs in an **object oriented programming** language, such as Java, we should use the principle of **object oriented design**. We start by identifying the **class**es we shall have in the program, by examining the **requirements statement** of the problem which

the program is to solve. This is recognizing the idea that problems inherently involve interactions between 'real world' objects. These will be modelled in our program, by it creating **object**s which are **instance**s of the classes we identify.

In this view then, an object is an entity which has some kind of **object state** which might change over time, and some kind of **object behaviour** which might be based on its state.

From the requirements, we think carefully about the state and the behaviour of the objects in the problem. Then we decide how to model their behaviour using **instance method**s, and their state using **instance variable**s. There may, in general, be a need for **class variable**s and **class method**s too.

13.4 Design: object oriented design: noun identification (page 185)

One way to analyse the **requirements statement** in order to decide what **class**es to have in the program, is to simply go through the requirements and list all the nouns and noun phrases we can find. This is called **noun identification** and is useful because the objects inherent in the solution to most problems actually appear as nouns in the description of the problem. Some of the nouns will relate to **object**s that will exist at **run time**, and some will relate to classes in the program.

It is not the case that every noun found will be a class or an object, of course, and sometimes we need classes that do not appear as nouns in the requirements. However, the technique is usually a good way of starting the process.

13.5 Design: object oriented design: encapsulation (page 187)

An important principle in **object oriented design** is the idea of **encapsulation**. A well designed **class** encapsulates the behaviour of the **object**s that can be created from it, in such a way that in order to use the class, one only needs to know about its **public method**s (including **constructor method**s) and what they mean, rather than how they work and what **instance variable**s the class may have. To help achieve good encapsulation, we follow the principle of **putting the logic where the data is** – all the code pertaining to the behaviour of particular objects are included in their class, rather than sprinkled around the various different classes of the program.

Encapsulation is an instance of **abstraction**. Abstraction is the process of ignoring detail which is not necessary for us to know (at the moment). We can use a class without having to know how it works, for example, if it is written by somebody else. Or, we can **design** the details of one class at a time for our programs, without at that moment being concerned with the details of how the other classes work.

For an example which has little to do with Java, assume you have just bought a cheap DVD TV

DRAFT April 15, 2011

recorder from your local supermarket. Do you need to know how it works in order to use it? Do you need to remove the case lid in order to use it? No, you only need to know about the buttons on the *outside* of the case. That is, until it breaks (after all it was a cheap one). Only at that point do you, or perhaps better still a TV gadget engineer, need to remove the case and poke around inside.

13.6 Design: Sorting a list (page 295)

A **list** of items, such as an **array**, contains those items in some, perhaps arbitrary, order. We often want to rearrange them into a *specific* order, without losing or gaining any. This is known as **sort**ing. For example, a list of numbers may be sorted into ascending or descending numerical order, a list of names may be sorted alphabetically, etc..

There are many different **algorithm**s for sorting lists, including **bubble sort**, **insertion sort**, **selection sort**, **quick sort**, **merge sort**, **tree sort** . . . .

13.7 Design: Sorting a list: bubble sort (page 296)

One **algorithm** for **sort**ing is known as **bubble sort**. This works by passing through the **list** looking at adjacent items, and swapping them over if they are in the wrong order. One pass through is not enough to ensure the list gets completely sorted, so more passes must be made until it is. However, after the first pass, the 'highest' item, that is, the one that should end up being furthest from the start of the list, must actually be at the end of the list.

For example suppose we start with the following list and wish to sort it into ascending order.

45	78	12	79	60	17

On the first pass, we compare 45 with 78, which are in order, and then 78 with 12 which need swapping. Next we compare 78 with 79, and so on. Eventually we end up with 79 being at the end of the list.

Start		45	78	12	79	60	17
45 <= 78	okay	45 <=	78	12	79	60	17
78 > 12	swap	45	12 <=	78	79	60	17
78 <= 79	okay	45	12	78 <=	79	60	17
79 > 60	swap	45	12	78	60 <=	79	17
79 > 17	swap	45	12	78	60	17 <=	79

The highest number, 79, is in place, but the preceding items are not yet sorted.

DRAFT April 15, 2011

After the second pass, the second highest item must be at the penultimate place in the list, and so on. It follows that, if there are N items in the list, then $N-1$ passes are enough to guarantee the whole list is sorted. Furthermore, the first pass needs to look at $N-1$ adjacent pairs, but the next pass can look at one less, because we know the highest item is in the right place at the end. The very last pass only needs to look at one pair, as all the other items must be in place by then.

Going back to our example, here are the results at the end of the next passes.

Pass						
2	12	45	60	17	78	79
3	12	45	17	60	78	79
4	12	17	45	60	78	79
5	12	17	45	60	78	79

Notice that pass 5 was actually unnecessary as the **array** became sorted after pass 4.

Here is some **pseudo code** for sorting `anArray` using bubble sort.

```
for passCount = 1 to anArray length - 1
  for pairLeftIndex = 0 to anArray length - 1 - passCount
    if items in anArray at pairLeftIndex and pairLeftIndex + 1
                           are out of order
      swap them over
```

This can be improved by observing that the list may get sorted before the maximum number of passes needed to guarantee it. For example it could be sorted to start with! Here is an alternative **design**.

```
int unsortedLength = anArray length
boolean changedOnThisPass
do
  changedOnThisPass = false
  for pairLeftIndex = 0 to unsortedLength - 2
    if items in anArray at pairLeftIndex and pairLeftIndex + 1
                           are out of order
      swap them over
      changedOnThisPass = true
    end-if
  end-for
  unsortedLength--
while changedOnThisPass
```

DRAFT April 15, 2011

13.8 Design: Sorting a list: total order (page 516)

A **total order** over some **data** is a relationship between pairs of that data which enables it to be **sort**ed. For example, **less than or equal** is a total order over numbers: we can choose to sort a **list** of numbers into ascending order. So, **greater than or equal** is also a total order – we can sort into descending order. We can sort strings using a **lexicographic order**ing (i.e. into dictionary order). A child might sort sweets into order by colour – unwittingly defining a total order of colours in the process.

More formally, every total order, $\preceq$, has three properties, for all values x, y and z:

[Antisymmetric:]	if $x \preceq y$ and $y \preceq x$, then $x = y$
[Transitive:]	if $x \preceq y$ and $y \preceq z$, then $x \preceq z$
[Total:]	$x \preceq y$ or $y \preceq x$

One way of modelling a total order is to provide a **function** that takes any pair, (x,y), of the data and yields one of three states as follows.

- x comes before y.
- x and y have the same placing.
- x comes after y.

In Java the function is typically implemented by an **instance method** `compareTo()` which compares the current **instance** (x) with a given other (y), and yields an **`int`** that is negative, zero, or positive to represent the three states respectively.

13.9 Design: Sorting a list: tree sort (page 554)

Another **algorithm** for **sort**ing is known as **tree sort**. In this approach, the items from the **list** are inserted into an **ordered binary tree**, and then the tree is scanned from left to right, that is smallest item to largest (according to the **total order** being used), to produce the result.

If the **data** to be sorted has no duplicates, or if it is desired to not include such multiple elements in the result, then in Java a tree sort can be achieved by using an **instance** of `java.util.TreeSet`. This is because its `iterator()` **instance method** produces an `Iterator` which gives access to the elements in order from smallest to largest. Any duplicate items will not appear in the result because a **set** is not changed by an attempt to add an element which is **equivalent** to one that is already present.

DRAFT April 15, 2011

13.10 Design: Searching a list: linear search (page 323)

The simplest way to find an item in a **list** of items, such as an **array**, is to perform a **linear search** – starting at the front and looking at each item in turn. For example, the following **array search method** finds the position of a given `int` in a given **array**, or **return**s `-1` if the number is not found.

```
private int posOfInt(int[] anArray, int toFind)
{
  int searchPos = 0;
  while (searchPos < anArray.length && anArray[searchPos] != toFind)
    searchPos++;
  if (searchPos == anArray.length) return -1;
  else                              return searchPos;
} // posOfInt
```

If the value of `toFind` is not in the array, then eventually the value of `searchPos` will reach `anArray.length`. At that point the first **conjunct** of the **while loop condition**, `searchPos < anArray.length` becomes **`false`** and hence so does the **conjunction** itself, without it evaluating the second conjunct, `anArray[searchPos] != toFind`. If on the other hand we swapped over the two conjuncts, when `searchPos` reaches that same value the (now) first conjunct would cause an `ArrayIndexOutOfBoundsException`.

```
  // Definitely silly code.
  while (anArray[searchPos] != toFind && searchPos < anArray.length)
    searchPos++;
```

13.11 Design: Searching a list: binary search (page 525)

When searching for a particular item in a **list** of items we can seriously improve efficiency compared with performing a **linear search** if the items are already **sort**ed into a known **total order** and we use a **binary search**. However, this more efficient approach is also more complicated than the simple, but slow, linear search. This exhibits the unfortunately typical trade-off between speed and simplicity.

In a binary search we have two indices `low` and `high` which start off as **index**ing the first and last elements of the **data** in the list. At any time, the item we are looking for lies somewhere between `low` and `high`, or it is not present. So we look at the item half way between them. If it is the one we are looking for, then the process is over. If it is **less than** the one we want, we move the `low` index up to that half way point plus one, otherwise we move the `high` index down to one less than it. If `low` and `high` meet or cross over then we can stop looking – the item is not there.

DRAFT April 15, 2011

```
list = ... items are stored in the list in ascending order
searchItem = .. the item we wish to find in list
int lowIndex = 0
int highIndex = list.length - 1
int midIndex = (lowIndex + highIndex) / 2
while lowIndex < highIndex && list[midIndex] != searchItem
  if list[midIndex] < searchItem
    lowIndex = midIndex + 1
  else
    highIndex = midIndex - 1
  midIndex = (lowIndex + highIndex) / 2
end-while
if list[midIndex] == searchItem
  ... you found it
else
  ... searchItem is not in the list
```

13.12 Design: UML (page 381)

Many professional Java programmers express their **design**s using the **Unified Modelling Language** (**UML**). This is a collection of diagram types which can be used to show various relationships between entities, such as **object**s and **class**es.

13.13 Design: UML: class diagram (page 381)

A **UML class diagram** can be used to represent an **inheritance hierarchy**. Each **class** appears as a box with its name, its **variable**s and its **method**s. Items which are **private** are marked with a - and **public** ones are marked with a +.

13.14 Design: Storing data (page 547)

Collections of **data** which need to be stored in the **computer memory** at **run time** are placed in a **data structure**. Perhaps the most obvious example is the **array**, in which data might be stored in an arbitrary or specific order. One common thing we want to do with stored data is find it, using some kind of search **algorithm** such as a **linear search** (see Section 14.6.2 on page 323) or a **binary search** (see Section 20.4.3 on page 525). The latter requires the data to be sorted in a particular order, using a **sort algorithm** such as **bubble sort** (see Section 14.3 on page 296).

DRAFT April 15, 2011

13.15 Design: Storing data: hash table (page 547)

One way of storing **data** so that it can be retrieved quickly is to use a **hash table**. This **data structure** places (**reference**s to) items in an **array**, where the **array index** is based on a **hash code** provided by each item that might need to be stored. The idea is that data items which are **equivalent** *must* have the same hash code, and items which are not equivalent try to have different hash codes. To insert an item into the structure, we take its hash code, which is an **integer**, and divide it by the size of the array we are using for the hash table, take the remainder and place the item at that array index. To see if an item is already in the hash table, we compute the corresponding array index and check the array.

The following diagram shows a hash table of size eleven, holding five items with hash codes 22, 223, 38, 30 and 119.

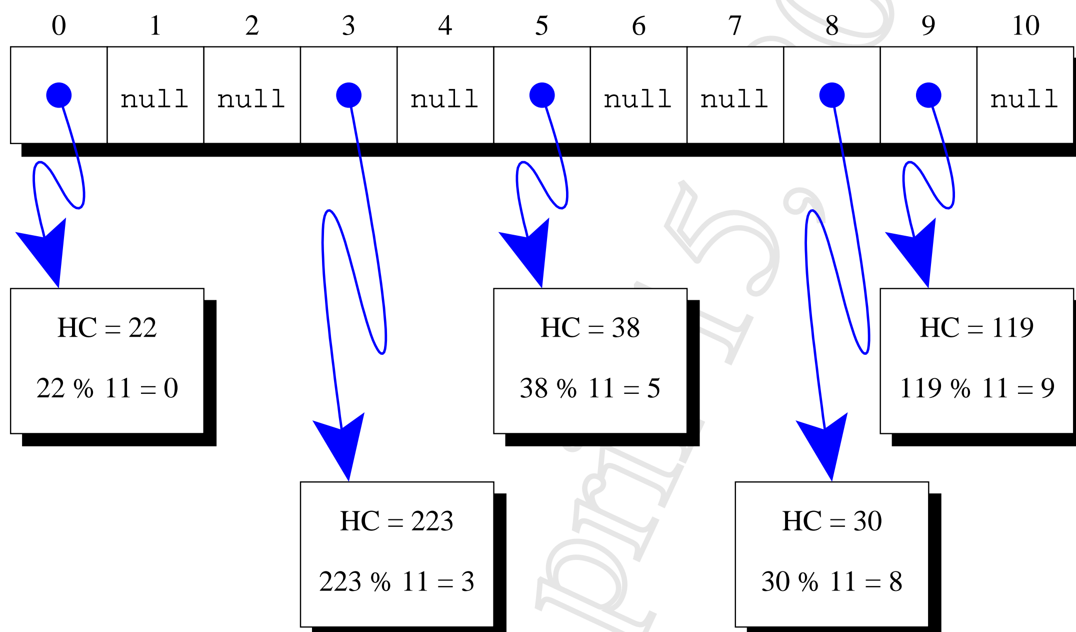

We may of course have clashes, caused by two items which are not equivalent having the same array index, and there are various strategies for coping with that, such as merely finding the next available free slot in the array. However that leads to a partial **linear search** to find such items later. To get the best efficiency from the hash table, it is important to minimize the occurrence of clashes, so we typically make the size of the array a **prime number** and, when **design**ing the **function** which computes the hash code of an element, we try to make non-equivalent items get different hash codes.

13.16 Design: Storing data: ordered binary tree (page 552)

An **ordered binary tree** (**OBT**) is a **data structure** which allows for quick storage and retrieval of **data**. The structure is so named because the data is stored in a tree, with each branch having

DRAFT April 5, 2011

a possible left subtree and/or a right subtree (binary) and the data is kept in some **total order** from left to right across the tree. That is, for every item in the tree, all items in its left subtree are **less than** it (according to whatever total order is being used) and all items in its right subtree are greater.

For example, the following diagram shows an OBT containing the ten **integer**s 12, 17, 19, 27, 34, 49, 53, 75, 81 and 99.

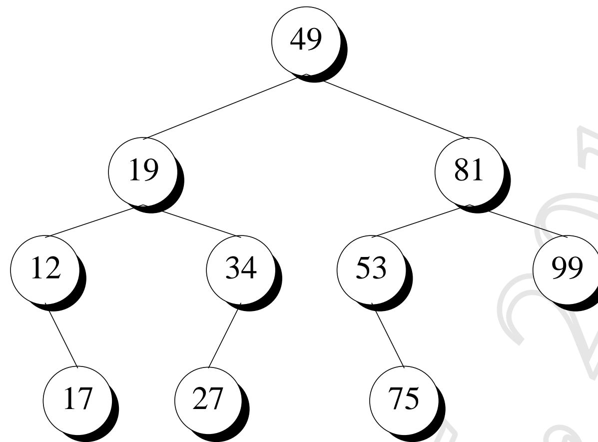

Because the data is ordered, we do not have to search the entire tree to find an item. Instead we start at the top and, if we have not yet found it, we go left if the item we want is less than the item where we are, or right otherwise, and repeat these steps until we either find the item or reach the bottom.

Searching an OBT has similar efficiency to a **binary search** (see Section 20.4.3 on page 525) because we (essentially) halve the search space at each stage as we proceed down the tree. Although this is not as fast as using a **hash table** with a *good* (and quick) **hash code function**, an OBT is useful in situations where we wish to retrieve the data from it in order.

13.17 Design: Storing data: linked list (page 557)

A **linked list** is a **data structure** which holds **data** in a chain of link **object**s, each containing a (**reference** to) one data element, and a reference to the next link object. In a **doubly linked list**, such as the one shown in the following diagram, each link also has a reference to the previous link.

DRAFT April 15, 2011

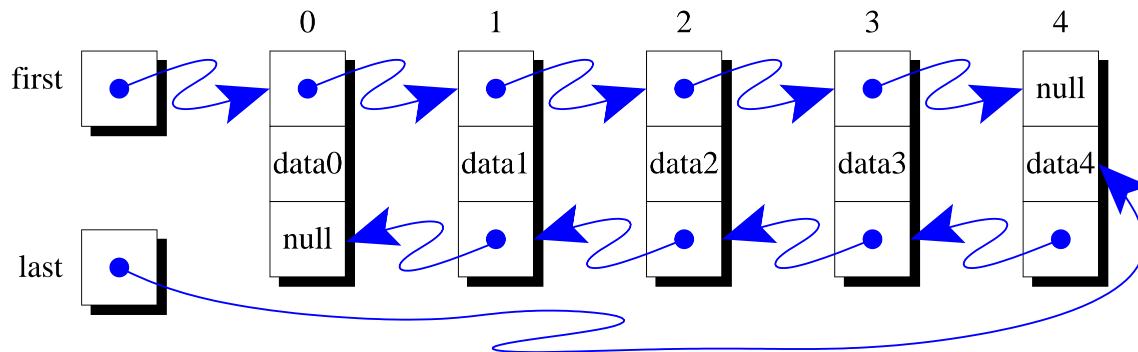

References are kept to the first and last links in the chain. To access an element at a particular **list index**, the chain must be followed from the front counting links until that index is reached, or from the back if that is nearer. So, this kind of **list** is not very efficient if many **random access**es of elements are needed. It can, however, be more efficient than using an **array** in other situations, such as adding at the back without ever having to use **array extension**, and adding or removing at the front or middle, without the need to shuffle the existing elements along.

14 Variable

14.1 Variable (page 36)

A **variable** in Java is an entity that can hold a **data** item. It has a name and a value. It is rather like the notion of a variable in algebra (although it is not quite the same thing). The name of a variable does not change – it is carefully chosen by the programmer to reflect the meaning of the entity it represents in relation to the problem being solved by the program. However, the *value* of a variable can (in general) be changed – we can vary it. Hence the name of the concept: a **variable** is an entity that has a (possibly) varying value.

The Java **compiler** implements variables by mapping their names onto **computer memory** locations, in which the values associated with the variables will be stored at **run time**.

So one view of a variable is that it is a box, like a pigeon hole, in which a value can be placed. If we wish, we can get the program to place a different value in that box, replacing the previous; and we can do this as many times as we want to.

Variables only have values at run time, when the program is **run**ning. Their names, created by the programmer, are already fixed by the time the program is **compile**d. Variables also have one more attribute – the **type** of the data they are allowed to contain. This too is chosen by the programmer.

DRAFT April 15, 2011

14.2 Variable: int variable (page 37)

In Java, **variable**s must be declared in a **variable declaration** before they can be used. This is done by the programmer stating the **type** and then the name of the variable. For example the code

```
int noOfPeopleLivingInMyStreet;
```

declares an **int variable**, that is a variable the value of which will be an `int`, and which has the name `noOfPeopleLivingInMyStreet`. Observe the semi-colon (`;`) which, according to the Java **syntax** rules, is needed to terminate the variable declaration. At **run time**, this variable is allowed to hold an **integer** (whole number). Its value can change, but it will always be an `int`. The name of a variable should reflect its intended meaning. In this case, it would seem from its name that the programmer intends the variable to always hold the number of people living in his or her street. The programmer would write code to ensure that this meaning is always reflected by its value at run time.

By convention, variable names start with a lower case letter, and consist of a number of words, with the first letter of each subsequent word capitalized.

14.3 Variable: a value can be assigned when a variable is declared (page 42)

Java permits us to assign a value to a **variable** at the same time as declaring it. You could regard this as a kind of **assignment statement** in which the variable is also declared at the same time. For example

```
int noOfHousesInMyStreet = 26;
```

14.4 Variable: double variable (page 54)

We can declare **double variable**s in Java, that is **variable**s which have the **type** `double`. For example the code

```
double meanAgeOfPeopleLivingInMyHouse;
```

declares a **variable** of type `double`, with the name `meanAgeOfPeopleLivingInMyHouse`. At **run time**, this variable is allowed to hold a `double` **data** item, that is a **real** (fractional decimal number). The value of this variable can change, but it will always be a `double`, including of course, approximations of *whole* numbers such as `40.0`.

DRAFT April 15, 2011

14.5 Variable: can be defined within a compound statement (page 92)

We can declare a **variable** within the body of a **method**, such as `main()`, (practically) anywhere where we can have a **statement**. The variable can then be used from that point onwards within the method body. The area of code in which a variable may be used is called its **scope**.

