

Introduction



CHAPTER OBJECTIVES

- Describe the relationship between hardware and software.
- Define various types of software and how they are used.
- Identify the core hardware components of a computer and explain their roles.
- Explain how the hardware components interact to execute programs and manage data.
- Describe how computers are connected into networks to share information.
- Introduce the Java programming language.
- Describe the steps involved in program compilation and execution.
- Present an overview of object-oriented principles.

This book is about writing well-designed software. To understand software, we must first have a fundamental understanding of its role in a computer system. Hardware and software cooperate in a computer system to accomplish complex tasks. The purpose of various hardware components, and the way those components are connected into networks, are important prerequisites to the study of software development. This chapter first discusses basic computer processing, and then begins our exploration of software development by introducing the Java programming language and the principles of object-oriented programming.

1.1 Computer Processing

We begin our exploration of computer systems with an overview of computer processing, defining some fundamental terminology and showing how the key pieces of a computer system interact.

A computer system is made up of hardware and software. The *hardware* components of a computer system are the physical, tangible pieces that support the computing effort. They include chips, boxes, wires, keyboards, speakers, disks, cables, plugs, printers, mice, monitors, and so on. If you can physically touch it and it can be considered part of a computer system, then it is computer hardware.

KEY CONCEPT

A computer system consists of hardware and software that work in concert to help us solve problems.

The hardware components of a computer are essentially useless without instructions to tell them what to do. A *program* is a series of instructions that the hardware executes one after another. *Software* consists of programs and the data those programs use. Software is the intangible counterpart to the physical hardware components. Together they form a tool that we can use to solve problems.

The key hardware components in a computer system are:

- central processing unit (CPU)
- input/output (I/O) devices
- main memory
- secondary memory devices

Each of these hardware components is described in detail in the next section. For now, let's simply examine their basic roles. The *central processing unit* (CPU) is the device that executes the individual commands of a program. *Input/output* (I/O) *devices*, such as the keyboard, mouse, and monitor, allow a human being to interact with the computer.

Programs and data are held in storage devices called memory, which fall into two categories: main memory and secondary memory. *Main memory* is the storage device that holds the software while it is being processed by the CPU. *Secondary memory* devices store software in a relatively permanent manner. The most important secondary memory device of a typical computer system is the hard disk that resides inside the main computer box. A floppy disk is similar to a hard disk, but it cannot store nearly as much information as a hard disk. Floppy disks have the advantage of portability; they can be removed temporarily or moved from computer to computer as needed. Other portable secondary memory devices include zip disks and compact discs (CDs).

Figure 1.1 shows how information moves among the basic hardware components of a computer. Suppose you have an executable program you wish to run. The program is stored on some secondary memory device, such as a hard disk.

When you instruct the computer to execute your program, a copy of the program is brought in from secondary memory and stored in main memory. The CPU reads the individual program instructions from main memory. The CPU then executes the instructions one at a time until the program ends. The data that the instructions use, such as two numbers that will be added together, are also stored in main memory. They are either brought in from secondary memory or read from an input device such as the keyboard. During execution, the program may display information to an output device such as a monitor.

The process of executing a program is fundamental to the operation of a computer. All computer systems basically work in the same way.

KEY CONCEPT

The CPU reads the program instructions from main memory, executing them one at a time until the program ends.

Software Categories

Software can be classified into many categories using various criteria. At this point we will simply differentiate between system programs and application programs.

The *operating system* is the core software of a computer. It performs two important functions. First, it provides a *user interface* that allows the user to interact with the machine. Second, the operating system manages computer resources such as the CPU and main memory. It determines when programs are allowed to run, where they are loaded into memory, and how hardware devices communicate. It is the operating system's job to make the computer easy to use and to ensure that it runs efficiently.

Several popular operating systems are in use today. Windows 2000 and Windows XP are two versions of the operating system developed by Microsoft for personal computers. Various versions of the Unix operating system are also quite popular, especially in larger

KEY CONCEPT

The operating system provides a user interface and manages computer resources.