However, if we declare a variable within a **compound statement**, its scope is restricted to the compound statement: it does not exist after the end of the compound statement. This is a good thing, as it allows us to localize our variables to the exact point of their use, and so avoid cluttering up other parts of the code with variables available to be used but which have no relevance.

Consider the following symbolic example.

```
public static void main(String[] args)
{
  ...
  int x = ...
  ... x is available here.
  while (...)
  {
    ... x is available here.
    int y = ...
    ... x and y are available here.
  } // while
  ... x is available here, but not y,
  ... so we cannot accidentally refer to y instead of x.
} // main
```

The variable `x` can be used from the point of its definition onwards up to the end of the method, whereas the variable `y` can only be used from the point of its definition up to the end of the compound statement which is the body of the **loop**.

14.6 Variable: local variables (page 124)

When we declare **variable**s inside a **method**, they are local to that method and only exist while that method is running – they cannot be accessed by other methods. They are known as **local variable**s or **method variable**s. Also, different methods can have variables with the same name – they are different variables.

DRAFT April 15, 2011

14.7 Variable: class variables (page 124)

We can declare **variable**s directly inside a **class**, outside of any **method**s. Such **class variable**s exist from the moment the class is loaded into the **virtual machine** until the end of the program, and they can be accessed by any method in the class. For example, the following are three class variables which might be used to store the components of today's date.

```
private static int presentDay;
private static int presentMonth;
private static int presentYear;
```

Notice that we use the **reserved word** **`static`** in their declaration. Also, class variables have a visibility **modifier** – the above have all been declared as being **private**, which means they can only be accessed by code inside the class which has declared them.

14.8 Variable: a group of variables can be declared together (page 129)

Java permits us to declare a group of **variable**s which have the same **type** in one declaration, by writing the type followed by a comma-separated list of the variable names. For example

```
int x, y;
```

declares two variables, both of type **`int`**. We can even assign values to the variables, as in the following.

```
int minimumVotingAge = 18, minimumArmyAge = 16;
```

This shorthand is not as useful as one might think, because of course, we typically have a **comment** before each variable explaining what its meaning is. However, we can sometimes have one comment which describes a group of variables.

14.9 Variable: boolean variable (page 133)

The **`boolean`** **type** can be used in much the same way as **`int`** and **`double`**, in the sense that we can have **boolean variable**s and **method**s can have **`boolean`** as their **return type**.

For example, consider the following code.

DRAFT April 15, 2011

```
if (age1 < age2 || age1 == age2 && height1 <= height2)
  System.out.println("You are in the correct order.");
else
  System.out.println("Please swap over.");
```

We could, if we wished, write it using a `boolean` variable.

```
boolean correctOrder = age1 < age2 || age1 == age2 && height1 <= height2;
if (correctOrder)
  System.out.println("You are in the correct order.");
else
  System.out.println("Please swap over.");
```

Some people would argue that this makes for more readable code, as in effect, we have named the **condition** in a helpful way. How appropriate that is would depend on how obvious the code is otherwise, which is context dependent and ultimately subjective. Of course, the motive for storing the condition value in a **variable** is less subjective if we wish to use it more than once.

```
boolean correctOrder = age1 < age2 || age1 == age2 && height1 <= height2;
if (correctOrder)
  System.out.println("You are in the correct order.");
else
  System.out.println("Please swap over.");

... Lots of stuff here.

if (!correctOrder)
  System.out.println("Don't forget to swap over!");
```

Many novice programmers, and even some so-called experts, when writing the code above may have actually written the following.

```
boolean correctOrder;
if (age1 < age2 || age1 == age2 && height1 <= height2)
  correctOrder = true;
else
  correctOrder = false;

if (correctOrder == true)
  System.out.println("You are in the correct order.");
else
  System.out.println("Please swap over.");

... Lots of stuff here.
```

DRAFT April 15, 2011

```
if (correctOrder == false)
  System.out.println("Don't forget to swap over!");
```

There are three *terrible* things wrong with this code (two of them are the same really) – identify them, *and do not write code like that*!

14.10 Variable: char variable (page 145)

We can declare **char variable**s in Java, that is **variable**s which have the **type** `char`. For example the code

```
char firstLetter = 'J';
```

declares a variable of type `char`, with the name `firstLetter`. At **run time**, this variable is allowed to hold a `char` **data** item, that is a single **character**.

14.11 Variable: instance variables (page 159)

The **variable**s that we wish to have inside **object**s are called **instance variable**s because they belong to the **instance**s of a **class**. We declare them in much the same way as we declare **class variable**s, except without the **reserved word** `static`. For example, the following code is part of the definition of a `Point` class with two instance variables to be used to store the components of a `Point` object.

```
public class Point
{
  private double x;
  private double y;
  ...
} // class Point
```

Like class variables, instance variables have a visibility **modifier** – the above variables have both been declared as being **private**, which means they can only be accessed by code inside the class which has declared them.

Class variables belong to the class in which they are declared, and they are created at **run time** in the **static context** when the class is loaded into the **virtual machine**. There is only one copy of each class variable. By contrast, instance variables are created dynamically, in a **dynamic context**, when the object they are part of is created during the **run** of the program. There are as many copies of each instance variable as there are instances of the class: each object has its own set of instance variables.

DRAFT April 15, 2011

14.12 Variable: instance variables: should be private by default (page 175)

Java allows us to give **public** visibility to our **instance variable**s if we wish, but generally it is a good idea to define them as **private**. This permits us to alter the way we implement the **class**, without it affecting the code in other classes. For example, the programmer who has the job of maintaining a `Point` class with instance variables `x` and `y`, might decide it was better to re-implement the class to use instance variables that store the polar coordinate radius and angle instead. This might be because some new **method**s being added to the class would work much more easily in the polar coordinate system. Because the `x` and `y` instance variables had originally been made private, the programmer would know that there could not be any mention of them in other classes. So it would be safe to replace them with ones of a different name and which work differently. To make the points behave the same as before, the values given to the **constructor method** would be converted from *x* and *y* values to polar values, before being stored, and the `toString()` method could convert them back again.

14.13 Variable: of a class type (page 161)

As a **class** is a **type**, we can use one in much the same way as we use the built-in types, such as **`int`**, **`double`** and **`boolean`**. This means we can declare a **variable** whose type is a class. For example, if we have a class `Point` then we can have variables of type `Point`.

```
Point p1;
Point p2;
```

The above defines two **local variable**s or **method variable**s of type `Point`. We also can have **class variable**s and even **instance variable**s whose type is a class.

14.14 Variable: of a class type: stores a reference to an object (page 162)

There is one important difference between a **variable** whose **type** is a built-in **primitive type**, such as **`int`** and one whose type is a **class**. With the former, Java knows from the type how much memory will be needed for the variable. For example, a **double variable** needs more memory than an **int variable**, but all variables of type **`int`** need the same amount of memory, as do those of type **`double`**. Java needs this information so that it knows how to allocate memory addresses for variables.

By contrast, it is not possible to calculate how much memory will be needed to store an **object**, because **instance**s of different classes will have different sizes, and in some cases it is possible for different instances of the same class to have different sizes! The only time the size of an object is reliably known is when it is created, at **run time**.

DRAFT April 15, 2011

To deal with this situation in a systematic way, variables which are of a class type do not store an object, but instead store a **reference** to an object. A reference to an object is essentially the memory address at which the object resides in memory, and is only known at run time when the object is created. Because they are really just memory addresses, the size of all references is the same, and is fixed. So by using references in variables of a class type, rather than actually storing objects, Java knows how much memory to allocate for any such variable.

Strictly speaking then, a type which is a class, is actually the **set** of possible *references* to instances of the class, rather than the set of actual instances themselves.

14.15 Variable: of a class type: stores a reference to an object: avoid misunderstanding (page 170)

Students new to the idea of **reference**s often fail to appreciate their significance, and make one or sometimes both of the following two mistakes.

1. Misconception: A **variable** is an **object**.
2. Misconception: A variable contains an object.

Neither of these are true, as we have already said: variables (of a **class type**) can contain a *reference* to an object. A common question is "why do we have to write `Date` twice in the following?".

```
Date someBirthday
= new Date(birthDate.day, birthDate.month, birthDate.year + 1);
```

It is because we are doing three things.

1. We are declaring a variable.
2. We are **construct**ing an object.
3. We are storing a reference to that object in the variable.

So we can have a variable without an object.

```
Date someBirthday;
```

And we can have an object without a variable – could that be useful?

DRAFT April 15, 2011

```
new Date(birthDate.day, birthDate.month, birthDate.year + 1);
```

Yes, it can be useful: for example, when we want to use objects just once, straight after constructing them.

```
System.out.println(new Point(3, 4).distanceFromPoint(new Point(45, 60)));
```

If we wish, we can have two variables referring to the same object.

```
Date theSameBirthday = someBirthday;
```

Also, we can change the value of a variable making it refer to a different object.

```
someBirthday = new Date(someBirthday.day, someBirthday.month,
                        someBirthday.year + 1);
```

This creates a **new** `Date` **object**, and stores the **reference** to it in `someBirthday` – overwriting the reference to the previous `Date` object. This is illustrated in the following diagram.

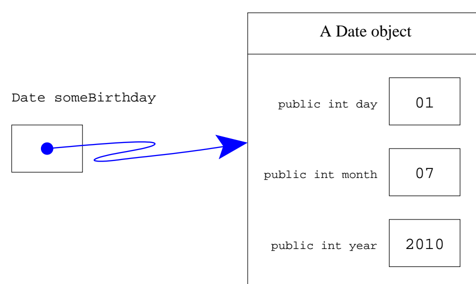


```
someBirthday = new Date(someBirthday.day, someBirthday.month, someBirthday.year + 1);
```

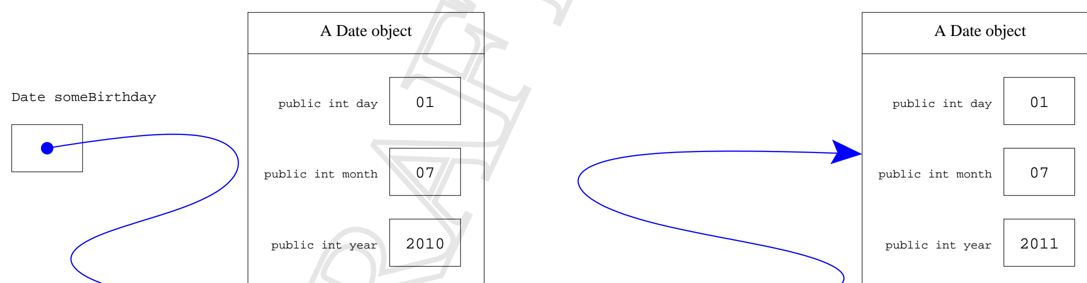

DRAFT April 15, 2011

14.16 Variable: of a class type: null reference (page 192)

When an **object** is created, the **constructor method return**s a **reference** to it, which is then used for all accesses to the object. Typically, this reference is stored in a **variable**.

```
Point p1 = new Point(75, 150);
```

There is a special reference value, known as the **null reference**, which does not refer to an object. We can talk about it using the **reserved word** **`null`**. It is used, for example, as a value for a variable when we do not want it to refer to any object at this moment in time.

```
Point p2 = null;
```

So, in the example code here we have two `Point` variables, `p1` and `p2`, but (at **run time**) only one `Point` object.

Suppose the `Point` **class** has **instance method**s `getX()` and `getY()` with their obvious implementations. Then obtaining the *x* value of the object referenced by `p1` is fine; the following code would print `75`.

```
System.out.println(p1.getX());
```

However, the similar code involving `p2` would cause a **run time error** (an **exception** called `NullPointerException`).

```
System.out.println(p2.getX());
```

This is because there is no object referenced by `p2`, and so any attempt to access the referenced object must fail.

14.17 Variable: of a class type: holding the same reference as some other variable (page 216)

A **variable** which is of a **class type** can hold a **reference** to any **instance** of that class (plus the **null reference**). There is nothing to stop two (or more) variables having the same reference value. For example, the following code creates one `Point` **object** and has it referred to by two variables.

DRAFT April 15, 2011

```
Point p1 = new Point(10, 30);

Point p2 = p1;
```

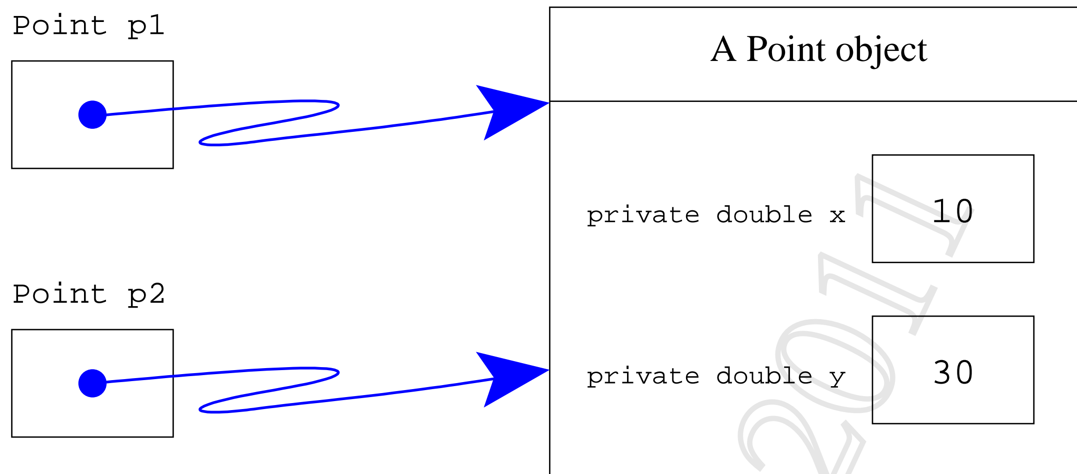

This reminds us that a variable is *not* itself an object, but merely a holder for a reference to an object.

Having two or more **variable**s refer to the same **object** can cause us no problems if it is an **immutable object** because we cannot change the object's state no matter which variable we use to access it. So, in effect, the object(s) referred to by the two variables behave the same as they would if they were two different objects. The following code has the same *effect* as the above fragment, almost no matter what we do with p1 and p2 subsequently.

```
Point p1 = new Point(10, 30);

Point p2 = new Point(10, 30);
```

The only behavioural difference between the two fragments is the **condition**s `p1 == p2` and `p1 != p2` which are **true** and **false** respectively for the first code fragment, and the other way round for the second one.

If, on the other hand, an **object referenced by more than one variable** is a **mutable object** we have to be careful because any change made via any one of the variables causes the change to occur in the (same) object referred to by the other variables. This may be, and often is, exactly what we want, or it may be a problem if our **design** is poor or if we have made a mistake in our code and the variables were not meant to share the object.

Consider the following simple example.

```
public class Employee
```

DRAFT April 15, 2011

```
{
  private final String name;
  private int salary;

  public Employee(String requiredName, int initialSalary)
  {
    name = requiredName;
    salary = initialSalary;
  } // Employee

  public String getName()
  {
    return name;
  } // getName

  public void setSalary(int newSalary)
  {
    salary = newSalary;
  } // setSalary

  public int getSalary()
  {
    return salary;
  } // getSalary

} // class Employee

  ...

  Employee debora = new Employee("Debs", 50000);
  Employee sharmane = new Employee("Shaz", 40000);

  ...

  Employee worstEmployee = debora;
  Employee bestEmployee = sharmane;

  ...
```

Now let us have an accidental piece of code.

```
  worstEmployee = bestEmployee;
```

Then we carry on with intentional code.

```
  ...
```

DRAFT April 15, 2011

```
bestEmployee.setSalary(55000);
worstEmployee.setSalary(0);

System.out.println("Our best employee, " + bestEmployee.getName()
                   + ", is paid " + bestEmployee.getSalary());
System.out.println("Our worst employee, " + worstEmployee.getName()
                   + ", is paid " + worstEmployee.getSalary());
```

The effect of the accidental sharing is to give Sharmane, who is our best employee, a pay increase to 55,000 immediately followed by a pay cut to zero because `worstEmployee` and `bestEmployee` are both referring to the same object, the one which is also referred to by `sharmane`. Meanwhile our worst employee, Debora, gets to keep her 50,000! Further more, the report only actually talks about Sharmane in both contexts!

```
Our best employee, Shaz, is paid 0
Our worst employee, Shaz, is paid 0
```

14.18 Variable: final variables (page 194)

When we declare a **variable** we can write the **reserved word `final`** as one of its **modifier**s before the **type** name. This means that once the variable has been given a value, that value cannot be altered.

If an **instance variable** is declared to be a **final variable** then it must be explicitly assigned a value by the time the **object** it belongs to has finished being **construct**ed. This would be done either by assigning a value in the **variable declaration**, or via an **assignment statement** inside the **constructor method**.

14.19 Variable: final variables: class constant (page 205)

A **class variable** which is declared to be a **final variable** (i.e. its **modifier**s include the **reserved word**s **`static`** and **`final`**) is also known in Java as a **class constant**. An example of this is the **variable** in the **class** `java.lang.Math` called `PI`.

```
public static final double PI = 3.14159265358979323846;
```

By convention, class constants are usually named using only capital letters with the words separated by underscores (_).

DRAFT April 15, 2011

14.20 Variable: final variables: class constant: a set of choices (page 308)

One use of **class constant**s is to define a set of options for the users of a **class**, without them having to know what values have been chosen to model each option – they instead use the name of one or more class constants to represent their choices.

For example, the following could be possible directions available in a class that is part of a game that permits simple movement of some game entity.

```
public static final int UP = 0;
public static final int DOWN = 1;
public static final int LEFT = 2;
public static final int RIGHT = 3;
```

Apart from leading to more readable code, this technique gives us more flexibility: the maintainer of the **source code** might decide for some reason to change the values (but not the names) of the four constants. This should not cause any code outside of the class to need rewriting.

14.21 Variable: final variables: class constant: a set of choices: dangerous (page 308)

The use of `int` **class constant**s to model a small set of options does have two dangers.

- The constants could be used for other purposes – e.g. they could be used inappropriately in some **arithmetic expression**.
- Someone may accidentally use another `int` value which is not one of the constants in places where a constant should be used. The **compiler** would accept it because it is an `int`.

14.22 Variable: of an array type (page 287)

We can declare **variable**s of an **array type** rather like we can of any other **type**. For example, here is a variable of type `int[]`.

```
int[] salaries;
```

As **array**s are **object**s, they are accessed via **reference**s. So an **array variable** at **run time** holds either a *reference* to an array or the **null reference**. The following diagram shows the above variable referring to an array of `int` values.

DRAFT April 15, 2011

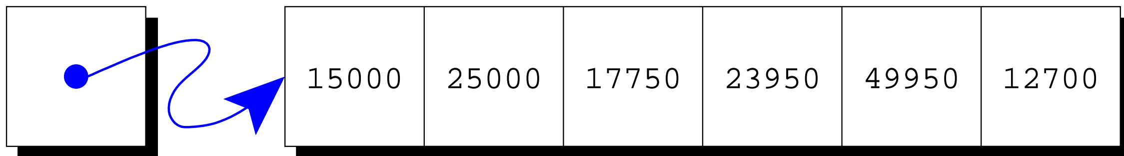

14.23 Variable: initial value (page 453)

When **class variable**s, **instance variable**s, and **array element**s are created, they are given a default initial value by the **virtual machine** (unless they are also final variables). In contrast, the **compiler** forces **local variable**s (**method variable**s) and **final variable**s to be initialized by our code.

It is dangerous to *quietly* rely on default values when they happen to be the initial values we desire, mainly because anyone looking at our code (including ourselves) cannot tell the difference between us doing that and having forgotten to initialize! Another reason is that sometimes you, or a reader of your program, may misremember what initial value there is for a **variable** of a particular **type**. So, one rule of thumb is to always perform our own initialization to make it clear we have not overlooked it. However, where that is non-trivial (e.g. for array elements), we instead write a clear **comment** stating that we are happy the default value is what we want, and what that value is.

15 Expression

15.1 Expression: arithmetic (page 38)

We can have **arithmetic expression**s in Java rather like we can in mathematics. These can contain **literal value**s, that is constants, such as the **integer literal**s `1` and `18`. They can also contain **variable**s which have already been declared, and **operator**s to combine sub-expressions together. Four common **arithmetic operator**s are **addition** (`+`), **subtraction** (`-`), **multiplication** (`*`) and **division** (`/`). Note the use of an asterisk for multiplication, and a forward slash for division – computer keyboards do not have multiply or divide symbols.

These four operators are **binary infix operator**s, because they take two **operand**s, one on either side of the operator. `+` and `-` can also be used as the **unary prefix operator**s, **plus** and **minus** respectively, as in `-5`.

When an **expression** is **evaluate**d (**expression evaluation**) Java replaces each variable with its current value and works out the result of the expression depending on the meaning of the operators. For example, if the variable `noOfPeopleLivingInMyStreet` had the value `47` then the expression `noOfPeopleLivingInMyStreet + 4` would evaluate to `51`.