FIGURE 1.1 A simplified view of a computer system

computer systems. A version of Unix called Linux was developed as an open source project, which means that many people contributed to its development and its code is freely available. Because of that, Linux has become a particular favorite among some users. Mac OS is the operating system used for computing systems developed by Apple Computers.

An *application* is a generic term for just about any software other than the operating system. Word processors, missile control systems, database managers, Web browsers, and games can all be considered application programs. Each application program has its own user interface that allows the user to interact with that particular program.

The user interface for most modern operating systems and applications is a *graphical user interface* (GUI), which, as the name implies, make use of graphical screen elements. These elements include:

- *windows*, which are used to separate the screen into distinct work areas
- *icons*, which are small images that represent computer resources, such as a file
- *pull-down menus*, which provide the user with lists of options
- *scroll bars*, which allow the user to move up and down in a particular window
- *buttons*, which can be “pushed” with a mouse click to indicate a user selection

The mouse is the primary input device used with GUIs; thus, GUIs are sometimes called *point-and-click interfaces*. The screen shot in Figure 1.2 shows an example of a GUI.

KEY CONCEPT

As far as the user is concerned, the interface is the program.

The interface to an application or operating system is an important part of the software because it is the only part of the program with which the user directly interacts. To the user, the interface *is* the program. Throughout this book we discuss the design and implementation of graphical user interfaces.

The focus of this book is the development of high-quality application programs. We explore how to design and write software that will perform calculations, make decisions, and control graphics. We use the Java programming language throughout the text to demonstrate various computing concepts.

Digital Computers

Two fundamental techniques are used to store and manage information: analog and digital. *Analog* information is continuous, in direct proportion to the source of the information. For example, a mercury thermometer is an analog device for

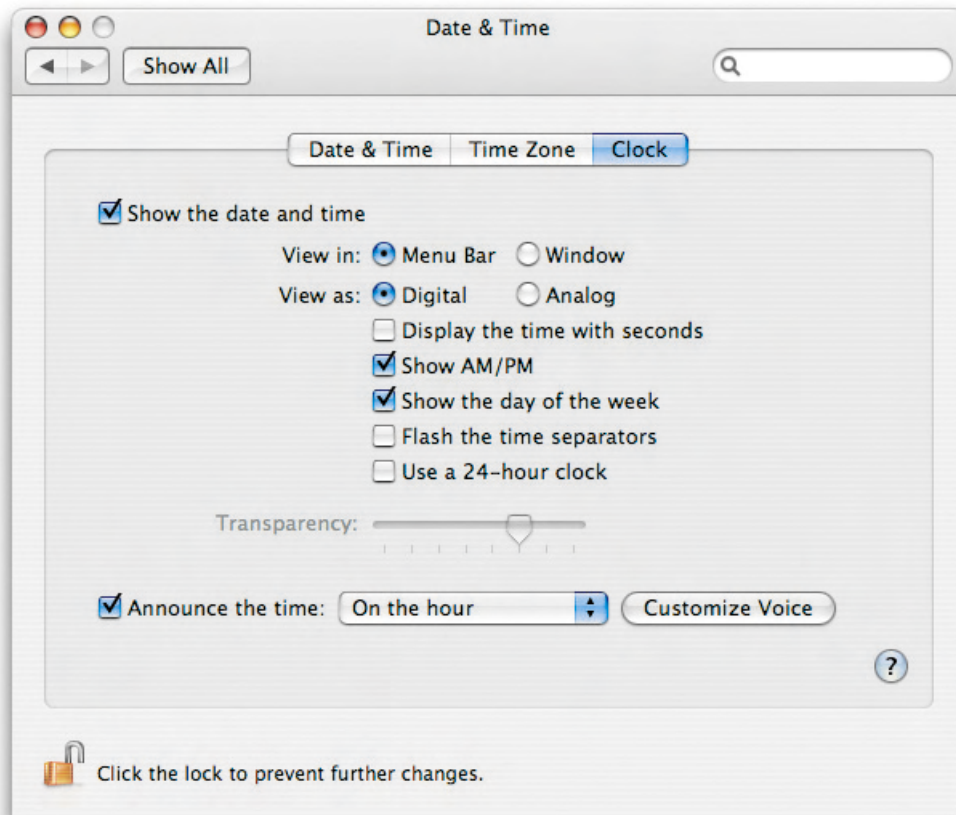


FIGURE 1.2 An example of a graphical user interface (GUI)