DRAFT April 15, 2011

15.2 Expression: arithmetic: int division truncates result (page 52)

The four **arithmetic operator**s, +, -, * and / of Java behave very similarly to the corresponding operators in mathematics. There is however one serious difference to look out for. When the **division operator** is given two **integer**s (whole numbers) it uses **integer division** which always yields an integer as its result, by throwing away any fractional part of the answer. So, `8 / 2` gives the answer `4` as you might expect, but `9 / 2` also gives `4` – not `4.5` as it would in mathematics. It does not round to the nearest whole number, it always rounds towards zero. In mathematics `15 / 4` gives `3.75`. In Java it yields `3` not `4`.

15.3 Expression: arithmetic: associativity and int division (page 52)

Like the **operator**s + and -, the operators * and / have equal **operator precedence** (but higher than + and -) and also have **left associativity**.

However, there is an extra complication to consider because the Java / operator truncates its answer when given two **integer**s. Consider the following two **arithmetic expression**s.

Expression	Implicit brackets	Value
`9 * 4 / 2`	`(9 * 4) / 2`	18
`9 / 2 * 4`	`(9 / 2) * 4`	16

In mathematics one would expect to get the same answer from both these **expression**s, but not in Java!

15.4 Expression: arithmetic: double division (page 55)

The Java **division operator**, /, uses **double division** and produces a **`double`** result if at least one of its **operand**s is a **`double`**. The result will be the best approximation to the actual answer of the division.

Expression	Result	Type of Result
`8 / 2`	`4`	`int`
`8 / 2.0`	`4.0`	`double`
`9 / 2`	`4`	`int`
`9 / 2.0`	`4.5`	`double`
`9.0 / 2`	`4.5`	`double`
`9.0 / 2.0`	`4.5`	`double`

DRAFT April 5, 2011

15.5 Expression: arithmetic: double division: by zero (page 291)

When using the **double division** operation in Java, if the numerator is not zero but the denominator is zero, the result we get is a model of **infinity**. This is represented, for example by `System.out.println()`, as `Infinity`.

However, if both the numerator and the denominator are zero, we instead get a model of the concept **not a number**, which is represented as `NaN`.

This behaviour of double division is in contrast to **integer division**, which produces an **exception** if the denominator is zero.

15.6 Expression: arithmetic: remainder operator (page 149)

Another **arithmetic operator** in Java is the **remainder operator**, also known as the **modulo** operator, `%`. When used with two **integer operands**, it yields the remainder obtained from dividing the first operand by the second. As an example, the following **method** determines whether a given `int` **method parameter** is an even number.

```
public static boolean isEven(int number)
{
  return number % 2 == 0;
} // isEven
```

15.7 Expression: arithmetic: shift operators (page 473)

Some more **arithmetic operators** in Java are the **shift operators**, `<<`, `>>` and `>>>`. The **left shift operator**, `<<`, yields the number obtained by shifting the first **operand** left by the number of **bits** given in the second operand, placing zeroes in that many rightmost places. The **unsigned right shift** operator, `>>>`, similarly shifts rightwards, placing zeroes on the left. The **signed right shift** operator, `>>`, is the same, except it places ones on the left if the number being shifted is negative.

For example, `1000` is `0001111101000` in **binary**.

4096	**2048**	**1024**	**512**	**256**	**128**	**64**	**32**	**16**	**8**	**4**	**2**	**1**
0	0	0	1	1	1	1	1	0	1	0	0	0
0+	0+	0+	512+	256+	128+	64+	32+	0+	8+	0+	0+	0 = 1000

When this is shifted left by three places, `1000 << 3`, we get the result `8000` which is `1111101000000` in binary.

DRAFT April 15, 2011

4096	2048	1024	512	256	128	64	32	16	8	4	2	1
1	1	1	1	1	0	1	0	0	0	0	0	0
4096+	2048+	1024+	512+	256+	0+	64+	0+	0+	0+	0+	0+	0 = 8000

Whereas, `1000 >> 3` and `1000 >>> 3` both yield `0000001111101` in binary, which is `125`.

4096	2048	1024	512	256	128	64	32	16	8	4	2	1
0	0	0	0	0	0	1	1	1	1	1	0	1
0+	0+	0+	0+	0+	0+	64+	32+	16+	8+	4+	0+	1 = 125

Shifting left by n bits, has the same effect as **multiplication** by 2^n and discarding any overflow. Signed shifting right by n bits has the same effect as dividing by 2^n and discarding any remainder.

15.8 Expression: arithmetic: integer bitwise operators (page 474)

The **operator**s `|`, `&`, and `^`, when applied to numeric **operand**s, have the effect of an **integer bitwise or**, **integer bitwise and** and **integer bitwise exclusive or**, respectively. The result is obtained by pairing the corresponding **bit**s of each operand according to the following table.

bit *n* **of** `op1`	**bit** *n* **of** `op2`	**bit** *n* **of** `op1 | op2`	**bit** *n* **of** `op1 & op2`	**bit** *n* **of** `op1 ^ op2`
0	0	0	0	0
0	1	1	0	1
1	0	1	0	1
1	1	1	1	0

For example, the value `1000` which is `1111101000` in **binary**, when anded with the value `23` which is `0000010111` in binary, yields `0000000000` – because they have no corresponding bit values in common. When they are instead or-ed together, we get `1111111111` in binary, which is `1023`. This happens to be the same as `1000 + 23`, but **integer bitwise or** is the same as **addition** only when the two numbers have no corresponding bits with the same value.

15.9 Expression: brackets and precedence (page 45)

In addition to **operator**s and **variable**s, **expression**s in Java can have round brackets in them. As in mathematics, brackets are used to define the structure of the expression by grouping parts of it into sub-expressions. For example, the following two expressions have different structures, and thus very different values.

```
(2 + 4) * 8
2 + (4 * 8)
```

DRAFT April 5, 2017

The value of the first expression is made from the **addition** of `2` and `4` and then **multiplication** of the resulting `6` by `8` to get `48`. The second expression is **evaluate**d by multiplying `4` with `8` to get `32` and then adding `2` to that result, ending up with `34`.

To help us see the structure of these two expressions, let us draw them as **expression tree**s.

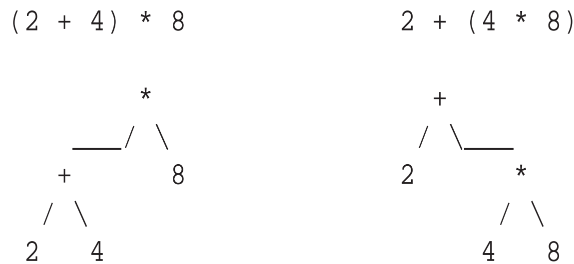

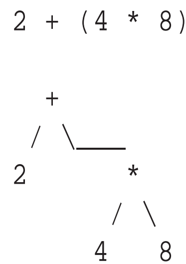

What if there were no brackets?

```
2 + 4 * 8
```

Java allows us to have expressions without any brackets, or more generally, without brackets around *every* sub-expression. It provides rules to define what the structure of such an expression is, i.e., where the missing brackets should go. If you look at the `4` in the above expression, you will see that it has an operator on either side of it. In a sense, the `+` operator and the `*` operator are both fighting to have the `4` as an **operand**. Rather like a tug of war, `+` is pulling the `4` to the left, and `*` is tugging it to the right. The question is, which one wins? Java, as in mathematics, provides the answer by having varying levels of **operator precedence**. The `*` and `/` operators have a higher precedence than `+` and `-`, which means `*` fights harder than `+`, so it wins! `2 + 4 * 8` evaluates to `34`.

15.10 Expression: associativity (page 48)

The principle of **operator precedence** is insufficient to disambiguate all **expression**s which are not fully bracketed. For example, consider the following expressions.

```
10 + 7 + 3
10 + 7 - 3
10 - 7 + 3
10 - 7 - 3
```

In all four expressions, the `7` is being fought over by two **operator**s which have the same precedence: either two `+`, two `-`, or one of each. So where should the missing brackets go? The **expression tree**s could have one of the two following structures, where `OP1` is the first operator, and `OP2` is the second.

DRAFT April 15, 2011

```
10 OP1 (7 OP2 3)          (10 OP1 7) OP2 3
```

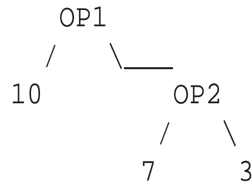

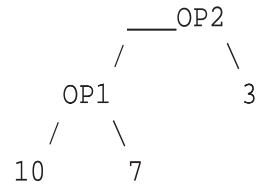

Let us see whether it makes a difference to the results of the expressions.

Expression	Value
(10 + 7) + 3	20
10 + (7 + 3)	20
(10 + 7) - 3	14
10 + (7 - 3)	14
(10 - 7) + 3	6
10 - (7 + 3)	0
(10 - 7) - 3	0
10 - (7 - 3)	6

As you can see, it does make a difference sometimes – in these cases when the first operator is **subtraction** (-). So how does Java resolve this problem? As in mathematics, Java operators have an **operator associativity** as well as a precedence. The operators +, -, * and / all have **left associativity** which means that when two of these operators of equal precedence are both fighting over one **operand**, it is the left operator that wins. If you like, the tug of war takes place on sloping ground with the left operator having the advantage of being lower down than the right one!

Expression	Implicit brackets	Value
10 + 7 + 3	(10 + 7) + 3	20
10 + 7 - 3	(10 + 7) - 3	14
10 - 7 + 3	(10 - 7) + 3	6
10 - 7 - 3	(10 - 7) - 3	0

The operators * and / also have equal precedence (but higher than + and -) so similar situations arise with those too.

15.11 Expression: boolean (page 60)

An **expression** which when **evaluate**d yields either `true` or `false` is known as a **condition**, and is typically used for controlling **conditional execution**. Conditions are also called **boolean expression**s.

DRAFT April 15, 2011

15.12 Expression: boolean: relational operators (page 60)

Java gives us six **relational operator**s for comparing values such as numbers, which we can use to make up **condition**s. These are all **binary infix operator**s, that is they take two **operand**s, one either side of the **operator**. They yield **`true`** or **`false`** depending on the given values.

Operator	Title	Description
`==`	Equal	This is the **equal** operator, which provides the notion of **equality**. `a == b` yields **`true`** if and only if the value of `a` is the same as the value of `b`.
`!=`	Not equal	This is the **not equal** operator, providing the the notion of not equality. `a != b` yields **`true`** if and only if the value of `a` is *not* the same as the value of `b`.
`<`	Less than	This is the **less than** operator. `a < b` yields **`true`** if and only if the value of `a` is less than the value of `b`.
`>`	Greater than	This is the **greater than** operator. `a > b` yields **`true`** if and only if the value of `a` is greater than the value of `b`.
`<=`	Less than or equal	This is the **less than or equal** operator. `a <= b` yields **`true`** if and only if the value of `a` is less than value of `b`, or is equal to it.
`>=`	Greater than or equal	This is the **greater than or equal** operator. `a >= b` yields **`true`** if and only if the value of `a` is greater than value of `b`, or is equal to it.

15.13 Expression: boolean: logical operators (page 128)

For some **algorithm**s, we need **condition**s on **loop**s etc. that are more complex than can be made simply by using the **relational operator**s. Java provides us with **logical operator**s to enable us to glue together simple conditions into bigger ones. The three most commonly used logical operators are **conditional and**, **conditional or** and **logical not**.

Operator	Title	Posh title	Description
`&&`	and	**conjunction**	`c1 && c2` is **`true`** if and only if both conditions `c1` and `c2` **evaluate** to **`true`**. Both of the two conditions, known as **conjunct**s, must be **`true`** to satisfy the combined condition.
`||`	or	**disjunction**	`c1 || c2` is **`true`** if and only if at least one of the conditions `c1` and `c2` evaluate to **`true`**. The combined condition is satisfied, unless both of the two conditions, known as **disjunct**s, are **`false`**.
`!`	not	**negation**	`!c` is **`true`** if and only if the condition `c` evaluates to **`false`**. This operator negates the given condition.

We can define these **operator**s using **truth table**s, where ? means the **operand** is not evaluated.

DRAFT April 15, 2011

c1	c2	c1 && c2
true	**true**	**true**
true	**false**	**false**
false	?	**false**

c1	c2	c1 || c2
true	?	**true**
false	**true**	**true**
false	**false**	**false**

c	!c
true	**false**
false	**true**

Using these operators, we can make up complex conditions, such as the following.

```
age1 < age2 || age1 == age2 && height1 <= height2
```

As with the **arithmetic operator**s, Java defines **operator precedence** and **operator associativity** to disambiguate complex conditions that are not fully bracketed, such as the one above. `&&` and `||` have a lower precedence than the relational operators which have a lower precedence than the arithmetic ones. `!` has a very high precedence (even more so than the arithmetic operators) and `&&` has a higher precedence than `||`. So the above example **expression** has implicit brackets as follows.

```
(age1 < age2) || ((age1 == age2) && (height1 <= height2))
```

This might be part of a program that **sort**s people standing in a line by age, but when they are the same age, it sorts them by height. Assuming that the **int variable**s `age1` and `height1` contain the age and height of one person, and the other two variables similarly contain that **data** for another, then the following code might be used to tell the pair to swap their order if necessary.

```
if (age1 < age2 || age1 == age2 && height1 <= height2)
  System.out.println("You are in the correct order.");
else
  System.out.println("Please swap over.");
```

We might have, perhaps less clearly, chosen to write that code as follows.

```
if (!(age1 < age2 || age1 == age2 && height1 <= height2))
  System.out.println("Please swap over.");
else
  System.out.println("You are in the correct order.");
```

You might find it tricky, but it's worth convincing yourself: yet another way of writing code with the same effect would be as follows.

```
if (age1 > age2 || age1 == age2 && height1 > height2)
  System.out.println("Please swap over.");
else
  System.out.println("You are in the correct order.");
```

DRAFT April 15, 2011

In mathematics, we are used to writing expressions such as $x \leq y \leq z$ to mean true, if and only if y lies in the range x to z, inclusive. In Java, such expressions need to be written as `x <= y && y <= z`.

Also, in everyday language we are used to using the words 'and' and 'or' where they have very similar meanings to the associated Java operators. However, we say things like "my mother's age is 46 or 47". In Java, we would need to write `myMumAge == 46 || myMumAge == 47` to capture the same meaning. Another example, "my brothers are aged 10 and 12", might be coded as `myBrother1Age == 10 && myBrother2Age == 12`.

However, there are times in everyday language when we say "and" when we really mean "or" in logic, and hence would use `||` in Java. For example, "the two possible ages for my dad are 49 *and* 53" is really the same as saying "my dad's age is 49 *or* my dad's age is 53".

15.14 Expression: boolean: logical operators: conditional (page 323)

The **logical operator**s `&&` and `||` in Java are called **conditional and** and **conditional or** because they have an important property, which distinguishes them from their classical logic counterparts. They are lazy. This means that if they can determine their result after evaluating their left **operand**, they will not **evaluate** their right one. That is, if the first **disjunct** of `||` evaluates to **`true`** it will not evaluate the second; and if the first **conjunct** of `&&` evaluates to **`false`** it will not evaluate the second. This allows us to safely write **condition**s such as the following. `data == null || data.length == 0`

15.15 Expression: conditional expression (page 94)

The **conditional operator** in Java permits us to write **conditional expression**s which have different sub-expressions **evaluate**d depending on some **condition**. The general form is

```
c ? e1 : e2
```

where `c` is some condition, and `e1` and `e2` are two **expression**s of some **type**. The condition is evaluated, and if the value is **`true`** then `e1` is evaluated and its value becomes the result of the expression. If the condition is **`false`** then `e2` is evaluated and its value becomes the result instead.

For example

```
int maxXY = x > y ? x : y;
```

is another way of achieving the same effect as the following.

DRAFT April 15, 2011

```
int maxXY;
if (x > y)
  maxXY = x;
else
  maxXY = y;
```

16 Package

16.1 Package (page 187)

There are hundreds of **class**es that come with Java in its **application program interface** (**API**), and even more that are available around the world for reusing in our programs if we wish. To help manage this huge number of classes, they are grouped into collections of related classes, called **package**s. But even this is not enough to make things manageable, so packages are grouped into a hierarchy in a rather similar way to how a well organized **file system** is arranged into directories and sub-directories. For example, there is one group of standard packages called `java` and another called `javax`.

16.2 Package: java.util (page 188)

One of the standard Java **package**s in the package group `java` is called `util`. This means its full name is `java.util` – the package addressing mechanism uses a dot (`.`) in much the same way as Unix uses a slash, or Microsoft Windows uses a backslash, to separate directories in a filename path. `java.util` contains many generally useful utility **class**es. For example, there is a class called `Scanner` which lives there, so its **fully qualified name** is `java.util.Scanner`. This fully qualified name is unique: if someone else was to create a class called `Scanner` then it would not be in the same package, so the two would not be confused.

We can refer to a class using its fully qualified name, for example the following declares a **variable** of **type** `java.util.Scanner` and creates an **instance** of the class too.

```
java.util.Scanner inputScanner = new java.util.Scanner(System.in);
```

16.3 Package: java.awt and javax.swing (page 245)

Inside the group of **package**s known as `java`, there is one called `awt`, so the the full name of the package is `java.awt`. It contains the **class**es that make up the original Java **graphical user interface** system known as the **Abstract Windowing Toolkit** (AWT). For example, there

DRAFT April 15, 2011

is a class that lives inside java.awt called Container, and so its **fully qualified name** is java.awt.Container.

Another group, javax contains a package called swing and this is the set of classes which make up the more modern **Java Swing** system, which is built on top of AWT. For example, there is a class that lives inside javax.swing called JFrame, and so its fully qualified name is javax.swing.JFrame.

Java programs that provide a **GUI** typically need to use classes from both these packages.

17 GUI API

17.1 GUI API: JFrame (page 245)

Each **instance** of the **class** javax.swing.JFrame corresponds to a window that appears on the screen.

17.2 GUI API: JFrame: setTitle() (page 246)

The **class** javax.swing.JFrame has an **instance method** called setTitle which takes a String to be used as the title of the window. This string typically appears in the title bar of the window, depending upon what window manager the user is using (in Unix worlds there is a massive variety of window managers to choose from).

17.3 GUI API: JFrame: getContentPane() (page 246)

The **class** javax.swing.JFrame has an **instance method** called getContentPane which **returns** the **content pane** of the JFrame. This is the part of the JFrame that holds the **graphical user interface** (**GUI**) components of the window. It is an **instance** of java.awt.Container.

17.4 GUI API: JFrame: setDefaultCloseOperation() (page 247)

The **class**
javax.swing.JFrame has an **instance method** called setDefaultCloseOperation which takes a **method parameter** that specifies what the JFrame should do when the end user presses the close button on the title bar of the window. There are four possible settings as follows.

DRAFT April 15, 2011

- **Do nothing on close** – Don't do anything.
- **Hide on close** – Hide the window, so that it is no longer visible, but do not destroy it.
- **Dispose on close** – Destroy the window.
- **Exit on close** – Exit the whole program.

The parameter is actually an `int`, but we do not need to know what exact value to give as a **method argument**, because there are four **class constant**s defined in `JFrame` which have the right values.

```
public static final int DO_NOTHING_ON_CLOSE = ?;
public static final int HIDE_ON_CLOSE = ?;
public static final int DISPOSE_ON_CLOSE = ?;
public static final int EXIT_ON_CLOSE = ?;
```

We simply use whichever class constant suits us, as in the following example.

```
setDefaultCloseOperation(DISPOSE_ON_CLOSE);
```

17.5 GUI API: JFrame: pack() (page 247)

The **class** `javax.swing.JFrame` has an **instance method** called `pack`. This makes the `JFrame` arrange itself ready for being shown on the screen. It works out the sizes and positions of all its components, and (in general) the size of the window itself. Typically `pack()` is called after all the **graphical user interface** (**GUI**) components have been added to the `JFrame`.

17.6 GUI API: JFrame: setVisible() (page 248)

The **class** `javax.swing.JFrame` has an **instance method** called `setVisible`. This takes a **`boolean` method parameter**, and if this value is **`true`** then it makes the `JFrame` **object** cause the window it represents to appear on the physical screen, or disappear otherwise.

17.7 GUI API: Container (page 246)

The **class** `java.awt.Container` implements part of a **graphical user interface** (**GUI**). An **instance** of the class is a component that is allowed to contain other components.

DRAFT April 15, 2011

17.8 GUI API: Container: add() (page 246)

The **class** `java.awt.Container` has an **instance method** called `add` which takes a **graphical user interface** (**GUI**) component and includes it in the collection of components to be displayed within the container.