measuring temperature. The mercury rises in a tube in direct proportion to the temperature outside the tube. Another example of analog information is an electronic signal used to represent the vibrations of a sound wave. The signal's voltage varies in direct proportion to the original sound wave. A stereo amplifier sends this kind of electronic signal to its speakers, which vibrate to reproduce the sound. We use the term analog because the signal is directly analogous to the information it represents. Figure 1.3 graphically depicts a sound wave captured by a microphone and represented as an electronic signal.

Digital technology breaks information into discrete pieces and represents those pieces as numbers. The music on a compact disc is stored digitally, as a series of numbers. Each number represents the voltage level of one specific instance of the recording. Many of these measurements are taken in a short period of time, perhaps 40,000 measurements every second. The number of measurements




FIGURE 1.3 A sound wave and an electronic analog signal that represents the wave

KEY CONCEPT

Digital computers store information by breaking it into pieces and representing each piece as a number.

per second is called the *sampling rate*. If samples are taken often enough, the discrete voltage measurements can be used to generate a continuous analog signal that is “close enough” to the original. In most cases, the goal is to create a reproduction of the original signal that is good enough to satisfy the human senses.

Figure 1.4 shows the sampling of an analog signal. When analog information is converted to a digital format by breaking it into pieces, we say it has been

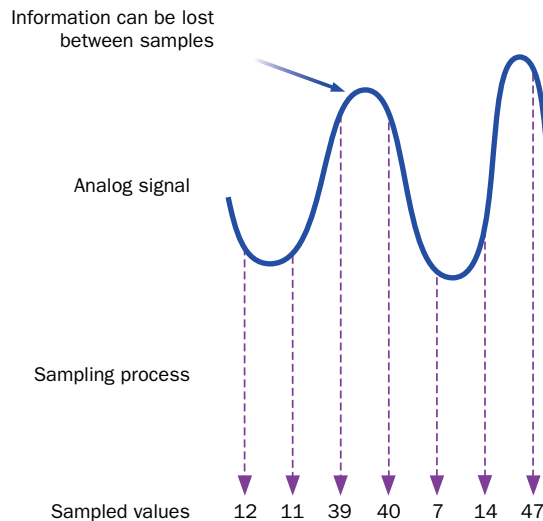


FIGURE 1.4 Digitizing an analog signal by sampling

the binary number system is base 2. Appendix B contains a detailed discussion of number systems.

KEY CONCEPT

Binary is used to store information in a computer because the devices that store and manipulate binary data are inexpensive and reliable.

Modern computers use binary numbers because the devices that store and move information are less expensive and more reliable if they have to represent only one of two possible values. Other than this characteristic, there is nothing special about the binary number system. Computers have been created that use other number systems to store information, but they aren't as convenient.

Some computer memory devices, such as hard drives, are magnetic in nature. Magnetic material can be polarized easily to one extreme or the other, but intermediate levels are difficult to distinguish. Therefore, magnetic devices can be used to represent binary values quite efficiently—a magnetized area represents a binary 1 and a demagnetized area represents a binary 0. Other computer memory devices are made up of tiny electrical circuits. These devices are easier to create and are less likely to fail if they have to switch between only two states. We're better off reproducing millions of these simple devices than creating fewer, more complicated ones.

Binary values and digital electronic signals go hand in hand. They improve our ability to transmit information reliably along a wire. As we've seen, an analog signal has continuously varying voltage, but a digital signal is *discrete*, which means the voltage changes dramatically between one extreme (such as +5 volts) and the other (such as -5 volts). At any point, the voltage of a digital signal is considered to be either "high," which represents a binary 1, or "low," which represents a binary 0. Figure 1.6 compares these two types of signals.