17.9 GUI API: Container: add(): adding with a position constraint (page 268)

The **class** `java.awt.Container` has another **instance method** called `add` which takes a **graphical user interface** (**GUI**) component and some other **object** constraining how the component should be positioned. This is intended for use with **layout managers** that use position constraints, such as `java.awt.BorderLayout`. For example, the following code makes the `JLabel` appear in the north position of `myContainer`.

```
myContainer.setLayout(new BorderLayout());
myContainer.add(new JLabel("This is in the north"), BorderLayout.NORTH);
```

17.10 GUI API: Container: setLayout() (page 250)

The **class** `java.awt.Container` has an **instance method** called `setLayout` which takes an **instance** of one of the **layout manager** classes, and uses that to lay out its **graphical user interface** (**GUI**) components each time a lay out is needed, for example, when the window it is part of is **pack**ed.

17.11 GUI API: JLabel (page 246)

The **class** `javax.swing.JLabel` implements a particular part of a **graphical user interface** (**GUI**) which simply displays a small piece of text, that is, a label. The label text is specified as a `String` **method argument** to one of the `JLabel` **constructor methods**.

17.12 GUI API: JLabel: setText() (page 258)

The **class** `javax.swing.JLabel` has an **instance method** called `setText` which takes a `String` **method argument** and changes the text of the label to it.

DRAFT April 15, 2011

17.13 GUI API: LayoutManager (page 249)

A **layout manager** is a **class** which contains the logic for laying out **graphical user interface** (**GUI**) components within an **instance** of `java.awt.Container` in some set pattern. There are various types of layout manager, including the following most common ones.

- `java.awt.FlowLayout` – arrange the components in a horizontal line.
- `java.awt.GridLayout` – arrange the components in a grid.
- `java.awt.BorderLayout` – arrange the components with one at the centre, and one at each of the four sides.

17.14 GUI API: LayoutManager: FlowLayout (page 249)

The **class** `java.awt.FlowLayout` is a **layout manager** which positions all the components within an **instance** of
`java.awt.Container` in a horizontal row. The components appear in the order they were added to the container.

17.15 GUI API: LayoutManager: FlowLayout: alignment (page 278)

The **class**
`java.awt.FlowLayout` can be given an alignment mode, passed as a **method argument** to one of its **constructor method**s. It affects the behaviour of the layout in cases when the component is larger than is needed to hold the components that are in it.

The argument is an **`int`** value, and should be an appropriate **class constant**, including the following.

- `FlowLayout.CENTER` – the laid out items are centred in the container.
- `FlowLayout.LEFT` – the laid out items are on the left of the container, with unused space on the right.
- `FlowLayout.RIGHT` – the laid out items are on the right of the container, with unused space on the left.

If we do not specify an alignment then centred alignment is used.

DRAFT April 15, 2011

17.16 GUI API: LayoutManager: GridLayout (page 251)

The **class** `java.awt.GridLayout` is a **layout manager** which positions all the components within an **instance** of
`java.awt.Container` in a rectangular grid. The container is divided into equal-sized rectangles, and one component is placed in each rectangle. The components appear in the order they were added to the container, filling up one row at a time.

When we create a `GridLayout` **object**, we provide a pair of `int` **method arguments** to the **constructor method**, the first specifies the number of rows, and the second the number of columns. One of these values should be zero. For example, the following **constructs** a `GridLayout` which has three rows, and as many columns as are needed depending upon the number of components being laid out.

```
new GridLayout(3, 0);
```

This next example constructs a `GridLayout` which has two columns, and as many rows as are needed depending upon the number of components being laid out.

```
new GridLayout(0, 2);
```

If both the rows and columns arguments are non-zero, then *the columns argument is totally ignored*! Neither values may be negative, and at least one of them must be non-zero, otherwise we get a **run time error**.

We can also specify the horizontal and vertical gaps that we wish to have between items in the grid. These can be given via a constructor method that takes four arguments.

```
new GridLayout(0, 5, 10, 20);
```

The above example creates a `GridLayout` that has five columns, with a horizontal gap of 10 pixels between each column, and a vertical gap of 20 pixels between each row. A pixel is the smallest unit of display position. Its exact size will depend on the resolution and physical size of the computer monitor.

17.17 GUI API: LayoutManager: BorderLayout (page 267)

The **class** `java.awt.BorderLayout` is a **layout manager** which has slots for five components, one at the centre, and one at each of the four sides around the centre. The names

DRAFT April 15, 2011

of these positions are modelled using five **class constant**s called BorderLayout.CENTER, BorderLayout.NORTH, BorderLayout.SOUTH, BorderLayout.WEST. and BorderLayout.EAST.

A BorderLayout is designed to be used when there is one **graphical user interface** (**GUI**) component which is in some sense the main component, for example a JTextArea which contains some result of the program. We can put this in the BorderLayout.CENTER position and some other component above in the BorderLayout.NORTH position, and/or below in the BorderLayout.SOUTH position, and/or to the left in the BorderLayout.WEST position and/or to the right in the BorderLayout.EAST position.

This is shown in the following diagram.

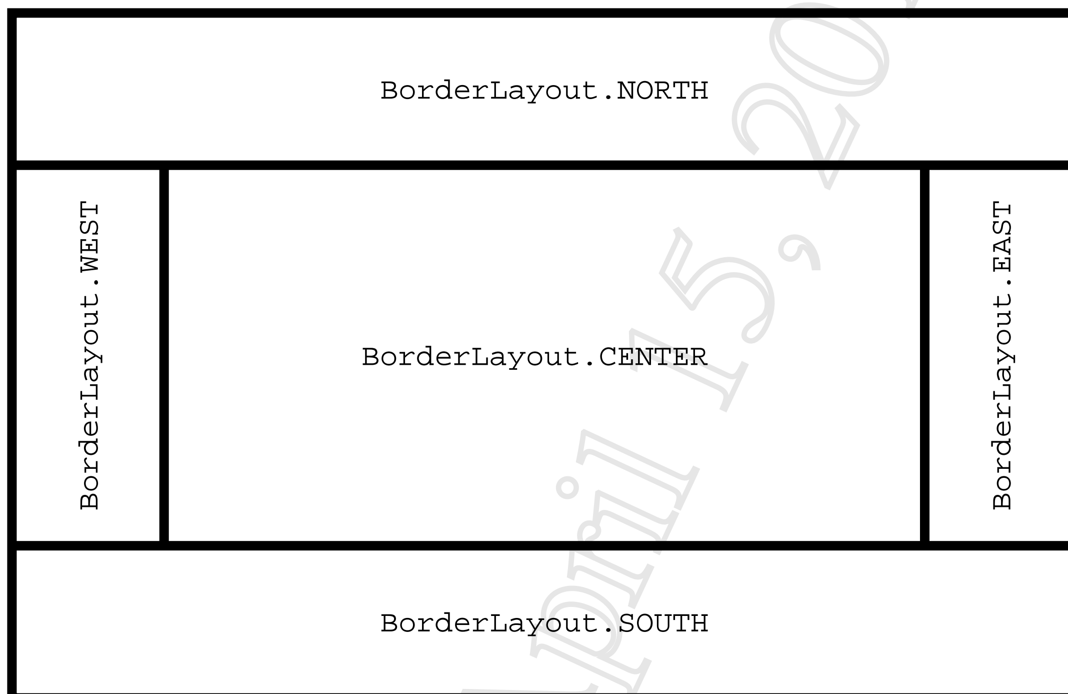

17.18 GUI API: Listeners (page 254)

Java uses a **listener** model for the processing of **graphical user interface** (**GUI**) **event**s. When something happens that needs dealing with, such as the end user pressing a GUI button, the **GUI event thread** creates an **object** representing the event before doing any processing that may be required. The event has an **event source**, which is some Java GUI object associated with the cause of the event. For example, an event created because the end user has pressed a button will have that button as its source. Each possible event source keeps a set of **listener** objects that have been registered as wishing to be 'told' if an event is created from that source. The GUI event thread processes the event by simply calling a particular **instance method** belonging to each of these listeners.

Let us consider an *abstract* example. Suppose we have some object that can be an event source,

DRAFT April 15, 2011

for example it might be a button. To keep it an abstract example, let us say it is an **instance** of `SomeKindOfEventSource`.

```
SomeKindOfEventSource source = new SomeKindOfEventSource(...);
```

Suppose also we wish events from that source to be processed by some code that we write. Let us put that in a **class** called `SomeKindOfEventListener` for this abstract example.

```
public class SomeKindOfEventListener
{
  public void processSomeKindOfEvent(SomeKindOfEvent e)
  {
    ... Code that deals with the event.
  } // processSomeKindOfEvent
} // class SomeKindOfEventListener
```

To link our code to the event source, we would make an instance of `SomeKindOfEventListener` and register it with the event source as a listener.

```
SomeKindOfEventListener listener = new SomeKindOfEventListener(...);

source.addSomeKindOfListener(listener);
```

The above code (or rather a concrete version of it) would typically be run in the **main thread** during the set up of the GUI. The following diagram illustrates the finished relationship between the source and listener objects.

DRAFT April 15, 2011

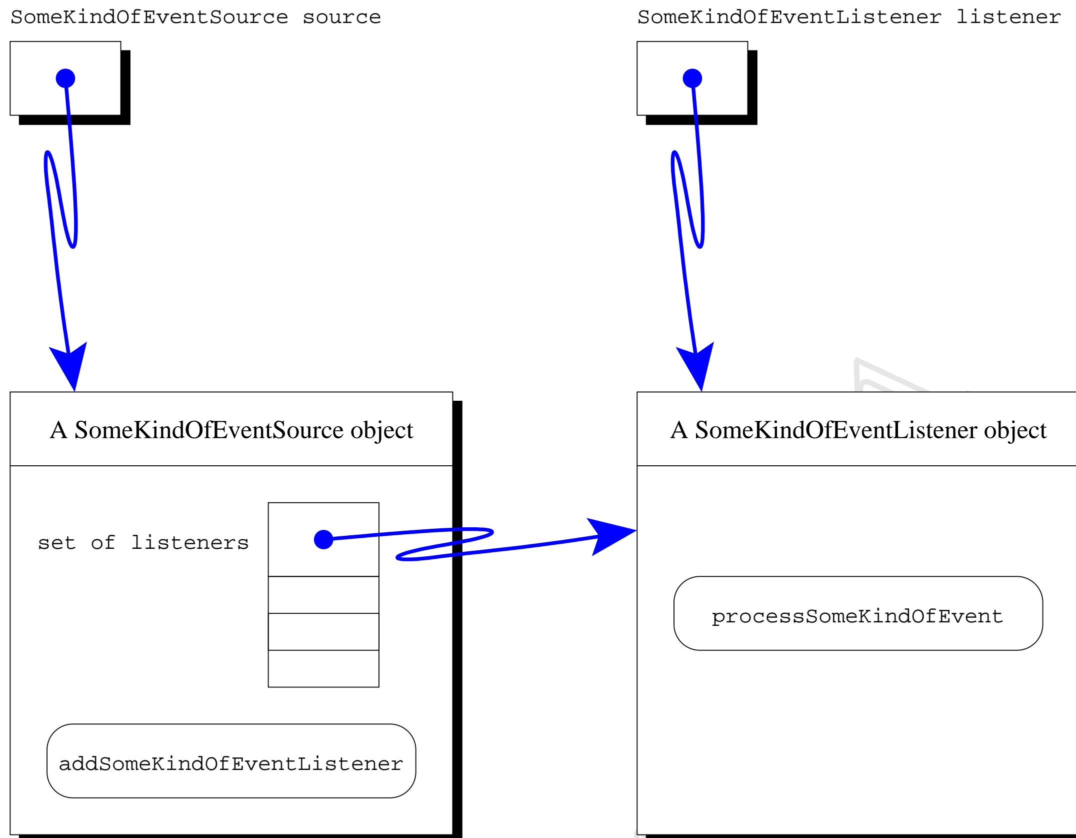

Now when an event happens, the GUI event thread can look at the set of listeners in the source object, and call the `processSomeKindOfEvent()` instance method belonging to each of them. So, when our `source` object generates an event, the `processSomeKindOfEvent()` instance method in our `listener` object is called.

Java Swing actually has several different kinds of listener for supporting different kinds of event. The above example is just an **abstraction** of this idea, so do *not* take the names `SomeKindOfEventSource`, `SomeKindOfEventListener`, `processSomeKindOfEvent` and `addSomeKindOfLi` literally – each type of event has corresponding names that are appropriate to it. For example, events generated by GUI buttons are known as `ActionEvents` and are processed by `ActionListener` objects which have an `actionPerformed()` instance method and are linked to the event source by an `addActionListener()` instance method.

17.19 GUI API: Listeners: ActionListener interface (page 257)

The standard **interface** called `java.awt.event.ActionListener` contains a body-less **instance method** which is called `actionPerformed`. The intention is that a full implementation of this instance method will contain code to process an **event** caused by the user doing something like pressing a **graphical user interface** (**GUI**) button.

DRAFT April

17.20 GUI API: Listeners: ActionListener interface: actionPerformed() (page 258)

After creating an **instance** of `java.awt.event.ActionEvent` when the end user has performed an 'action' such as pressing a button, the **GUI event thread** finds out from that **event source** which `ActionListener` **object**s have registered with it as wanting to be told about the **event**. The GUI event thread then invokes the **instance method** called `actionPerformed` belonging to each of those registered `ActionListeners`, passing the `ActionEvent` object as a **method argument**.

So, the heading of the `actionPerformed()` instance method is as follows.

```
public void actionPerformed(ActionEvent event)
```

Each implementation of the method will perform whatever task is appropriate as a response to the particular action in a particular program.

17.21 GUI API: JButton (page 256)

The **class** `javax.swing.JButton` implements a particular part of a **graphical user interface** (**GUI**) which offers a button for the end user to 'press' using the mouse. The text to be displayed on the button is specified as a `String` **method argument** to the `JButton` **constructor method**.

17.22 GUI API: JButton: addActionListener() (page 256)

The **class** `javax.swing.JButton` has an **instance method** called `addActionListener`. This takes as its **method parameter** an `ActionListener` **object**, and remembers it as being a **listener** interested in processing the **event** caused by an end-user pressing this button.

```
public void addActionListener(ActionListener listener)
{
  ... Remember that listener wants to be informed of action events.
} // addActionListener
```

17.23 GUI API: JButton: setEnabled() (page 266)

The **class** `javax.swing.JButton` has an **instance method** called `setEnabled`, which takes a **`boolean`** **method parameter**. If it is given the value **`false`**, the button becomes disabled, that is any attempt to press it has no effect. If instead the parameter is **`true`**, the button becomes enabled. When in the disabled state, the button will typically look 'greyed out'.

DRAFT April 15, 2011

17.24 GUI API: JButton: setText() (page 267)

The **class** `javax.swing.JButton` has an **instance method** called `setText` which takes a `String` and changes the text label displayed on the button, to the given **method argument**.

17.25 GUI API: ActionEvent (page 258)

When the **GUI event thread** detects that the end user has performed an 'action', such as pressing a button, it creates an **instance** of the **class** `java.awt.event.ActionEvent` in which it stores information about the **event**. For example, it stores a **reference** to the **event source object**, such as the button that was pressed.

17.26 GUI API: ActionEvent: getSource() (page 280)

The **class** `java.awt.event.ActionEvent` has an **instance method** called `getSource` which **return**s a **reference** to the **object** that caused the **event**.

17.27 GUI API: JTextField (page 265)

The **class** `javax.swing.JTextField` implements a particular part of a **graphical user interface** (**GUI**) which allows a user to enter a small piece of text. One of the **constructor method**s of the class takes a single **`int`** **method parameter**. This is the minimum number of **character**s of text we would like the field to be wide enough to display.

We can also use a `JTextField` to display a small piece of text generated from within the program.

17.28 GUI API: JTextField: getText() (page 265)

The **class** `javax.swing.JTextField` has an **instance method** called `getText` which takes no **method argument**s and **return**s the text contents of the text field, as a `String`.

17.29 GUI API: JTextField: setText() (page 265)

The **class** `javax.swing.JTextField` has an **instance method** called `setText` which takes a `String` as its **method argument** and changes the text of the text field to the given value.

DRAFT April 15, 2011

17.30 GUI API: JTextField: setEnabled() (page 267)

The **class** `javax.swing.JTextField` has an **instance method** called `setEnabled`, which takes a **`boolean`** **method parameter**. If it is given the value **`false`**, the text field becomes disabled, that is any attempt to type into it has no effect. If instead the parameter is **`true`**, the text field becomes enabled. When in the disabled state, the text field will typically look 'greyed out'.

17.31 GUI API: JTextField: initial value (page 274)

The **class** `javax.swing.JTextField` has a **constructor method** which takes a `String` **method parameter** to be used as the initial value for the text inside the text field.

```
JTextField nameJTextField = new JTextField("Type your name here.");
```

17.32 GUI API: JTextArea (page 267)

The **class** `javax.swing.JTextArea` implements a particular part of a **graphical user interface** (**GUI**) which displays a larger piece of text, consisting of multiple lines. The size of the text area can be specified as **method argument**s to the **constructor method**, as the number of rows (lines) and the number of columns (characters per line).

17.33 GUI API: JTextArea: setText() (page 269)

The **class** `javax.swing.JTextArea` has an **instance method** called `setText` which takes a `String` as a **method argument** and changes the text of the text area to the given value. This `String` may contain **new line character**s in it, and the text area will display the text appropriately as separate lines.

17.34 GUI API: JTextArea: append() (page 269)

The **class** `javax.swing.JTextArea` has an **instance method** called `append` which takes a `String` and appends it onto the end of the text already in the text area. Any required line breaks must be made by including explicit **new line character**s.

DRAFT April 15, 2011

17.35 GUI API: JPanel (page 270)

The **class** `javax.swing.JPanel` is an **extension** of the older `java.awt.Container`, which means that it is a component that is allowed to contain other components, and it has `add()` **instance method**s allowing us to add components to it. `JPanel` is designed to work well with the rest of the **Java Swing package**, and is the recommended kind of container to use when we wish to group a collection of components so that they are treated as one for layout purposes.

17.36 GUI API: JScrollPane (page 274)

The **class** `javax.swing.JScrollPane` implements a particular part of a **graphical user interface** (**GUI**) which provides a scrolling facility over another component.

The simplest way to use it is to invoke the **constructor method** which takes a GUI component as a **method parameter**. This creates a `JScrollPane` **object** which provides a scrollable view of the given component.

As an example, consider the following code which adds a `JTextArea` to the **content pane** of a `JFrame`.

```
Container contents = getContentPane();
contents.add(new JTextArea(15, 20));
```

To make the `JTextArea` scrollable, we would replace the above with the following code instead.

```
Container contents = getContentPane();
contents.add(new JScrollPane(new JTextArea(15, 20)));
```

17.37 GUI API: Color (page 400)

The **class** `java.awt.Color` implements colours to be used in **graphical user interface**s. Each `Color` **object** comprises four values in the range 0 to 255, one for each of the primary colours red, green and blue, and a fourth component (alpha) for opacity.

For convenience, the class includes a number of **class constant**s containing **reference**s to `Color` objects which represent some common colours.

```
public static final Color black     = new Color(0,     0,   0, 255);
```

DRAFT April 15, 2011

```
public static final Color white     = new Color(255, 255, 255, 255);
public static final Color red       = new Color(255,   0,   0, 255);
public static final Color green     = new Color(0,   255,   0, 255);
public static final Color blue      = new Color(0,   0,   255, 255);

public static final Color lightGray = new Color(192, 192, 192, 255);
public static final Color gray      = new Color(128, 128, 128, 255);
public static final Color darkGray  = new Color(64,   64,  64, 255);

public static final Color pink      = new Color(255, 175, 175, 255);
public static final Color orange    = new Color(255, 200,   0, 255);
public static final Color yellow    = new Color(255, 255,   0, 255);
public static final Color magenta   = new Color(255,   0, 255, 255);
public static final Color cyan      = new Color(0,   255, 255, 255);
```

Among many other features, there is an **instance method** `getRGB()` which **return**s a unique **`int`** for each **equivalent** colour, based on the four component values.

18 Interface

18.1 Interface (page 257)

An **interface** is like a **class**, except all the **instance method**s in it must have no bodies. It is used as the basis of a kind of contract, in the sense that it may be declared that some class **implement**s an interface. This means that it supplies full definitions for all the body-less instance methods listed in the interface. For example, the following code

```
public class MyClass implements SomeInterface
{
  ...
} // MyClass
```

says that the class being defined, `MyClass`, provides full definitions for all the instance methods listed in the interface `SomeInterface`. So, for example, if a **method** somewhere has a **method parameter** of **type** `SomeInterface`, then an **instance** of `MyClass` could be supplied as a corresponding **method argument**, as it satisfies the requirements of being of type `SomeInterface`.

18.2 Interface: definition (page 511)

An **interface** is like a **class**, except all the **instance method**s in it must be **abstract method**s, that is, they have no bodies. Only the **method interface**s are declared, i.e. the **method sig-**

DRAFT April 15, 2011

natures and **return type**s. The **method implementation**s must be provided by each non-**abstract class** that **implement**s the interface. For example, the following code says that the class `StopClock` is both a **subclass** of `JFrame` and implements `ActionListener`.

```
import java.awt.event.ActionListener;
import javax.swing.JFrame;

public class StopClock extends JFrame implements ActionListener
{
  ...
  public void actionPerformed(ActionEvent event)
  {
    ...
  } // actionPerformed
  ...
} // class StopClock
```

This means that an **instance** of `StopClock` is **polymorphic** – it **is a** `StopClock`, **is a** `JFrame` and also **is an** `ActionListener`.

The definition of an interface has the **reserved word** **`interface`** in its heading, instead of the reserved word **`class`**. It can contain a list of instance method headings, each of which has no body – just a semi-colon (`;`). If we wish, we can write the reserved word **`abstract`** in the heading of the interface, like we would for an **abstract class**. We can also write it in the instance method headings, like we would for abstract methods appearing in abstract classes. However, we are discouraged from doing so by the Java language standard[5], because all the instance methods *must* be abstract methods. Similarly, all the instance methods *must* be **public**, and so we do not need to write that visibility **modifier** either, and are discouraged from doing so.

The following is what you might expect the `ActionListener` interface to look like.

```
public interface ActionListener
{
  void actionPerformed(ActionEvent e);
} // interface ActionListener
```

An interface cannot contain **constructor method**s nor **class method**s (**`static`** **method**)s. What is more, if it has any **variable**s defined, they must be **public**, **`static`** and **final variable**s, although we can omit those modifiers if we wish.

There can be no **private** instance methods or variables in an interface – obviously.

DRAFT April 15, 2011

18.3 Interface: is a type (page 512)

An **interface** is a **type**, the **set** of all (**reference**s to) **object**s that can be created which are **instance**s of any **class** that **implement**s the interface. It has operations, which are the **method implementation**s of the **instance method**s of the interface, provided by each class which implements the interface. And each of these operations has an **operation interface**, which is the **method interface** defined in the interface (and, in effect, redefined in each class that implements the interface).

So, an interface defines only the operation interfaces of the type, not the actual operations. That is why this code construct is called an *interface*. We can think of it as being an **interface contract** – any class that claims to implement it is obliged to supply operation implementations.

18.4 Interface: method implementation (page 513)

A non-**abstract class** which **implement**s an **interface** must supply **method implementation**s for the **abstract method**s defined in that interface. As when making an **override** of an **instance method** defined in a **superclass**, there is a danger of getting a **method parameter type** wrong, and introducing an **overloaded method** instead, or mistyping the method name. The **override annotation**, `@Override`, introduced in Java 5.0, was extended in Java 6.0 to enable us to tell the **compiler** that we believe an instance method is an override or an implementation of one from a superclass *or* an interface. One situation this detects is when we have indeed got the method implementation correct, but forgot to say that our **class** implements the interface we had in mind!

18.5 Interface: generic interface (page 520)

A **generic interface** is an **interface** which has one or more **type parameter**s written within angled brackets (`<>`) just after its name in the interface heading. Such type parameters may be used as **type**s in the declaration of the **abstract method**s in the interface. The feature works in the same way as for **generic class**es: the generic interface itself is a **raw type** and when we supply **type argument**s for the type parameters, we identify a **parameterized type**.

18.6 Interface: extending another interface (page 526)

An **interface** can **extend** another interface. This means that the **abstract method**s and **class constant**s specified in the **superinterface** are **inherit**ed in the **subinterface**. From a **polymorphism** point of view, it also means that **instance**s of a **class** which **implement**s the subinterface, are members of the **type** denoted by the superinterface in addition to the type denoted by the subinterface.

DRAFT April 15, 2011

Unlike **class**es which can only extend one other class, interfaces can extend many other interfaces.

18.7 Interface: a class can implement many interfaces (page 530)

A **class** can **extend** at most one other class, but is permitted to **implement** any number of **interface**s. The implemented interfaces are listed, with commas between, after the **reserved word `implements`**.

For example, we could imagine a `StopClock` program, which automatically stopped and started the clock when the mouse is moved out of and back in to the window. This would probably implement `MouseListener` as well as `ActionListener`.

```
import java.awt.ActionListener;
import java.awt.MouseListener;
import javax.swing.JFrame;
...
public class StopClock extends JFrame
                       implements ActionListener, MouseListener
{
  ...
  // actionPerformed is defined in the interface ActionListener
  public void actionPerformed(ActionEvent event)
  {
    ...
  } // actionPerformed

  ... Various methods here, as specified in MouseListener.

} // class StopClock
```

19 Array

19.1 Array (page 286)

An **array** is a fixed size, ordered collection (**list**) of items of some particular **type**. The items are stored next to each other in **computer memory** at **run time**. As an example, the following is a representation of an array of 8 `int` values, which happen to be the first 8 **prime number**s (excluding 1).

DRAFT April 15, 2011

2	3	5	7	11	13	17	19

Each box, or **array element**, contains a value, which can be changed if desired. In other words, each element is a separate **variable**. At the same time, the array as a whole is a single entity. This is rather similar to the idea of an **object** having **instance variable**s, except that the elements of an array must all be of the same type.

Indeed, arrays in Java *are* **object**s.

19.2 Array: array creation (page 287)

We can create an **array** in Java using the **reserved word** **`new`**, like we do with other **object**s. However, instead of following this with the name of a **class**, we can state the **array base type** and then, in square brackets, the size of the array. For example, the following creates an array of ten `double` values.

```
new double[10]
```

At **run time**, this code yields a **reference** to the **new**ly created array, which we typically would want to store in a **variable**.

```
double[] myFingerLengths = new double[10];
```

Thanks to the use of references, the size of an array does not need to be known at **compile time**, because the **compiler** does not need to allocate memory for it. This means at **run time** we can create an array which is the right size for the actual **data** being processed.

```
int noOfEmployees = Integer.parseInt(args[0]);

String[] employeeNames = new String[noOfEmployees];
```

19.3 Array: array creation: initializer (page 320)

When we declare an **array variable** we can at the same time create the actual array by listing the **array element**s which are to be placed in it, using an **array initializer**. This is *instead*

DRAFT April 15, 2011

of saying how big the array is. Java counts this **list**, creates an array that big, and assigns the elements in the order listed. For example, the following code declares an **array variable** which refers to an array containing the first eight **prime number**s (excluding 1).

```
int[] smallPrimes = {2, 3, 5, 7, 11, 13, 17, 19};
```

This is just a shorthand for the following.

```
int[] smallPrimes = new int[8];
...
smallPrimes[0]=2;  smallPrimes[1]=3;  smallPrimes[2]=5;
smallPrimes[3]=7;  smallPrimes[4]=11; smallPrimes[5]=13;
smallPrimes[6]=17; smallPrimes[7]=19;
```

19.4 Array: element access (page 288)

The **array element**s in an **array** can be accessed individually via an **array index**. This is a whole number **greater than or equal** to zero. The first element in an array is indexed by zero, the second by one, and so on. To access an element, we write a **reference** to the array, followed by the index within left and right square brackets.

For example, assuming we have the array

```
double[] myFingerLengths = new double[10];
```

and somehow we have placed the lengths of my fingers and thumbs into the ten elements of `myFingerLengths`, then the following code would compute the total length of my fingers and thumbs.

```
double myTotalFingerLength = 0;
for (int index = 0; index < 10; index++)
  myTotalFingerLength += myFingerLengths[index];
```

So, arrays are a bit like ordinary **object**s with the array elements being **instance variable**s, except that the number of instance variables is chosen when the array is created, they are all the same **type**, they are ‘named’ by indices rather than names, and they are accessed using a different **syntax**.

DRAFT April 15, 2011

19.5 Array: element access: in two-dimensional arrays (page 330)

Each grid element in a **two-dimensional array** is indexed by two indices – the first **array index** accesses the row **array**, and the second accesses the **array element** within that row. For example, given the code

```
int[][] my2DArray = new int[5][4];
```

then my2DArray[0] is a **reference** to the first row, and so my2DArray[0][0] is the first element in the first row. Similarly, my2DArray[4][3] is the last element in the last row.

19.6 Array: length (page 292)

Every **array** in Java has a **public instance variable** called length, of **type** **int**, which contains the **array length** or size of the array. It is, of course, a **final variable**, so we cannot change its value.

```
int[] myArray = new int[25];
int myArrayLength = myArray.length;
```

In the above code fragment, the **variable** myArrayLength will have the value 25.

19.7 Array: empty array (page 292)

When we create an **array** we say how many **array element**s it should have, and this number can be zero. Although such an **empty array** may not seem of much use, it still exists – we can access its **array length** for example.

```
int[] myEmptyArray = new int[0];
System.out.println(myEmptyArray.length);
```

The above code will output zero, whereas the following code will cause a **run time error** (in fact a NullPointerException), because there is no array so we cannot ask for its length.

```
int[] myNonArray = null;
System.out.println(myNonArray.length);
```

DRAFT April 15, 2011

19.8 Array: of objects (page 301)

An **array** can contain values of any **type**, including **object**s. Of course, as with any other kind of **variable**, the **array element**s of an array with an **array base type** which is a **class**, actually contain **reference**s to the objects.

The most obvious example of an array of objects, is the **command line argument**s passed to the **main method**.

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

The following diagram shows the above **method parameter** referring to an array, with the array elements themselves referring to String objects.

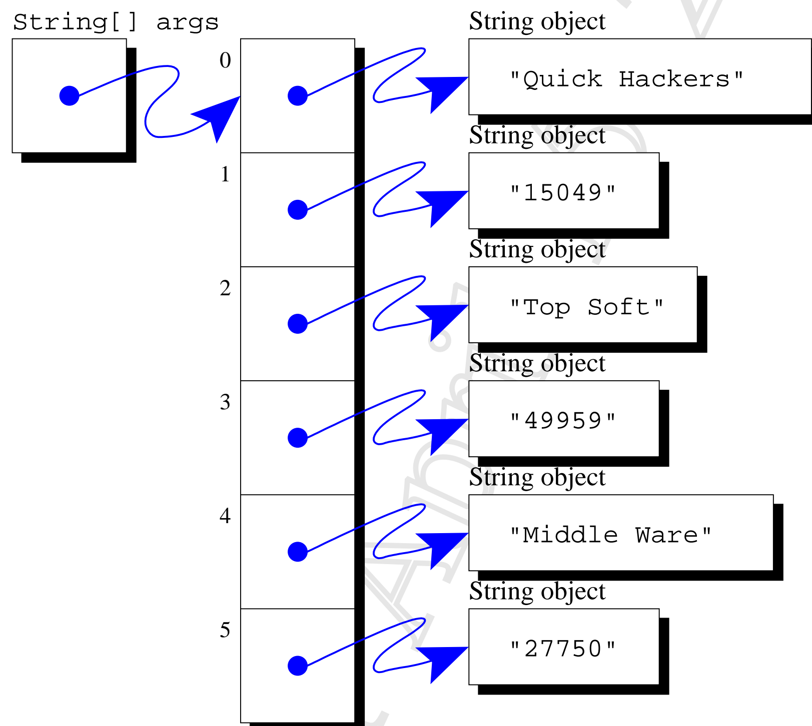

19.9 Array: partially filled array (page 310)

An **array** has a fixed size, specified when it is created. A **partially filled array** is one in which not all of the **array element**s are used, only a leading portion of them. The size of this portion is typically stored in a separate **variable**.

For example, suppose we have an array of 100 elements, of which initially none are in use.

DRAFT April 2011

```
private final int MAX_NO_OF_ITEMS = 100;
private int noOfItemsInArray = 0;
private SomeType[] anArray = new SomeType[MAX_NO_OF_ITEMS];
```

We can add another item into the array, or do nothing if it is full, as follows.

```
if (noOfItemsInArray < MAX_NO_OF_ITEMS)
{
  anArray[noOfItemsInArray] = aNewItem;
  noOfItemsInArray++;
} // if
```

19.10 Array: partially filled array: deleting an element (page 404)

The simplest way to delete an **array element** from a **partially filled array** with an arbitrary order, is to replace the unwanted item with the one at the end of the used portion and decrement the count of items.

```
int indexToBeDeleted = ...
noOfItemsInArray--;
anArray[indexToBeDeleted] = anArray[noOfItemsInArray];
```

19.11 Array: array extension (page 311)

If we are using a **partially filled array** then we may need to worry about the problem of it becoming full when we still wish to add more items into it. The principle of **array extension** deals with this by making a **new**, bigger **array** and copying items from the original into it.

We start by making an array of a certain size, with no items in it.

```
private static final int INITIAL_ARRAY_SIZE = 100;
private static final int ARRAY_RESIZE_FACTOR = 2;
private int noOfItemsInArray = 0;
private SomeType[] anArray = new SomeType[INITIAL_ARRAY_SIZE];
```

When we come to add an item, we make a bigger array if required.

```
if (noOfItemsInArray == anArray.length)
{
```

DRAFT April 15, 2011

```
    SomeType[] biggerArray
      = new SomeType[anArray.length * ARRAY_RESIZE_FACTOR];
    for (int index = 0; index < noOfItemsInArray; index++)
      biggerArray[index] = anArray[index];
    anArray = biggerArray;
  } // if

  anArray[noOfItemsInArray] = aNewItem;
  noOfItemsInArray++;
```

The new array does not need to be twice as big as the original, just at least one element bigger. However, increasing the size by only one at a time would be slow due to the need for copying the existing elements across.

19.12 Array: shallow copy (page 314)

When we copy an **array** containing **reference**s to **object**s, we can either make a **shallow copy** or a **deep copy**. A shallow copy contains the same references, so the objects end up being shared between the two arrays. A deep copy contains references to *copies* of the original objects.

19.13 Array: array of arrays (page 329)

The **array element**s of an **array** may be of any **type**, including arrays. This means the elements of the array are **reference**s to other arrays. For example, the following diagram shows an **array variable** which contains a reference to an array of arrays of `int` values.

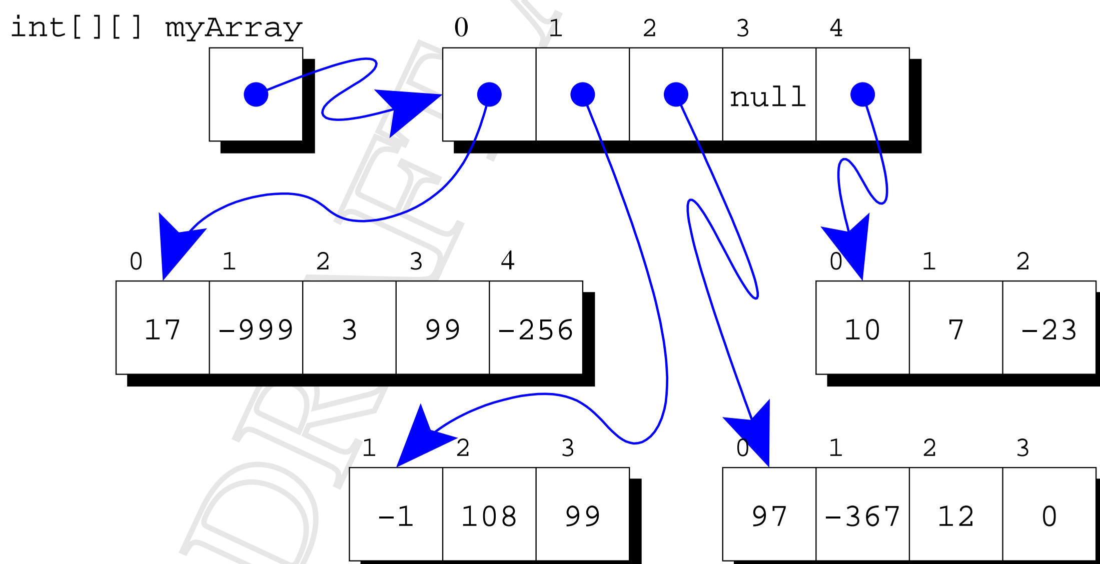

DRAFT April 15, 2011

The type of the variable is `int[][]`, that is **`int`** array, array. The variable references an array of 5 values, the first of which is a reference to an array of 5 numbers, the second a reference to an array of 3 numbers, the third is a reference to an array of 4 numbers, the fourth is the **null reference** and the final element is a reference to an array of 3 numbers. These arrays could be created, ready for the numbers to be put in them, as follows.

```
int[][] myArray = new int[5][];
myArray[0] = new int[5];
myArray[1] = new int[3];
myArray[2] = new int[4];
myArray[3] = null;
myArray[4] = new int[3];
```

19.14 Array: array of arrays: two-dimensional arrays (page 330)

A very common situation when we have an **array** of arrays, is that none of the **array element**s are the **null reference** and all of the arrays they **reference** are the same length. This is known as a **two-dimensional array**, and is essentially a model of a rectangular grid. For example, the following diagram shows a **variable** which contains a reference to a two-dimensional array of **`int`** values.

DRAFT April 15, 2011

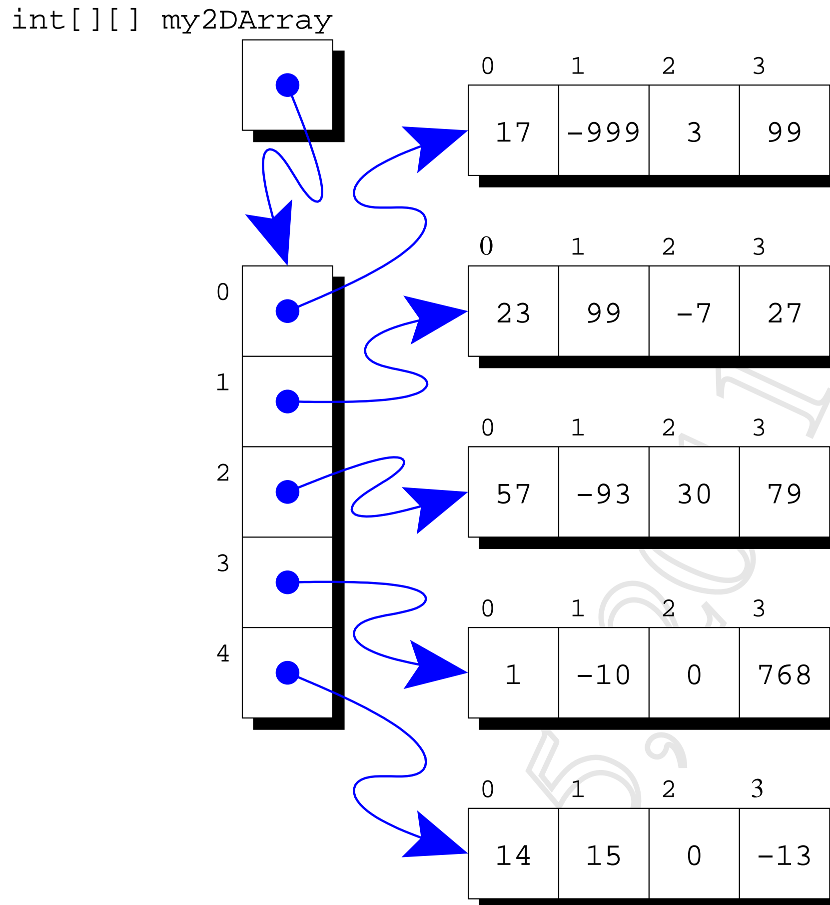

The above two-dimensional array could be created (without the numbers being assigned into it yet) by the following code.

```
int[][] my2DArray = new int[5][];
my2DArray[0] = new int[4];
my2DArray[1] = new int[4];
my2DArray[2] = new int[4];
my2DArray[3] = new int[4];
my2DArray[4] = new int[4];
```

Two-dimensional arrays are so common, that Java provides a shorthand notation for defining them. The shorthand for the above example is as follows.

```
int[][] my2DArray = new int[5][4];
```

The code `new int[5][4]` makes an array of length 5 get created at **run time**, and also 5 arrays of length 4, which are capable of holding `int` values, with these latter 5 arrays being referenced by the 5 elements in the first array.

DRAFT April 5, 2019

20 Exception

20.1 Exception (page 340)

A **run time error** is called an **exception** in Java. There is a standard **class** called `java.lang.Exception` which is used to record and handle exceptions. When an exceptional situation happens, an **instance** of this class is created, containing information about the error, stored in its **instance variables**. In particular, it includes a **stack trace** containing the source line number, **method** name and class name at which the error occurred. This stack also contains the same information for the method that called the one that failed, and so on, right back up to the main method (for an error occurring in the **main thread**).

20.2 Exception: getMessage() (page 345)

When an **instance** of `java.lang.Exception` is created, it may be given a text message helping to describe the reason for the error. This may be retrieved from an `Exception` **object** via its `getMessage()` **instance method**.

20.3 Exception: there are many types of exception (page 347)

The **class** `java.lang.Exception` is a general model of **exceptions**. Java also has many classes for modelling exceptions which are more specific to a particular kind of error. Here are a few of the ones from the `java.lang` **package**, each listed with an example error situation which causes an **instance** of the exception class to be created.