As a signal moves down a wire, it gets weaker and degrades due to environmental conditions. That is, the voltage levels of the original signal change slightly. The trouble with an analog signal is that as it fluctuates, it loses its original information. Since the information is directly analogous to the signal, any change in




FIGURE 1.6 An analog signal vs. a digital signal

the signal changes the information. The changes in an analog signal cannot be recovered because the degraded signal is just as valid as the original. A digital signal degrades just as an analog signal does, but because the digital signal is originally at one of two extremes, it can be reinforced before any information is lost. The voltage may change slightly from its original value, but it still can be interpreted as either high or low.

The number of bits we use in any given situation determines the number of unique items we can represent. A single bit has two possible values, 0 and 1, and therefore can represent two possible items or situations. If we want to represent the state of a light bulb (off or on), one bit will suffice, because we can interpret 0 as the light bulb being off and 1 as the light bulb being on. If we want to represent more than two things, we need more than one bit.

Two bits, taken together, can represent four possible items because there are exactly four combinations of two bits: 00, 01, 10, and 11. Suppose we want to represent the gear that a car is in (park, drive, reverse, or neutral). We would need only two bits, and could set up a mapping between the bit combinations and the gears. For instance, we could say that 00 represents park, 01 represents drive, 10 represents reverse, and 11 represents neutral. In this case, it wouldn't matter if we switched that mapping around, though in some cases the relationships between the bit combinations and what they represent is important.

Three bits can represent eight unique items, because there are eight combinations of three bits. Similarly, four bits can represent 16 items, five bits can represent 32 items, and so on. Figure 1.7 shows the relationship between the number of bits used and the number of items they can represent. In general, N bits can represent 2^N unique items. For every bit added, the number of items that can be represented doubles.

KEY CONCEPT

There are exactly 2^N combinations of N bits. Therefore, N bits can represent up to 2^N unique items.

We've seen how a sentence of text is stored on a computer by mapping characters to numeric values. Those numeric values are stored as binary numbers. Suppose we want to represent character strings in a language that contains 256 characters and symbols. We would need to use eight bits to store each character because there are 256 unique permutations of eight bits (2^8 equals 256). Each bit permutation, or binary value, is mapped to a specific character.

Ultimately, representing information on a computer boils down to the number of items there are to represent and determining the way those items are mapped to binary values.

1 bit 2 items	2 bits 4 items	3 bits 8 items	4 bits 16 items	5 bits 32 items
0	00	000	0000	00000 10000
1	01	001	0001	00001 10001
	10	010	0010	00010 10010
	11	011	0011	00011 10011
		100	0100	00100 10100
		101	0101	00101 10101
		110	0110	00110 10110
		111	0111	00111 10111
			1000	01000 11000
			1001	01001 11001
			1010	01010 11010
			1011	01011 11011
			1100	01100 11100
			1101	01101 11101
			1110	01110 11110
			1111	01111 11111

FIGURE 1.7 The number of bits used determines the number of items that can be represented

SELF-REVIEW QUESTIONS *(see answers in Appendix N)*

- SR 1.1** What is hardware? What is software?
- SR 1.2** What are the two primary functions of an operating system?
- SR 1.3** The music on a CD is created using a sampling rate of 40,000 measurements per second. Each measurement is stored as a number that represents a specific voltage level. How many such numbers are used to store a three-minute long song? How many such numbers does it take to represent one hour of music?
- SR 1.4** What happens to information when it is stored digitally?
- SR 1.5** How many unique items can be represented with the following?
- 2 bits
 - 4 bits
 - 5 bits
 - 7 bits
- SR 1.6** Suppose you want to represent each of the 50 states of the United States using a unique combination of bits. How many bits would be needed to store each state representation? Why?

1.2 Hardware Components

Let's examine the hardware components of a computer system in more detail. Consider the computer described in Figure 1.8. What does it all mean? Is the system capable of running the software you want it to? How does it compare to other systems? These terms are explained throughout this section.

Computer Architecture

The architecture of a house defines its structure. Similarly, we use the term *computer architecture* to describe how the hardware components of a computer are put together. Figure 1.9 illustrates the basic architecture of a generic computer system. Information travels between components across a group of wires called a *bus*.

The CPU and the main memory make up the core of a computer. As we mentioned earlier, main memory stores programs and data that are in active use, and the CPU methodically executes program instructions one at a time.

Suppose we have a program that computes the average of a list of numbers. The program and the numbers must reside in main memory while the program runs. The CPU reads one program instruction from main memory and executes it. If an instruction needs data, such as a number in the list, to perform its task, the CPU reads that information as well. This process repeats until the program ends. The average, when computed, is stored in main memory to await further processing or long-term storage in secondary memory.

KEY CONCEPT

The core of a computer is made up of main memory, which stores programs and data, and the CPU, which executes program instructions one at a time.

- 2.8 GHz Intel Pentium 4 processor
- 512 MB RAM
- 160 GB Hard Drive
- 48x CD-RW / DVD-ROM Combo Drive
- 17" Flat Screen Video Display with 1280 x 1024 resolution
- 56 Kb/s Modem

FIGURE 1.8 The hardware specification of a particular computer

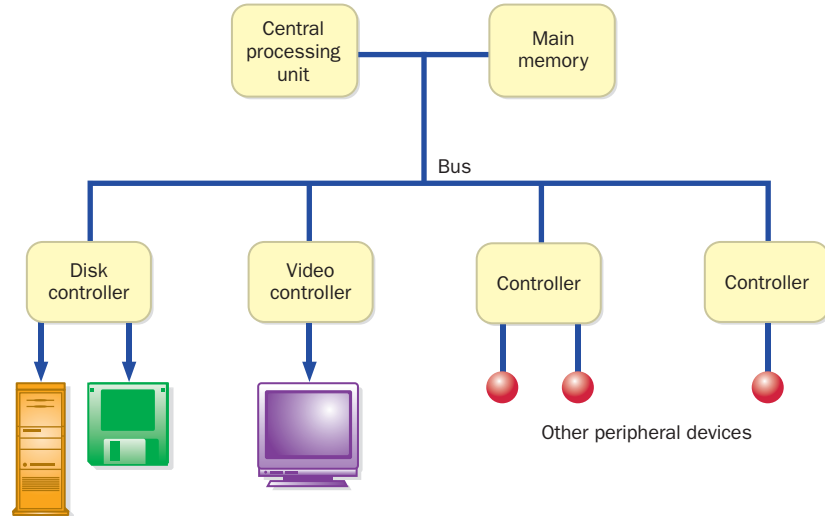


FIGURE 1.9 Basic computer architecture

Almost all devices in a computer system other than the CPU and main memory are called *peripherals*; they operate at the periphery, or outer edges, of the system (although they may be in the same box). Users don't interact directly with the CPU or main memory. Although they form the essence of the machine, the CPU and main memory would not be useful without peripheral devices.

Controllers are devices that coordinate the activities of specific peripherals. Every device has its own particular way of formatting and communicating data, and part of the controller's role is to handle these idiosyncrasies and isolate them from the rest of the computer hardware. Furthermore, the controller often handles much of the actual transmission of information, allowing the CPU to focus on other activities.

Input/output (I/O) devices and secondary memory devices are considered peripherals. Another category of peripherals includes *data transfer devices*, which allow information to be sent and received between computers. The computer specified in Figure 1.8 includes a data transfer device called a *modem*, which allows information to be sent across a telephone line. The modem in the example can transfer data at a maximum rate of 56 *kilobits* (Kb) per second, or approximately 56,000 *bits per second* (bps).

In some ways, secondary memory devices and data transfer devices can be thought of as I/O devices because they represent a source of information (input)

and a place to send information (output). For our discussion, however, we define I/O devices as those devices that allow the user to interact with the computer.