DRAFT April 15, 2011

Exception class	**Example use**
`ArrayIndexOutOfBoundsException`	When some code tries to access an **array element** using an **array index** which is not in the range of the **array** being indexed.
`IllegalArgumentException`	When a **method** is passed a **method argument** which is inappropriate in some way.
`NumberFormatException`	In the `parseInt()` method of the `java.lang.Integer` class when it is asked to interpret an invalid `String` method argument as an **`int`**. (Actually, `NumberFormatException` is a particular kind of the more general `IllegalArgumentException`.)
`ArithmeticException`	When an **integer division** has a denominator which is zero.
`NullPointerException`	When we have code that tries to access the **object referenced** by a **variable**, but the variable actually contains the **null reference**.

20.4 Exception: creating exceptions (page 350)

The standard **class** `java.lang.Exception` has a number of **constructor methods** enabling us to create **instances** of it. One of these takes no **method arguments**, and creates an `Exception` that has no message associated with it. A second constructor method takes a `String` which is to be used as the message. The other kinds of **exception**, such as `ArrayIndexOutOfBoundsException`, `IllegalArgumentException`, `NumberFormatException`, `ArithmeticException` and `NullPointerExcept`
also have these two constructor methods.

20.5 Exception: creating exceptions: with a cause (page 357)

The standard **class** `java.lang.Exception` also has two more **constructor methods** enabling us to create **instances** which know about another **exception** that caused this one to be created. One of these takes the message and the **exception cause**, the other just takes the cause (and hence has no message). Whenever we **throw** a **new** exception inside a **catch clause**, it is good practice to include the caught exception as the cause of the new one.

Many of the other kinds of exception also have these two constructor methods.

DRAFT April 15, 2011

20.6 Exception: getCause() (page 366)

The **exception cause** stored inside an `Exception` may be retrieved via its `getCause()` **instance method**. This will **return** the **null reference** if no cause was given.

20.7 Exception: inheritance hierarchy (page 434)

All **exception**s in Java are modelled as **instance**s of **class**es. For example, the class `java.lang.Exception` models a very general idea of exception, and `java.lang.ArrayIndexOutOfBoundsException` a much more specific kind. The different kinds of exception are arranged in an **inheritance hierarchy**, with those classes near the top being models of quite general exceptions, and those at the bottom being very specific. An instance of `ArrayIndexOutOfBoundsException` is created when an **array index** is out of the legal range for the **array**. This class is a **subclass** of the more general `java.lang.IndexOutOfBoundsException`. A different subclass of `IndexOutOfBoundsException` is called `java.lang.StringIndexOutOfBoundsException`. Instances of this are created in circumstances such as supplying an illegal **method argument** to the `charAt()` **instance method** of a `String`. The class `IndexOutOfBoundsException` is itself a subclass of `java.lang.RuntimeException`, the kind of exception that Java does not *require* us to **catch**, although we sometimes do, and this class is a subclass of `Exception`.

We can show this relationship in a **UML class diagram**, including the **constructor method**s and some of the **public** instance methods.

DRAFT April 15, 2011

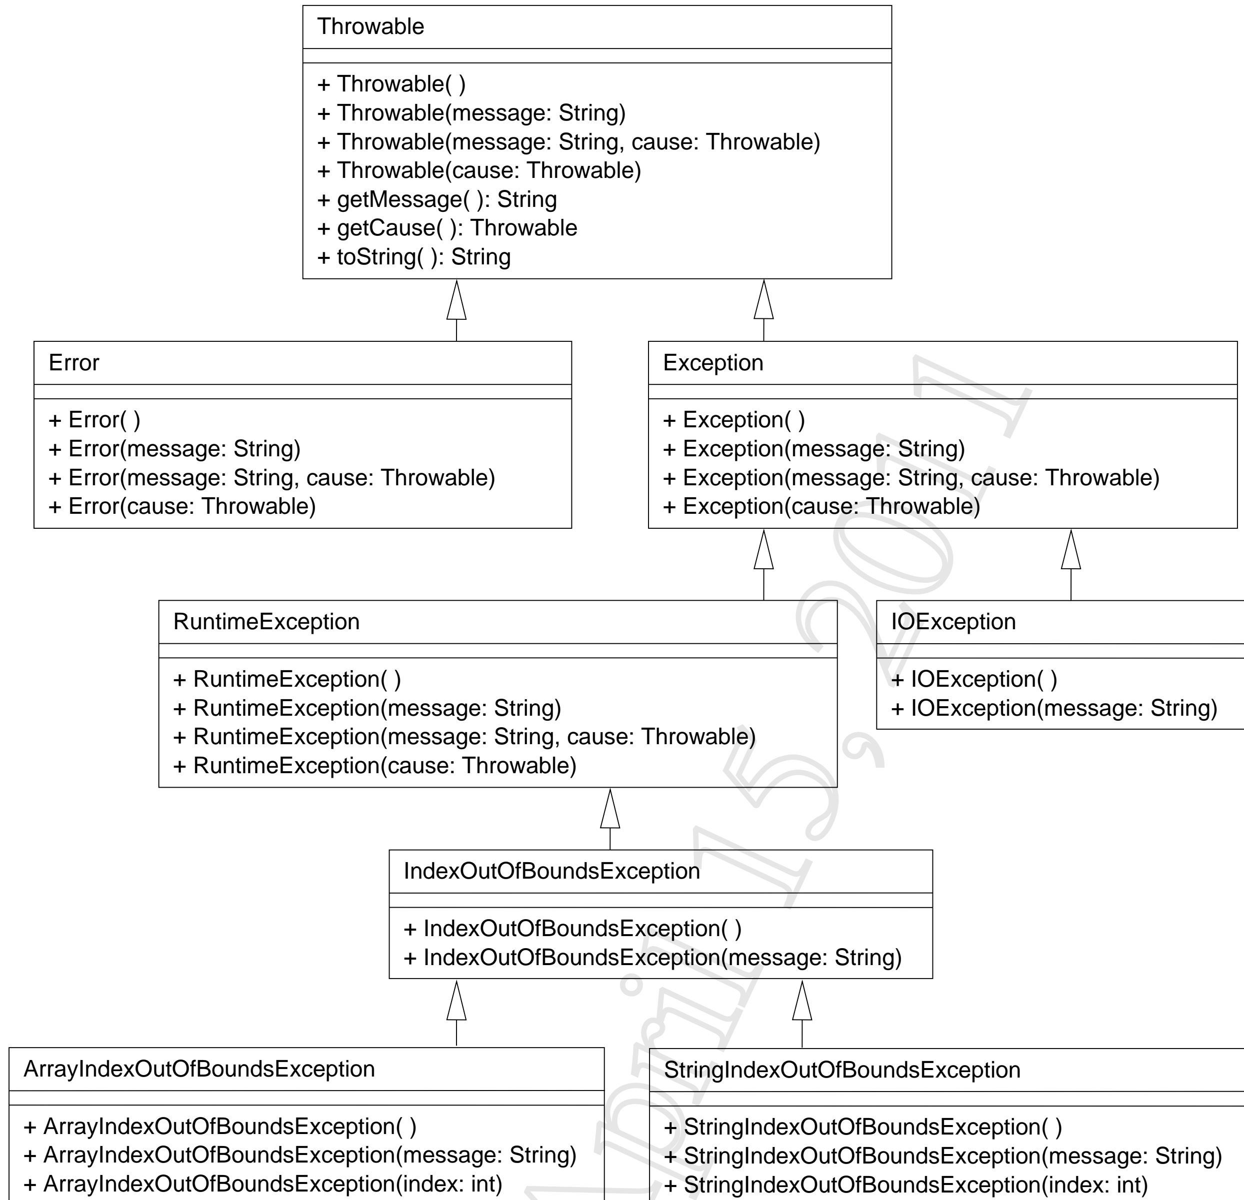

DRAFT April 15, 2011

You can see that `Exception` is itself a subclass of something even more general called `java.lang.Throwable`, and there is a separate subclass of `Throwable` called `java.lang.Error`. The class `Throwable` is the **type** of all **object**s that can be **throw**n and handled by catches of a **try statement**. `Error` is the type of `Throwable`s which represent such serious conditions, that it is unlikely a program would bother trying to catch them. For example, `java.lang.OutOfMemoryError` is a subclass of `Error`, and an instance of it is thrown when the **virtual machine** has run out of memory to create any more objects. Catching this kind of condition is unlikely to be helpful in most situations, and so Java does not force us to. They are examples of **unchecked exception**s. However, ultimately the programmer knows best, so `Error`s *can* be caught if desired.

`Exception` is the type of `Throwable` which represents conditions that should typically be caught at some point. If a **method** contains code that could cause an `Exception`, or one of its subclasses, to be thrown, then the **compiler** forces the exception to either be caught within

the method, or declared in the **throws clause** of the method – they are **checked exception**s.

However, the `RuntimeException` class (and its subclasses) represents the kind of possible exception which programmers usually avoid in the first place. For example, when **loop**ing an array index over an array, the code would probably be written to use the correct values, and so avoid an `ArrayIndexOutOfBoundsException` exception. It would be highly inconvenient to *have* to write a **catch clause** or a throws clause even though we know the exceptions are avoided, and so Java relaxes the rule for this subclass – they too are **unchecked exception**s. Of course, this means we must discipline ourselves: especially in code intended for **software reuse**, we *should* write catch or throws clauses if we have not eliminated the possibility of these exceptions!

The diagram above is only a sample. There are almost 80 direct subclasses of `Exception` in the standard classes in Java 6.0, including `java.io.IOException` – instances of that can be thrown when processing **file**s. There are nearly 50 direct subclasses of `RuntimeException`.

One advantage of this **inheritance hierarchy** is that when we catch exceptions, we can decide how general or specific we need to be. For example, the following fragment of code could cause an `ArrayIndexOutOfBoundsException` to be thrown in some circumstances, and in other cases a `StringIndexOutOfBoundsException`.

```
int arrayIndex, stringIndex;
String[] listOfStrings;

... Code here to populate the above array,
... and set arrayIndex and stringIndex.

char c = listOfStrings[arrayIndex].charAt(stringIndex)
```

We can catch any exceptions of type `ArrayIndexOutOfBoundsException`, caused by `arrayIndex` having a bad value. Alternatively we can catch exceptions caused by the value of `stringIndex` being unsuitable, that is `StringIndexOutOfBoundsException` exceptions. If we wish, we can have two **catch clause**s, one for each. However, the exception inheritance hierarchy allows us the option of having one catch clause to deal with both, if that is appropriate, by catching `IndexOutOfBoundsException`.

20.8 Exception: making our own exception classes (page 435)

Another advantage of **exception**s being arranged in an **inheritance hierarchy** is that we can easily make our own exception **class**es. Sometimes, the leaf classes at the bottom of the standard exception inheritance hierarchy tree are not quite specific enough to suit the errors that can occur in our own code. They are, obviously, **design**ed to be appropriate to the standard classes. So, whenever we wish to **throw** an exception, we should ask ourselves whether there

DRAFT April 15, 2011

is a standard exception that nicely captures the meaning of the error, and if not, we should make our own exception class that does.

Making a new exception class is very easy. All we need to do is choose one of the standard classes which is closest to characterizing what we want, and make a **subclass** of it. Often this standard class will be either `java.lang.Exception` itself or `java.lang.RuntimeException`. We would choose the former if we want ours to be **checked exception**s, or the latter if we want them to be **unchecked exception**s because we believe the circumstances leading to them can be and typically should be avoided.

Most often, our own exception classes contain nothing but four **constructor method**s, one with no **method parameter**s, one which takes a `String` for the message associated with the exception, one which has both a message and a `Throwable` **exception cause**, and one which has only a cause. These simply invoke the corresponding constructor method from the **superclass**.

21 Inheritance

21.1 Inheritance (page 373)

A **class** can be used to model a category of **object**s with certain characteristics that exist in some way in the requirements of the program. However, sometimes the requirements exhibit sub-categories of objects. For example, a program which is designed to simulate traffic movement to help with road planning would probably have a class called `Vehicle`, representing the category of all road vehicles. This would contain properties which are common to all vehicles, such as average speed, and the relationship between their position and traffic lights, etc.. Sub-categories of vehicle might be bicycle, private car, taxi, bus, lorry etc.. These all have different specific properties – for example bicycles can be secured to many suitable fixed objects, such as railings and of course bicycle stands, whereas cars need car parks and metered side streets, etc.. Lorries need specific access and unloading points at specific places, such as shops that require regular deliveries. The road simulation would probably want to model people wishing to move about on the roads, and in this respect, bicycles, private cars, taxis and buses have a current and maximum number of passengers. Lorries might instead have a current and maximum load capacity. The behaviour of taxis and buses respectively link to the properties of taxi ranks and bus stops. And so on.

We would want to model these sub-categories as separate classes, each with whatever properties they specifically need, and yet still model the idea that they are all vehicles with the general properties. In **object oriented programming** we signify this relationship by having **superclass**es and **subclass**es. A superclass is something which models the general category of certain objects, and a subclass models a sub-category of those objects. So, we might decide that `Vehicle` is the superclass of all road vehicles, and that the class `Bicycle` models the sub-category of bicycles, and have the classes `PrivateCar`, `Taxi`, `Bus`, `Lorry`, etc. for the other specific sub-categories.

DRAFT April 15, 2011

By saying that a class is a subclass of another, its superclass, we are modelling the **is a** relationship. So, in the above example, a bicycle **is a** vehicle, that is, an **instance** of `Bicycle` is also an instance of `Vehicle`.

The relationship between superclasses and their subclasses is known as **inheritance** because the subclasses **inherit** the general properties from the superclass, as well as adding any specific properties of their own.

21.2 Inheritance: a subclass extends its superclass (page 378)

A **subclass** is said to be an **extension** of its **superclass**, because, in addition to **inherit**ing the properties of the superclass, it may have more properties that the superclass does not have. We state the relationship by declaring in the heading for the subclass that it **extend**s the superclass. For example, in a program to simulate traffic flow we might have the following.

```
public class Bicycle extends Vehicle
{
  ...
  public void chainToRailings(Railings railings)
  {
    ...
  } // chainToRailings
  ...
} // class Bicycle
```

So a `Bicycle` **object** has all the properties of a `Vehicle`, but also has the feature of being able to be chained to railings.

As well as being used to represent **is a** relationships between the model **class**es of our programs, subclasses are commonly used in the **graphical user interface** parts of our programs. For example, the following says that the `HelloWorld` class is a subclass of the `javax.swing.JFrame` class. This means `HelloWorld` is an extension of `JFrame`, that is, an **instance** of `HelloWorld` is a `JFrame` **object** too, but with extra properties that a plain `JFrame` object does not have.

```
import javax.swing.JFrame;
public class HelloWorld extends JFrame
{
  ... Code to add a JLabel with the text "Hello World!" in it.
} // class HelloWorld
```

DRAFT April 15, 2011

21.3 Inheritance: invoking the superclass constructor (page 379)

In the body of the **constructor method** of a **subclass** we typically start by invoking a constructor method of its **superclass**. This is done by writing the **reserved word** `super` followed by the appropriate **method argument**s in brackets. Such a **superclass constructor call** must be the first **statement** in the body of the constructor method, and furthermore, the superclass must have a constructor method which matches the supplied arguments.

For example, in a traffic flow simulation program, when a vehicle is added to the simulation it probably would always be given a position, direction and current speed.

```
public class Vehicle
{
  ...
  public Vehicle(Position requiredPosition,
                 Direction requiredDirection, Speed requiredSpeed)
  {
    ... Code that does something with requiredPosition,
    ... requiredDirection and requiredSpeed.
  } // Vehicle
  ...
} // class Vehicle
```

Instead of creating plain `Vehicle` **object**s we would make **instance**s of a subclass, such as `Bicycle`. We would still supply the position, direction and current speed information to the constructor method of `Bicycle`, and it would most likely simply pass it on to the constructor method of `Vehicle`.

```
public class Bicycle extends Vehicle
{
  ...
  public Bicycle(Position position, Direction direction, Speed speed)
  {
    super(position, direction, speed);
    ... Code specific to making a Bicycle, if any, goes here.
  } // Bicycle
  ...
} // class Bicycle
```

21.4 Inheritance: invoking the superclass constructor: implicitly (page 423)

In the body of a **constructor method**, if the first **statement** is not a **superclass constructor call** (using `super`), nor is it an **alternative constructor call** (using `this`), then a call to the

DRAFT April 15, 2011

constructor method of the **superclass** which has no **method argument**s, is assumed. This is because the first work which is done by a constructor method must be to actually create the **object**, that is, allocate memory for it, and this is done inside the constructor method of the `java.lang.Object` **class**.

21.5 Inheritance: overriding a method (page 380)

The **instance method**s of a **superclass** are **inherit**ed by its **subclass**es. Sometimes, the definition of an instance method needs to be changed in a subclass, in which case the subclass simply redefines it. The subclass version **override**s the inherited definition. To override an instance method, the redefinition must have the same name and **type**s of **method parameter**s otherwise it is a definition of a different **method**! It must also still be an instance method, and have the same **return type**. (Actually, the return type of the new instance method can be a subclass of the return type of the one in the superclass.)

For example, in a traffic flow simulation program, most kinds of vehicle probably perform an emergency stop in much the same way. However, a bicycle probably does it differently to most.

```
public class Vehicle
{
  ...
  public void emergencyStop()
  {
    ... General code for most vehicles.
  } // emergencyStop
  ...
} // class Vehicle


public class Bicycle extends Vehicle
{
  ...
  public void emergencyStop()
  {
    ... Specific code for bicycles.
  } // emergencyStop
  ...
} // class Bicycle
```

21.6 Inheritance: overriding a method: @Override annotation (page 430)

Java 5.0 introduced an idea called **annotation**s. These allow us to provide additional information to the **compiler** which can then be used to help in various ways. In particular, the **over-**

DRAFT April 15, 2011

ride annotation, `@Override`, can be written immediately before the heading of an **instance method** that we believe **override**s one from the **superclass**, or is a **method implementation** of an **abstract method** in the superclass. The compiler will complain if this is not the case, thus protecting us from accidentally getting the **method signature** wrong – perhaps misspelling the method name or mis-ordering the **method parameter type**s and creating an **overloaded method**, etc..

21.7 Inheritance: abstract class (page 385)

If we wish that no **instance**s of a particular **class** should be made, we can declare it as an **abstract class**. This is done by including the **reserved word** **`abstract`** before the word **`class`** in its heading. The **compiler** will produce an error if any code attempts to create a direct instance of an abstract class.

For example, in a program that simulates traffic flow, it is likely that we do not wish any direct instances of the class `Vehicle` to be made, only **subclass**es of it.

```
public abstract class Vehicle
{
  ...
} // class Vehicle

public class Bicycle extends Vehicle
{
  ...
} // class Bicycle
```

The following code would produce an error message from the compiler.

```
Vehicle v = new Vehicle(...);
```

Whereas this code would be allowed.

```
Bicycle b = new Bicycle(...);
```

21.8 Inheritance: abstract method (page 386)

An **abstract class** is permitted to have **abstract method**s declared in it. These are **instance method**s which have **modifier**s (such as **`public`** – but not **`static`**), **return type**, name and

DRAFT April 15, 2011

method parameters as usual, but also include the **reserved word `abstract`** and instead of a body defined within braces, the heading is followed by a semi-colon (`;`). This declares only the **method interface**, i.e. the **method signature** and **return type**, and not the **method implementation**.

For example, in a traffic flow simulation program, the abstract class `Vehicle` might have an abstract method that decides whether the vehicle can pass down a particular route. It may well be that each kind of vehicle needs to implement this in a different way.

```
public abstract class Vehicle
{
  ...
  public abstract boolean canPassDown(Route r);
  ...
} // class Vehicle
```

All **subclass**es of the abstract class must either provide a method implementation of all the abstract methods, or themselves be abstract classes. When we write an abstract method, we are saying that all (non-abstract) subclasses of the abstract class contain an instance method with the given method interface (name, method parameters and return type), but the implementations of the instance method are provided by the subclasses, rather than one being defined here. This saves us having to provide an implementation that is never used, in cases where *every* subclass would **override** it with their own version.

```
public class Bicycle extends Vehicle
{
  ...
  public boolean canPassDown(Route r)
  {
    ... Code for deciding if this bicycle can pass down the route.
  } // canPassDown
  ...
} // class Bicycle
```

When a subclass defines a non-abstract instance method which is also defined in its **superclass**, we say that it **override**s the one from the superclass. When it defines an instance method which is declared as an abstract method in its superclass, we say it provides a **method implementation**. We can think of an override as *replacing* the method implementation from the superclass.

21.9 Inheritance: polymorphism (page 390)

An **instance** of a **subclass** is also an instance of its **superclass**. For example, in a traffic flow simulation program, if the **class** `Bicycle` is a subclass of `Vehicle`, then an instance of

DRAFT April 15, 2011

`Bicycle` **is a** `Bicycle` and also it **is a** `Vehicle`. It may be treated as a `Bicycle`, because that is its **type**. However, it also may be treated as a `Vehicle` because that is also its type. It has *both* these forms. We say that it is **polymorphic**, which means 'has many forms'. Java supports **polymorphism** via the use of **inheritance**.