Input/Output Devices

Let's examine some I/O devices in more detail. The most common input devices are the keyboard and the mouse. Others include:

- *bar code readers*, such as the ones used at a grocery store checkout
- *joysticks*, often used for games and advanced graphical applications
- *microphones*, used by voice recognition systems that interpret simple voice commands
- *virtual reality devices*, such as gloves that interpret the movement of the user's hand
- *scanners*, which convert text, photographs, and graphics into machine-readable form

Monitors and printers are the most common output devices. Others include:

- *plotters*, which move pens across large sheets of paper (or vice versa)
- *speakers*, for audio output
- *goggles*, for virtual reality display

Some devices can provide both input and output capabilities. A *touch screen* system can detect the user touching the screen at a particular place. Software can then use the screen to display text and graphics in response to the user's touch. Touch screens are particularly useful in situations where the interface to the machine must be simple, such as at an information booth.

The computer described in Figure 1.8 includes a monitor with a 17-inch diagonal display area. It is a flat screen, which makes use of newer liquid crystal display (LCD) technology rather than the older cathode ray tube (CRT) monitors that take up much more space on a desk. A picture is represented in a computer by breaking it up into separate picture elements, or *pixels*. The monitor can display a grid of 1280 by 1024 pixels. Representing and managing graphical data is discussed in more detail in Chapter 2.

Main Memory and Secondary Memory

Main memory is made up of a series of small, consecutive *memory locations*, as shown in Figure 1.10. Associated with each memory location is a unique number called an *address*.

KEY CONCEPT

An address is a unique number associated with each memory location.

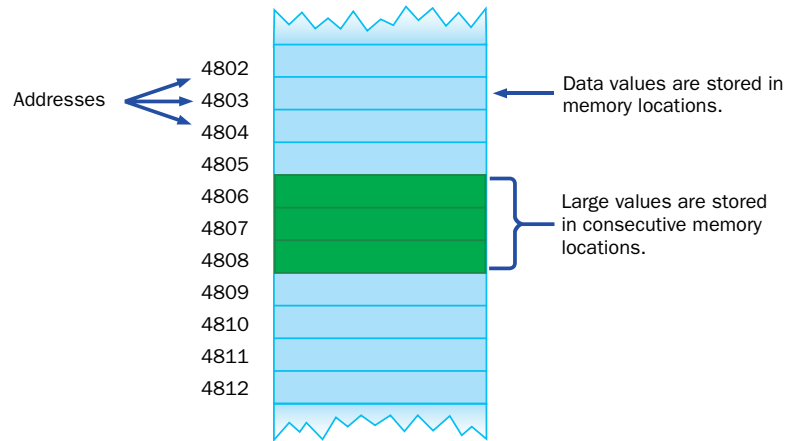


FIGURE 1.10 Memory locations

When data is stored in a memory location, it overwrites and destroys any information that was previously stored at that location. However, data is read from a memory location without affecting it.

On many computers, each memory location consists of eight bits, or one *byte*, of information. If we need to store a value that cannot be represented in a single byte, such as a large number, then multiple, consecutive bytes are used to store the data.

The *storage capacity* of a device such as main memory is the total number of bytes it can hold. Devices can store thousands or millions of bytes, so you should become familiar with larger units of measure. Because computer memory is based on the binary number system, all units of storage are powers of two. A *kilobyte* (KB) is 1024, or 2^{10} , bytes. Some larger units of storage are a *megabyte* (MB), a *gigabyte* (GB), and a *terabyte* (TB), as listed in Figure 1.11. It's usually easier to think about these capacities by rounding them off. For example, most computer users think of a kilobyte as approximately one thousand bytes, a megabyte as approximately one million bytes, and so forth.

Many personal computers have 512 megabytes or 1 gigabyte of main memory, or RAM, such as the system described in Figure 1.8 (we discuss RAM in more detail later in this chapter). A large main memory allows large programs, or multiple programs, to run efficiently because they don't have to retrieve information from secondary memory as often.

KEY CONCEPT

Main memory is volatile, meaning the stored information is maintained only as long as electric power is supplied.

Main memory is usually *volatile*, meaning that the information stored in it will be lost if its electric power supply is turned off. When you are working on a computer, you should often save your work onto

Unit	Symbol	Number of Bytes
byte		$2^0 = 1$
kilobyte	KB	$2^{10} = 1024$
megabyte	MB	$2^{20} = 1,048,576$
gigabyte	GB	$2^{30} = 1,073,741,824$
terabyte	TB	$2^{40} = 1,099,511,627,776$

FIGURE 1.11 Units of binary storage

a secondary memory device such as a disk in case the power goes out. Secondary memory devices are usually *nonvolatile*; the information is retained even if the power supply is turned off.

The most common secondary storage devices are hard disks and floppy disks. A high-density floppy disk can store 1.44 MB of information. The storage capacities of hard drives vary, but on personal computers, capacities typically range between 40 and 160 GB, such as in the system described in Figure 1.8.

A disk is a magnetic medium on which bits are represented as magnetized particles. A read/write head passes over the spinning disk, reading or writing information as appropriate. A hard disk drive might actually contain several disks in a vertical column with several read/write heads, such as the one shown in Figure 1.12.

To get an intuitive feel for how much information these devices can store, consider that all the information in this book, including pictures and formatting, requires about 7 MB of storage.

Magnetic tapes are also used as secondary storage but are considerably slower than disks because of the way information is accessed. A disk is a *direct access device* since the read/write head can move, in general, directly to the information needed. The terms *direct access* and *random access* are often used interchangeably. However, information on a tape can be accessed only after first getting past the intervening data. A tape must be rewound or fast-forwarded to get to the appropriate position. A tape is therefore considered a *sequential access device*. Tapes are usually used only to store information when it is no longer used frequently, or to provide a backup copy of the information on a disk.

Two other terms are used to describe memory devices: *random access memory* (RAM) and *read-only memory* (ROM). It's important to understand these terms




FIGURE 1.12 A hard disk drive with multiple disks and read/write heads

because they are used often, and their names can be misleading. The terms RAM and main memory are basically interchangeable, describing the memory where active programs and data are stored. ROM can refer to chips on the computer motherboard or to portable storage such as a compact disc. ROM chips typically store software called BIOS that provide the preliminary instructions needed when the computer is turned on initially. After information is stored on ROM, generally it is not altered (as the term *read-only* implies) during typical computer use. Both RAM and ROM are direct (or random) access devices.

KEY CONCEPT

The surface of a CD has both smooth areas and small pits. A pit represents a binary 1 and a smooth area represents a binary 0.

A *CD-ROM* is a portable secondary memory device. CD stands for compact disc. It is accurately called ROM because information is stored permanently when the CD is created and cannot be changed. Like its musical CD counterpart, a *CD-ROM* stores information in binary format. When the CD is initially created, a microscopic pit is pressed into the disc to represent a binary 1, and the disc is left smooth to represent a binary 0. The bits are read by shining a low-intensity laser beam onto the spinning disc. The laser beam reflects strongly from a smooth area on the disc but weakly from a pitted area. A sensor receiving the reflection determines whether each bit is a 1 or a 0 accordingly. A typical *CD-ROM*'s storage capacity is approximately 650 MB.

Variations on basic CD technology have emerged quickly. It is now common for a home computer to be equipped with a *CD-Recordable* (*CD-R*) drive. A *CD-R* can be used to create a CD for music or for general computer storage. Once created, you can use a *CD-R* disc in a standard CD player, but you can't change

the information on a CD-R disc once it has been “burned.” Music CDs that you buy in a store are pressed from a mold, whereas CD-Rs are burned with a laser.

A *CD-Rewritable* (CD-RW) disc can be erased and reused. They can be reused because the pits and flat surfaces of a normal CD are simulated on a CD-RW by coating the surface of the disc with a material that, when heated to one temperature becomes amorphous (and therefore nonreflective) and when heated to a different temperature becomes crystalline (and therefore reflective). The CD-RW media doesn’t work in all players, but CD-RW drives can create both CD-R and CD-RW discs.