21.10 Inheritance: polymorphism: dynamic method binding (page 391)

In general, a **class** might have a **subclass**, which might **override** some of its **instance method**s. Also, **abstract method**s are designed to have different **method implementation**s in different subclasses. Thus, when the **compiler** produces the **byte code** for a **method call** on an instance method, it does not know which actual **method implementation** will get used – the same call could invoke different versions of the method at different moments, depending on the value of the **object reference** at **run time**.

For example, assume we have the class `Vehicle` with the instance method `emergencyStop()`, and subclass `PoshCar` that does not override it, and another subclass `Bicycle` that does. Which version of the method is called by the second line in the following code?

```
Vehicle funRide = Math.random() < 0.5 ? new PoshCar(...) : new Bicycle(...);
funRide.emergencyStop();
```

Only at run time can the answer be determined: the reference stored in `funRide` refers either to a `PoshCar` object, in which case the version from `Vehicle` is used, or a `Bicycle` object, in which case the version from `Bicycle` is used. The process of determining at run time which actual method to invoke is known as **dynamic method binding**.

As a programmer, we have to be aware of this principle, because it means that our code might not behave as we expected it to in some subclass where some of our instance methods have been replaced with ones that do something different to what we were expecting. Instance methods which are declared as **private** are safe – they cannot be overridden because they are not even visible in any subclass.

21.11 Inheritance: final methods and classes (page 391)

If we wish that no **subclass** may **override** a particular **public instance method**, we can declare it as a **final method** by including the **reserved word** **`final`** in its heading. This should be used with care – it may be that future requirements dictate that a subclass which has not yet been written needs its own version of the instance method, but it would not be able to have one without us removing the **`final`** **modifier** in the **superclass**.

Similarly, we can state that a **class** is a **final class** and cannot have any subclasses at all, by including **`final`** in the class heading.

DRAFT April 15, 2011

21.12 Inheritance: adding more object state (page 393)

A **subclass** is said to be an **extension** of its **superclass**, because in general it may add more properties that the superclass does not have. One way of **extend**ing is to add more **object state**, that is, additional **instance variable**s.

21.13 Inheritance: adding more instance methods (page 395)

Another way of **extend**ing the **superclass** in a **subclass** is to add more **instance method**s. This is especially likely to be desired if the subclass also has additional **instance variable**s.

21.14 Inheritance: testing for an instance of a class (page 397)

The **reserved word** **`instanceof`** is a **binary infix operator** which takes an **object reference** as its left **operand**, and a **class** name as its right operand. It yields **`true`** if the reference refers to an object which **is an** **instance** of the named class (including being an instance of a **subclass** of the named class), **`false`** otherwise.

For example, in a traffic flow simulation program, if the class `Tandem` is a subclass of `Bicycle` which is a subclass of `Vehicle`, then the following code might be found.

```
Vehicle vehicle = new Tandem(...);
... Code that might change what vehicle refers to.
if (vehicle instanceof Bicycle)
  ... Code that is only run if vehicle is still referring to a Bicycle,
  ... perhaps still the original Tandem.
```

21.15 Inheritance: casting to a subclass (page 397)

An **instance** of a **subclass** **is a**n instance of its **superclass** too. This means something which is of the subclass **type** can always be used wherever the superclass type is required. For example, in a traffic flow simulation program, if `Bicycle` is a subclass of `Vehicle`, then the following would be permitted.

```
Vehicle vehicle1 = new Bicycle(...);
```

However, obviously not every instance of a superclass is also an instance of a particular one of its subclasses, and so something of the superclass type cannot automatically be used where something of a subclass type is required.

DRAFT April 15, 2011

For example, the following is not permitted.

```
Vehicle vehicle1 = new Bicycle(...);
...
Bicycle bicycle1 = vehicle1;
```

The problem is in the last line – `vehicle1` is definitely of type `Vehicle`, but as far as Java is concerned, its value might not be of type `Bicycle`, and so a **compile time error** will result.

If we are convinced that it is safe to treat something of the superclass type as though it is of a particular subclass type, then we can **cast** the value to that subclass, by preceding the value with the name of the subclass in brackets. For example, the following is appropriate if we are sure that after the code represented as `...` has been **execute**d, the value of the **variable** `vehicle1` is still a **reference** to a `Bicycle` **object**.

```
Vehicle vehicle1 = new Bicycle(...);
...
Bicycle bicycle1 = (Bicycle)vehicle1;
```

The **compiler** will accept this on face value, but the type cast is checked at **run time**. If it turns out that the value being cast to a subtype is not a reference to an object of that type, then a `ClassCastException` object is **throw**n.

A common misunderstanding is that a **class** cast somehow changes the object that is being cast. Rather, it merely *checks* that the object is already of the stated type. This is in contrast to a **primitive type** cast, such as converting a `double` into an `int`, which really does create a new value from the old one.

21.16 Inheritance: is a versus has a (page 406)

When a **class**, `A`, is a **subclass** of another class, `B`, we say that an **object** of **type** `A` **is a** `B`.

If, on the other hand, a class, `C`, has an **instance variable** of type `D`, we say that an object of type `C` **has a** `D`.

21.17 Inheritance: using an overridden method (page 414)

A **subclass** can **override** an **instance method** defined in a **superclass**, but sometimes the behaviour of the new version is based on that of the one it is overriding. This means we need

DRAFT April 15, 2011

to have a **method call** to the *superclass* version, which we can do by prepending the instance method name with the **reserved word** **`super`** and a dot.

For example, in a traffic flow simulation program where most kinds of vehicle probably perform an emergency stop in much the same way, perhaps a bicycle's behaviour is based on the more general one.

```
public class Vehicle
{
  ...
  public void emergencyStop()
  {
    ... General code for most vehicles.
  } // emergencyStop
  ...
} // class Vehicle


public class Bicycle extends Vehicle
{
  ...
  public void emergencyStop()
  {
    ... Specific code for bicycles.
    super.emergencyStop();
    ... More specific code for bicycles.
  } // emergencyStop
  ...
} // class Bicycle
```

This `super.` notation can be used in any instance method of the subclass, not just in the overriding method.

21.18 Inheritance: constructor chaining (page 423)

Whenever a **constructor method** is invoked, the first thing done is either a call to another constructor method in the same **class**, or to a constructor method in the **superclass**. This in turn does the same, all the way up the **inheritance hierarchy** until eventually the constructor method of the `java.lang.Object` class is called. This process is known as **constructor chaining**.

Such chaining must always be possible for every class we write, or else we would not be able to have **object**s created at **run time** – it is the constructor method of `Object` that actually creates an object. So, one rule is that at least one constructor method of every class must *not* start with a call to another constructor method of the same class!

DRAFT April 15, 2011

21.19 Inheritance: multiple inheritance (page 509)

By saying that a **class** is a **subclass** of another we are modelling the **is a** relationship. Sometimes, it can appear natural to view a class as being a subclass of more than one **superclass**. This results in the subclass **inherit**ing properties from each of its superclasses, which is known as **multiple inheritance**.

Whilst the idea can sound attractive, it brings with it a complication when two or more of these superclasses contain an **instance method** with the same name and **method parameter**s. This problem is best illustrated by an abstract example. Suppose we have the class `Super1`, with the instance method `methodA()`.

```
public class Super1
{
  ...
  public void methodA()
  {
    ...
  } // methodA
  ...
} // class Super1
```

Suppose we, quite separately, have the class `Super2`, which also has an instance method `methodA()`.

```
public class Super2
{
  ...
  public void methodA()
  {
    ...
  } // methodA
  ...
} // class Super2
```

At some later date, somebody could make a subclass, `Sub`, of both `Super1` and `Super2`.

```
public class Sub extends Super1, Super2
{
  ...
  public void methodB()
  {
    ...
    methodA();
```

DRAFT April 15, 2011

```
      ...
    } // methodB
    ...
  } // class Sub
```

There are two, related, issues here. The first is about ambiguity: which `methodA()` should the call inside `methodB()` invoke? Many people regard the potential for this problem as being the basis for the view that multiple inheritance is a bad idea – it leads to problematic **inheritance hierarchy** designs. No doubt when the class `Super1` was written, the name `methodA` was a good name for the **method**. And the same was true when `Super2` was being written. But the two methods may have completely unrelated functions, written by different people at different times.

The second issue is concerned with **run time** efficiency. When the **virtual machine** is performing **dynamic method binding** for a **method call**, it needs to search the inheritance hierarchy for every superclass, to find the method, perhaps hoping there is no conflict, but somehow dealing with it if there is. This takes more time than searching up the tree in a single inheritance hierarchy.

In practice, *full* multiple inheritance is not very often required anyway. So, for all these reasons, Java does not permit a class to have more than one superclass. Every class, except `java.lang.Object`, has exactly one superclass, and `Object` has none because it is at the top of the inheritance hierarchy.

22 File IO API

22.1 File IO API: IOException (page 450)

When processing **file**s, there is much potential for things to go wrong. For example, attempting to read a file that does not exist, or the end user running out of file space while writing a file, or the **operating system** experiencing a disk or network filestore problem, and so on. As a result, most of the operations we can perform on files in Java are capable of **throw**ing an **exception**, of the **type** `java.io.IOException`. As you might expect, there are many **subclass**es of `IOException`, including `java.io.FileNotFoundException`.

`IOException` is itself a direct **subclass** of `java.lang.Exception`, rather than `java.lang.RuntimeException` and thus **instance**s of it are **checked exception**s, that is, we must write **catch clause**s or **throws clause**s for them. This is because the errors which cause them are not generally avoidable by writing code.

DRAFT April 15, 2011

22.2 File IO API: InputStream (page 451)

The basic building block for reading **data** in Java, is the **class** `java.io.InputStream`. This provides a view of the data as a **byte stream** – a continuous sequence of **byte**s.

The simplest way to access these bytes, one by one, is via the `read()` **instance method**. This takes no **method argument**s and **return**s the next byte from the stream. However, if there are no more bytes, because all of them have been read (or there was none in the first place), then it returns the number `-1` instead. If something goes wrong during the read, then an `IOException` is **thrown**.

The value returned by `read()` must be able to distinguish `-1` from the byte value `255`, which is the same as `-1` in 8-bit number representation. For this reason, the result is actually an **`int`** rather than a **`byte`**.

As an example, here is possible skeleton code to process all the data in an `InputStream`. This is another appropriate use of treating an **assignment statement** as an **expression**: we have a **loop** which terminates when the result of some expression is a certain value, and we also want to use that result inside the body of the loop. Notice that we need to put brackets around the assignment statement; this is because `=` has a lower **operator precedence** than the `!=` **operator**.

```
InputStream inputData;
try
{
  inputData = ... Code to set up inputData.
  int currentByte;
  while ((currentByte = inputData.read()) != -1)
  {
    ... Code to do something with currentByte.
  } // while
} // try
catch (IOException exception)
{
  System.err.println("Ooops -- that didn't work! " + exception.getMessage());
} // catch
finally
{
  try { if (inputData != null) inputData.close(); }
  catch (IOException exception)
    { System.err.println("Could not close input " + exception); }
} // finally
```

Notice how we have used a **try finally statement** to make sure that there is an attempt to **close** the `InputStream` even if something else goes wrong. It is a good idea to ensure we close input and/or output streams when we have finished with them. For example, on some **operating system**s that do not separate the notions of **file** name from file contents, a file cannot be deleted

DRAFT April 15, 2011

or renamed if a program has it open for reading or writing. Additionally, if we do not close output streams then the data might never get written to its destination!

22.3 File IO API: InputStreamReader (page 456)

If we wish to treat an `InputStream` as a sequence of **character**s, rather than a sequence of **byte**s, we can wrap it up in an **instance** of the **class** `java.io.InputStreamReader`. This provides an **instance method** called `read`, which **return**s the next *character* from the wrapped up `InputStream`, or `-1` if there are no more to be read. To achieve this, the instance method reads one or more bytes from the underlying `InputStream` for each character.

`InputStreamReader` has two **constructor method**s, one takes just an `InputStream` which it wraps up. It will (usually) use the the default **file encoding** in operation on the computer where the program is **run**. The second constructor method takes both an `InputStream` and the character encoding which is to be used – permitting us to read character streams that were generated under a different **locale**.

22.4 File IO API: BufferedReader (page 459)

Whilst the **class** `java.io.InputStreamReader` converts **byte**s into **character**s, it does not provide an **instance method** to read a whole line of characters in one go. Instead, this functionality is provided by `java.io.BufferedReader`. This class wraps up an `InputStreamReader` **object** and provides the instance method `readLine()`, as well as `read()` for a single character (and other **method**s). We can create a `BufferedReader` object by providing the **constructor method** with an **instance** of `InputStreamReader`, which we wish it to wrap up.

The instance method `readLine()` takes no **method argument**s and **return**s a `String`, containing the next line of the input from the underlying `InputStreamReader`; or the **null reference** if there are no more lines to be read.

22.5 File IO API: FileInputStream (page 462)

To read **byte**s from a **file**, we use an **instance** of the **class** `java.io.FileInputStream`. This is a **subclass** of `java.io.InputStream` which reads its input bytes from a file.

22.6 File IO API: FileReader (page 462)

To read **character**s instead of **byte**s from a **file**, we can wrap a `FileInputStream` in an `InputStreamReader`. For convenience we can instead create an **instance** of `java.io.FileReader`,

DRAFT April 15, 2011

which then creates the required FileInputStream and InputStreamReader internally for us. FileReader is a **subclass** of java.io.InputStreamReader, and so has a read() **instance method** to read a **character**, and can be wrapped inside a BufferedReader to obtain a readLine() instance method. One of the **constructor methods** of FileReader takes the name of the file to be accessed.

Here is a possible skeleton use of FileReader.

```
FileReader fileReader;
try
{
  fileReader = new FileReader("my-data.txt");
  int currentCharacter;
  while ((currentCharacter = fileReader.read()) != -1)
  {
    ... do something with currentCharacter.
  } //while
} //try
catch (IOException exception)
{
  System.err.println(exception.getMessage());
} // catch
finally
{
  try { if (fileReader != null) fileReader.close(); }
  catch (IOException exception)
    { System.err.println("Could not close input file " + exception); }
} // finally
```

22.7 File IO API: OutputStream (page 462)

The basic building block for writing **data** in Java, is the **class** java.io.OutputStream. Like java.io.InputStream, this provides a view of the data as a **byte stream**. OutputStream has, amongst others, an **instance method** write() to write a single **byte**.

22.8 File IO API: OutputStreamWriter (page 462)

If we wish to treat an OutputStream as a sequence of **characters**, rather than a sequence of **bytes**, we can wrap it up in an **instance** of the **class** java.io.OutputStreamWriter. This is analogous to
java.io.InputStreamReader for InputStream **objects**. OutputStreamWriter has, amongst others, an **instance method** write() to write a single character.

DRAFT April 15, 2011

22.9 File IO API: FileOutputStream (page 463)

To write **byte**s to a **file**, we use an **instance** of the **class** `java.io.FileOutputStream`. This is a **subclass** of `java.io.OutputStream` which writes its output bytes to a file.

22.10 File IO API: FileWriter (page 463)

To write **character**s instead of **byte**s to a **file**, we can wrap a `FileOutputStream` in an `OutputStreamWriter`. For convenience we can instead create an **instance** of `java.io.FileWriter`, which then creates the required `FileOutputStream` and `OutputStreamWriter` internally for us. `FileWriter` is a **subclass** of `java.io.OutputStreamWriter`, and so has a `write()` **instance method** to write a character. One of the **constructor method**s of `FileWriter` takes the name of the file to be written to.

Here is a possible skeleton use of `FileWriter`. Notice the call to the `close()` instance method in the **finally block** – it is a good idea to **close** files, especially for output files, when we have finished with them. If we do not, then it is possible that **data** written into the `FileWriter` might still be waiting in memory buffers, and never get written into the physical file.

```
FileWriter fileWriter;
try
{
  fileWriter = new FileWriter("my-results.txt");
  boolean iFeelLikeIt = ...
  while (iFeelLikeIt)
  {
    int currentCharacter = ...
    fileWriter.write(currentCharacter);
    ...
    iFeelLikeIt = ...
  } // while
} // try
catch (IOException exception)
{
  System.err.println(exception.getMessage());
} // catch
finally
{
  try { if (fileWriter != null) fileWriter.close(); }
      catch (IOException exception)
        { System.err.println("Could not close output file " + exception); }
} // finally
```

Notice that the **variable** to hold each character is an `int`. Only the lowest 16 **bit**s, which is the

DRAFT April 15, 2011

size of **char**, are used by write(). This avoids the need for us to **cast** the value to a **char** if we have in fact just obtained it from read() of an InputStream.

22.11 File IO API: PrintWriter (page 463)

Whilst the **class** java.io.OutputStreamWriter (and its **subclass** java.io.FileWriter) converts **character**s into **byte**s, it does not provide **instance method**s to print whole lines of text, or decimal representations of numbers, etc.. Instead, this functionality is provided by java.io.PrintWriter. This class wraps up an OutputStreamWriter **object** and provides instance methods println(), and print() for a range of possible **method argument**s. Since Java 5.0 it also has printf(). We can create a PrintWriter object by providing the **constructor method** with an **instance** of OutputStreamWriter, which we wish it to wrap up.

22.12 File IO API: PrintWriter: checkError() (page 464)

Curiously, the **instance method**s of the java.io.PrintWriter **class** never **throw** any **exception**s! (However, some of its **constructor method**s do.) So, to find out whether something has gone wrong with the printing, we can use its checkError() instance method. This **return**s a **boolean** which is **true** if there has been an error, **false** otherwise.

Hence, a typical use of PrintWriter might be as follows.

```
PrintWriter printWriter;
try
{
  printWriter = ...
  while (...)
  {
    ...
    printWriter.write(...);
    ...
  } // while
} // try
catch (IOException exception)
{
  System.err.println(exception.getMessage());
} // catch
finally
{
  if (printWriter != null)
  {
    // printWriter.close() does not throw an exception.
    printWriter.close();
```

DRAFT April 15, 2011

```
      if (printWriter.checkError())
        System.err.println("Something went wrong with the output");
    } // if
  } // finally
```

22.13 File IO API: PrintWriter: versus PrintStream (page 468)

An often asked question is, what is the difference between `java.io.PrintStream` and `java.io.PrintWriter`? `PrintStream` is a **subclass** of `OutputStream`, and so has `write()` **instance method**s for writing **byte**s, but also has `print()`, `println()` and `printf()` instance methods for printing representations of things as **character**s, (e.g. decimal representations of **int**s, `String`s as lines, etc.). A `PrintWriter` is a wrapper around an **instance** of `java.io.OutputStreamWriter` and provides `print()`, `println()` and `printf()` instance methods for printing representations as characters via that `OutputStreamWriter`. It does not have any way to write bytes.

The desire to write a *mixture* of bytes and characters to the same stream is highly unusual – we nearly always want either all bytes or all characters, the latter sometimes with the ability to print representations. `PrintStream` primarily exists for `System.out` and `System.err`, so that the **standard output** and the **standard error** are each available as a stream of bytes, but can also be conveniently treated as 'printable' – e.g. for error messages, debugging messages, or very simple programs.

Programs that need to produce representations as a stream of characters should use `PrintWriter` rather than `PrintStream`, *because* `PrintWriter` does not have instance methods to write bytes; we cannot accidentally use them. (And programs that wish to produce a stream of bytes should use `OutputStream` (including `java.io.FileOutputStream`) rather than `PrintStream`.)

22.14 File IO API: PrintWriter: can also wrap an OutputStream (page 468)

`System.out` is an `OutputStream` (actually its **subclass**, `PrintStream`). If we wish to treat it as a `PrintWriter`, then we can wrap it up inside an `OutputStreamWriter` and then inside a `PrintWriter`.

```
  PrintWriter systemOut = new PrintWriter(new OutputStreamWriter(System.out));
```

However, for convenience one of the **constructor method**s of `PrintWriter` can take an `OutputStream` directly, and **construct** the intermediate `OutputStreamWriter` internally for us.

```
  PrintWriter systemOut = new PrintWriter(System.out);
```

DRAFT April 15, 2017

All **instance**s of output **class**es which act as wrappers around some other output class **object** may typically store their output in an internal buffer before sending it to the wrapped up object, in an effort to speed up overall operation of our programs. Such buffers are **flush**ed by calls to the `flush()` **instance method**, or when the output is **close**d, via the `close()` **instance method**. For a `PrintWriter` which is wrapping up `System.out`, it is likely we would want to enable **automatic flushing**. This ensures that **data** is sent all the way through to appearing at the final destination (e.g. the screen) whenever one of the `println()` or `printf()` instance methods has finished producing its result (but not `print()`). Automatic flushing can be enabled by using a separate constructor method which takes an additional **`boolean`** **method argument**.

```
PrintWriter systemOut = new PrintWriter(System.out, true);
```

22.15 File IO API: File (page 469)

The **class** `java.io.File` allows us to examine properties of **file**s. Although the class is called `File`, it is really all about file *names*, and properties of any files of those names. One **constructor method** of the `File` class takes the path name of a file as its single **method argument**. There are a number of **instance method**s, including `exists()` which **return**s a **`boolean`** indicating whether or not the `File` **object** represents a file that actually exists. In other words, whether or not the path name given to the constructor method is the name of a file that currently exists.

22.16 File IO API: DataOutputStream (page 479)

If we wish to write values of any **primitive type**, rather than just **`byte`**, to a **binary file**, we can use the `java.io.DataOutputStream` **class**. This is a **subclass** of `java.io.OutputStream` and an **instance** of it is also a wrapper around an `OutputStream` (including its subclasses such as `java.io.FileOutputStream`). For example, a `DataOutputStream` **object** which writes to the **file** `out.dat` can be **construct**ed with the following code.

```
DataOutputStream out = new DataOutputStream(new FileOutputStream("out.dat"));
```

`DataOutputStream` has **instance method**s to write all the kinds of primitive type, such as `writeInt()` to write an **`int`** value in four **byte**s, and `writeShort()` to write a **`short`** value in two bytes. The *most* significant byte of numbers is written first, although if we intend to read the **data** back using the corresponding `readXXX()` instance method of `java.io.DataInputStream`, we do not really need to worry about the byte order.

Instances of `java.lang.String` can also be written, using the `writeUTF()` instance method. This records the information in (a slight variant of) a **file encoding** known as **8-bit Unicode Transformation Format**. **UTF-8** allows for all **Unicode**[20] **character**s to be represented.

DRAFT April 15, 2011

22.17 File IO API: DataInputStream (page 479)

If we wish to read values from a **binary file** which was written using a `DataOutputStream`, we can use the `java.io.DataInputStream` **class**. This is a **subclass** of `java.io.InputStream` and an **instance** of it is also a wrapper around an `InputStream` (including its subclasses such as `java.io.FileInputStream`). For example, a `DataInputStream` **object** which reads from the **file** `in.dat` can be **construct**ed with the following code.

```
DataInputStream in = new DataInputStream(new FileInputStream("in.dat"));
```

`DataInputStream` has **instance method**s to read all the kinds of **primitive type**, such as `readInt()` to read an `int` value from four **byte**s, and `readShort()` to read a `short` value from two bytes. The *most* significant byte of numbers is read first, although if we are just reading **data** back which was written using the corresponding `writeXXX()` instance method of `DataOutputStream`, we do not really need to worry about the byte order.

Instances of `java.lang.String` which were written using `writeUTF()` of `DataOutputStream`, can be read using the `readUTF()` instance method.

23 Collections API

23.1 Collections API (page 538)

The need to store **collection**s of **data** is very common in programming, and so in addition to the **array type** built-in to Java, the standard Java **application program interface** (**API**) provides the **collections framework**. This is a group of **class**es and **interface**s designed to store collections of data in various different ways. These collections typically allow elements to be added to them without us worrying about memory allocation – that is, they automatically grow big enough to hold the elements that are added to them.

23.2 Collections API: Lists (page 538)

One of the kinds of **collection** supported by the **collections framework** is the **list collection**. These are collections of **data** which are essentially **list**s or sequences. This means that duplicate elements are permitted, they are stored in some order, and each element occurs at a particular **list index** position, starting at index zero. Lists are, in principle, similar to **array**s.

DRAFT April 15, 2011

23.3 Collections API: Lists: List interface (page 538)

The **interface** `java.util.List` is part of the **collections framework**. It specifies the **instance methods** needed to support a **list collection**. These include the following.

Method definitions in interface `List` (some of them).			
Method	**Return**	**Arguments**	**Description**
`size`	**`int`**		Returns the size of this `List`, that is, the number of elements in it.
`add`	**`boolean`**	`Object`	Appends the given `Object` to the end of the `List`. Returns **`true`**.
`get`	`Object`	**`int`**	Returns the `Object` at the specified **list index**, which must be legal (`0 <= index < size()`) to avoid an `IndexOutOfBoundsException`.
`set`	`Object`	**`int`**, `Object`	Overwrites an existing element with a new one: i.e. it replaces the `Object` at the given **`int`** list index with the given other `Object`. Returns the original `Object`. The index must be legal to avoid an `IndexOutOfBoundsException`.

Since Java 5.0, `List` is a **generic interface** with a single **type parameter** representing the **type** of **objects** that can be stored in it. So, when we use a **parameterized type** of `List` rather than its **raw type**, all the occurrences of `Object` in the above table of instance methods are replaced by the **type argument**.

23.4 Collections API: Lists: List interface: iterator() (page 553)

The **instance method** `iterator()`, specified in the **interface** `java.util.List`, **return**s an **object** that **implement**s `java.util.Iterator`, supporting an **iteration** of the elements in the `List`, in ascending order of their **list index** in the `List`.

For example, the following code prints out all the elements of a `List`, from the one indexed by zero, up to the last one, indexed by `size()` minus one.

```
public static <ListType> void printList(List<ListType> list)
{
  Iterator<ListType> iterator = list.iterator();
  while (iterator.hasNext())
  {
    ListType item = iterator.next();
```

```
    System.out.println(item);
  } // while
} // printList
```

For an `ArrayList`, this way of scanning through the elements is just as efficient as using the list index of each element. However, for some kinds of `Lists`, accessing by index is not efficient, whereas scanning using an `Iterator` always will be, because it is designed for that purpose. So, as a rule of thumb, whenever you need to scan through the elements of a **list** in an arbitrary order, or from first to last, you should use an `Iterator` rather than the indices.

23.5 Collections API: Lists: List interface: extends Collection (page 556)

The **interface** `java.util.List` is an **extension** of the more general interface `java.util.Collection`.

```
public interface List<E> extends Collection<E>
{
  ...
} // interface List
```

23.6 Collections API: Lists: ArrayList (page 539)

The **class** `java.util.ArrayList` is part of the **collections framework**, and is one **implementation** of a **list collection**. It **implements** the `java.util.List` **interface**. As the name suggests, this kind of **list** is implemented using a **private instance variable**, which is an **array** of **type** `java.lang.Object[]`. This array is grown (by **array extension**) automatically as required.

Since Java 5.0, `ArrayList`, and the other classes in the collections framework are **generic classes**. The **type parameter** of an `ArrayList` is the **type** of **objects** that can be stored in it.

```
public class ArrayList<E> implements List<E>
{ ... }
```

23.7 Collections API: Lists: add(index) and remove(index) (page 557)

The **interface** `java.util.List` specifies **instance methods** for adding and removing an element at a particular **list index**, in addition to those defined in `java.util.Collection` for adding an element (at the end in **lists**), or removing an element **equivalent** to a given one.

DRAFT April 15, 2011

Method definitions in interface `List` (some more of them).			
Method	**Return**	**Arguments**	**Description**
add		**int**, Object	Inserts the given Object at the specified list index, shifting any elements after that position up by one place. To avoid an IndexOutOfBoundsException, the index must be legal (0 <= index <= size()).
remove	Object	**int**	Removes the element at the given list index, shifting elements after that position down by one place. To avoid an IndexOutOfBoundsException, the index must be legal (0 <= index < size()).

DRAFT April 15, 2011

23.8 Collections API: Lists: LinkedList (page 558)

The **class** java.util.LinkedList is part of the **collections framework**, and is another implementation of a **list collection**. It **implement**s the java.util.List **interface** by using a **doubly linked list**.

Since Java 5.0, LinkedList, and the other classes in the collections framework are **generic class**es. The **type parameter** of a LinkedList is the **type** of **object**s that can be stored in it.

```
public class LinkedList<E> implements List<E>
{ ... }
```

23.9 Collections API: Collections class (page 543)

The standard **class** java.util.Collections provides various **class method**s to perform complex manipulations of **collection**s. One of these is called sort, and takes a List of Objects which it **sort**s into their **natural ordering**. For this to work without **throw**ing an **exception**, the items in the List must all be of **type** java.lang.Comparable and be **mutually comparable**. The **algorithm** used is called **merge sort**, which is far more efficient than **bubble sort** (but less simple).

Since Java 5.0, many of the class methods in Collections have become **generic method**s. The sort() class method has a single **type parameter** which is the **type** of the items in the given List. These must be Comparable with themselves, and so you would probably expect the heading of the class method to be as follows.

```
public static <T extends Comparable<T>>
              void sort(List<T> list)
```

In fact, it is defined in this way instead.

```
public static <T extends Comparable<? super T>>
              void sort(List<T> list)
```

The code <? super T> means “any type that is T or a **superclass** (or **superinterface**) of it”. So here, this means any supplied **type argument** must **implement** Comparable with itself, or a superclass of itself. Many of the **type parameter**s in the standard **application program interface** (**API**) are expressed in that sort of way, because it leads to more flexibility and convenience.

23.10 Collections API: Sets (page 546)

Another of the kinds of **collection** supported by the **collections framework** is the **set collection**. These are collections of **data** which are essentially **set**s, which means that adding an element to them has no effect if the set already contains an element that is **equivalent** to the new one. Also, the order in which the elements are added to the collection is *not* preserved.

For the purposes of determining whether two Objects are equivalent, sets are intended to use the equals() **instance method** of the elements in them.

23.11 Collections API: Sets: Set interface (page 546)

The **interface** java.util.Set is part of the **collections framework**. It specifies the **instance method**s needed to support a **set collection**. These include the following.

Method definitions in interface Set (some of them).

Method	Return	Arguments	Description
size	**int**		Returns the size of this Set, that is, the number of elements in it.
add	**boolean**	Object	Inserts the given Object into the Set, unless an **equivalent** one is already present. Returns **true** if it gets added, **false** otherwise.
contains	**boolean**	Object	Return **true** if the Set contains an Object which is equivalent to the given one, **false** otherwise.

Since Java 5.0, Set is a **generic interface**. The **type parameter** of a Set is the **type** of **object**s

that can be stored in it. So, when we use a **parameterized type** of `Set` rather than its **raw type**, all the occurrences of `Object` in the above table of instance methods are replaced by the **type argument**.

23.12 Collections API: Sets: Set interface: iterator() (page 554)

The **instance method** `iterator()`, specified in the **interface** `java.util.Set`, **return**s an **object** that **implement**s `java.util.Iterator`, supporting an **iteration** of the elements in the `Set`. The order of the iteration will depend on the kind of **set**, and may be in some arbitrary order.

23.13 Collections API: Sets: Set interface: extends Collection (page 557)

The **interface** `java.util.Set` is an **extension** of the more general interface `java.util.Collection`.

```
public interface Set<E> extends Collection<E>
{
  ...
} // interface Set
```

23.14 Collections API: Sets: HashSet (page 548)

The **class** `java.util.HashSet` is part of the **collections framework**, and is one implementation of a **set collection**. It **implement**s the `java.util.Set` **interface**. This kind of **set** uses a **hash table**, with the **hash code**s being obtained from the `hashCode()` **instance method** of the items stored in it. For this to work, any **object**s which are **equivalent** *must* have the same hash code, otherwise multiple copies of equivalent items will be allowed in the set! For efficiency, non-equivalent objects should tend to have different hash codes.

Since Java 5.0, `HashSet`, and the other classes in the collections framework are **generic class**es. The **type parameter** of a `HashSet` is the **type** of **object**s that can be stored in it.

```
public class HashSet<E> implements Set<E>
{ ... }
```

23.15 Collections API: Sets: TreeSet (page 552)

The **class** `java.util.TreeSet` is part of the **collections framework**, and is another implementation of a **set collection**. It **implement**s the `java.util.Set` **interface**. This kind of **set**

DRAFT April 15, 2017

uses an **ordered binary tree** and so it has to be possible to order the elements which are stored in it. The simplest way of providing such an ordering is to ensure that the **class** of the elements implements `java.lang.Comparable`.

Since Java 5.0, `TreeSet`, and the other classes in the collections framework are **generic class**es. The **type parameter** of a `TreeSet` is the **type** of **object**s that can be stored in it.

```
public class TreeSet<E> implements Set<E>
{ ... }
```

23.16 Collections API: Sets: TreeSet: iterator() (page 554)

The `iterator()` **instance method** of the `java.util.TreeSet` **class return**s an **object** that **implement**s
`java.util.Iterator`, which supports an **iteration** through the elements in the order they appear in the tree, from left to right. With the simplest use of a `TreeSet` we thus get the **natural ordering** of elements as provided by **method implementation**s of `compareTo()`.

As a rule of thumb, `java.util.HashSet` should be used in preference to `TreeSet` when it is not desired to obtain the values from the **set collection** in a specific order. If there is little or no **hash code** clashing, a `HashSet` operates in nearly constant time per addition and membership test. By contrast, a `TreeSet` operates in time which is proportional to the logarithm of the size of the **set**.

23.17 Collections API: Iterator interface (page 553)

The **interface** `java.util.Iterator` is part of the **collections framework**. It specifies the **instance method**s needed to support a way of accessing all the elements in a **collection** one by one.

Method definitions in interface `Iterator` (some of them).

Method	Return	Arguments	Description
`hasNext`	**`boolean`**		Returns **`true`** if the **iteration** has more elements, **`false`** otherwise.
`next`	`Object`		Returns the next element in the iteration, and moves the iteration on to the element following that one.

When a **new** `Iterator` **object** is obtained from a collection, `hasNext()` will **return** **`true`**,

DRAFT April 15, 2011

unless the collection is empty. The first time we call `next()`, we get the first element from the iteration if there is one, then the second time we get the second element, and so on. Sooner or later `hasNext()` will return **`false`** because we have called `next()` as many times as there are elements. Typically we use `hasNext()` to control a **loop** and call `next()` inside the loop only if there is another element.

All **list collection**s and **set collection**s support the instance method `iterator()` which returns some **object** that is an **instance** of some **class** that **implement**s `Iterator`. The object returned supports an iteration through the elements of the collection, in an order which depends on the kind of collection.

Since Java 5.0, `Iterator` is a **generic interface**. The **type parameter** of an `Iterator` is the **type** of objects that are stored in the corresponding collection. In other words, if the collection was given a **type argument**, then the `next()` instance method of an `Iterator` over that collection returns an object of that type.

23.18 Collections API: Collection interface (page 556)

The **interface** `java.util.Collection` is part of the **collections framework**. It specifies the **instance method**s needed to support a **collection**, such as a **list collection** or a **set collection**. These include the following.

Method definitions in interface `Collection` (some of them).

Method	Return	Arguments	Description
`size`	**`int`**		Returns the size of this `Collection`, that is, the number of elements in it.
`add`	**`boolean`**	`Object`	Ensures that this `Collection` contains the given `Object`, or an **equivalent** one if appropriate. It **return**s **`true`** if the `Collection` was modified, **`false`** otherwise. For example, a `List` always appends the element on the end and returns **`true`**, whereas a `Set` will do nothing if it already contains an equivalent element.
`remove`	**`boolean`**	`Object`	Removes one element equivalent to the given `Object`, and returns **`true`** if the `Collection` was changed (i.e. there was at least one element matching the given one).
`addAll`	**`boolean`**	`Collection`	Adds all the elements of the given `Collection` to this one, and returns **`true`** if this collection was changed. (E.g. the given collection could be empty, or this one could be a `Set` and already contain the elements.)

DRAFT April 15, 2011

Method definitions in interface `Collection` (some of them).

Method	Return	Arguments	Description
`removeAll`	**`boolean`**	`Collection`	Removes all the elements of the given `Collection` from this one, and returns **`true`** if this collection was changed.
`retainAll`	**`boolean`**	`Collection`	Removes all elements of this collection which are *not* contained in the given `Collection`, and returns **`true`** if this collection was changed.
`contains`	**`boolean`**	`Object`	Returns **`true`** if the `Collection` contains at least one `Object` which is equivalent to the given one, **`false`** otherwise.
`containsAll`	**`boolean`**	`Collection`	Returns **`true`** if this `Collection` contains at least one equivalent `Object` for each element in the given collection, **`false`** otherwise.
`iterator`	`Iterator`		Returns an **object** that **implement**s `java.util.Iterator`, giving access to all the elements of the `Collection`. The order depends on the kind of collection.

Since Java 5.0, `Collection` is a **generic interface** with a single **type parameter** which represents the **type** of **object**s that can be stored in it. So, when we use a **parameterized type** of `Collection` rather than its **raw type**, all occurrences of `Object` in the above table of instance methods are replaced by the **type argument**.

23.19 Collections API: Collection interface: constructor taking a Collection (page 568)

The **application program interface** (**API**) documentation of the `java.util.Collection` **interface** states that any **class** which **implement**s it should provide two **constructor method**s, one which takes no **method argument**s and builds an empty `Collection`, and one which takes an existing `Collection` and builds a **new** one containing the same elements.

Interestingly, there is no way for this requirement to be enforced in Java, as interfaces cannot specify constructor methods! It could be argued that this is a deficiency in the use of interfaces as a means of contractual obligation.

Anyway, all the standard implementations of `java.util.Collection` do satisfy the requirement.

DRAFT April 15, 2014

23.20 Collections API: Maps (page 559)

Another kind of **collection** supported by the **collections framework** is the **map**. One view of **array**s and **list collection**s is that they are functions from a **key** to a corresponding element, where the key is the **array index** or **list index**. A map is more general, in the sense that the key can be any **type** of **object**, rather than always an **`int` index**. For every key in the map, there is an associated value. Two different keys may map on to the same value, but every possible key maps on to at most one value. To put it another way, you can think of a map as being a **set** of pairs, each containing a key and a value. The keys are all unique within a particular map – every pair has a key which is not **equivalent** to the key in any other pair. By contrast, the values may be duplicated – any number of different pairs may have values which are equivalent. Thus, a map is a **many-to-one association**, otherwise known in Mathematics as a **function**.

23.21 Collections API: Maps: Map interface (page 560)

The **interface** `java.util.Map` is part of the **collections framework**. It specifies the **instance method**s needed to support a **map**. These include the following.

Method definitions in interface `Map` (some of them).

Method	Return	Arguments	Description
`put`	`Object`	`Object, Object`	Takes a **key** and a value, and adds that association to the map. If the map previously contained a mapping for this key (or an **equivalent** one), the old value is replaced with the new one. Returns the **null reference**, if this is a new key, or **return**s the old value otherwise.
`get`	`Object`	`Object`	Takes a key and returns the value associated with it, or the null reference if the map does not contain a mapping with a key which equivalent to the given one.
`values`	`Collection`		Returns a `Collection` of the values (not keys) in the map. The `iterator()` **instance method** of the resulting `Collection` may support iterating through the values in a particular order, or not, depending on the kind of `Map`.
`keySet`	`Set`		Returns a `Set` of the keys (not values) in the map.

Since Java 5.0, `Map` is a **generic interface**. There are *two* **type parameter**s for a `Map`, first the

DRAFT April 15, 2011

type of **object**s that can be used as keys, and then the type of objects that can be used as values. So, when we use a **parameterized type** of `Map` rather than its **raw type**, occurrences of `Object` in the above table of instance methods are replaced by the corresponding **type argument** as appropriate.

23.22 Collections API: Maps: TreeMap (page 560)

The **class** `java.util.TreeMap` is part of the **collections framework**, and is one implementation of a **map**. It **implement**s the `java.util.Map` **interface**. This kind of map is implemented using an **ordered binary tree**. This means that there must be an ordering on the **key**s of the map, and the simplest way of providing such an ordering is to ensure that the keys implement `java.util.Comparable`.

The `values()` **instance method** of `TreeMap` gives a `Collection`; the `iterator()` of this gives an **object** that **implement**s `java.util.Iterator`, and supports **iteration** over the values of the map in key order.

Since Java 5.0, `TreeMap`, and the other classes in the collections framework are **generic class**es. There are *two* **type parameter**s for a `TreeMap`, first the **type** of **object**s that can be used as keys, and then the type of objects that can be used as values.

```
public class TreeMap<K, V> implements Map<K, V>
{ ... }
```

(Actually `TreeMap` implements an interface called `java.util.SortedMap` which **extend**s `Map`.)

23.23 Collections API: Maps: HashMap (page 567)

The **class** `java.util.HashMap` is part of the **collections framework**, and is another implementation of a **map**. It **implement**s the `java.util.Map` **interface**. This kind of map is implemented using a **hash table** and so each element must have an appropriate implementation of `hashCode()` so that the `HashMap` works correctly.

The `values()` **instance method** of `HashMap` gives a `Collection` containing the values of the map, which can yield an **object implement**ing `java.util.Iterator` that supports **iteration** over these values in no specific order.

As a rule of thumb, `HashMap` should be used in preference to `java.util.TreeMap` when it is not desired to obtain the values from the map in **key** order. If there is little or no **hash code** clashing, a `HashMap` operates in nearly constant time per look up and addition. By contrast, a `TreeMap` operates in time which is proportional to the logarithm of the size of the map.

DRAFT April 15, 2011

Since Java 5.0, `HashMap`, and the other classes in the collections framework are **generic class**es. There are *two* **type parameter**s for a `HashMap`, first the **type** of **object**s that can be used as keys, and then the type of objects that can be used as values.

```
public class HashMap<K, V> implements Map<K, V>
{ ... }
```

DRAFT April 15, 2011