KEY CONCEPT

A rewritable CD simulates the pits and smooth areas of a regular CD by using a coating that can be made amorphous or crystalline as needed.

CDs were initially a popular format for music; they later evolved to be used as a general computer storage device. Similarly, the *DVD* format was originally created for video and is now making headway as a general format for computer data. DVD once stood for digital video disc or digital versatile disc, but now the acronym generally stands on its own. A DVD has a tighter format (more bits per square inch) than a CD and can therefore store much more information. It is likely that DVD-ROMs eventually will replace CD-ROMs completely because there is a compatible migration path, meaning that a DVD drive can read a CD-ROM. Similar to CD-R and CD-RW, there are DVD-R and DVD-RW discs. The drive listed in Figure 1.8 allows the user to read and write CD-RW discs and read DVD-ROMs. This, of course, includes the ability to play music CDs and watch DVD videos.

The speed of a CD or DVD is expressed in multiples of x , which represents a data transfer speed of 153,600 bytes of data per second for a CD, and nine times that speed, or about 1.32 megabytes of data per second for a DVD. The drive described in Figure 1.8 has a maximum data access speed of 48 x , though it probably writes data at much slower speeds.

The capacity of storage devices changes continually as technology improves. A general rule in the computer industry suggests that storage capacity approximately doubles every 18 months. However, this progress eventually will slow down as capacities approach absolute physical limits.

The Central Processing Unit

The central processing unit (CPU) interacts with main memory to perform all fundamental processing in a computer. The CPU interprets and executes instructions, one after another, in a continuous cycle. It is made up of three important components, as shown in Figure 1.13. The *control unit* coordinates the processing steps, the *registers* provide a small amount of storage space in the CPU itself, and the *arithmetic/logic unit* performs calculations and makes decisions.

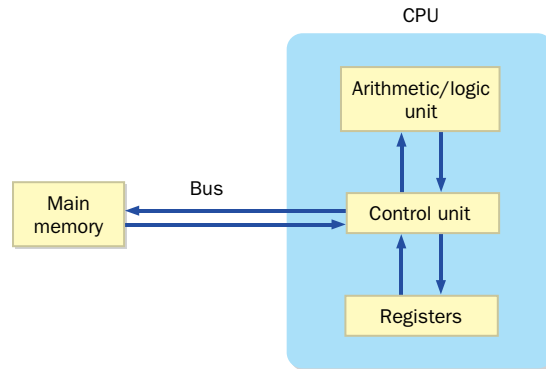


FIGURE 1.13 CPU components and main memory

The control unit coordinates the transfer of data and instructions between main memory and the registers in the CPU. It also coordinates the execution of the circuitry in the arithmetic/logic unit to perform operations on data stored in particular registers.

In most CPUs, some registers are reserved for special purposes. For example, the *instruction register* holds the current instruction being executed. The *program counter* is a register that holds the address of the next instruction to be executed. In addition to these and other special-purpose registers, the CPU also contains a set of general-purpose registers that are used for temporary storage of values as needed.

The concept of storing both program instructions and data together in main memory is the underlying principle of the *von Neumann architecture* of computer design, named after John von Neumann, who first advanced this programming concept in 1945. These computers continually follow the *fetch-decode-execute* cycle depicted in Figure 1.14. An instruction is fetched from main memory at the address stored in the program counter and is put into the instruction register. The program counter is incremented at this point to prepare for the next cycle. Then the instruction is decoded electronically to determine which operation to carry out. Finally, the control unit activates the correct circuitry to carry out the instruction, which may load a data value into a register or add two values together, for example.

KEY CONCEPT

The fetch-decode-execute cycle forms the foundation of computer processing.

The CPU is constructed on a chip called a *microprocessor*, a device that is part of the main circuit board of the computer. This board also contains ROM chips and communication sockets to which device controllers, such as the controller that manages the video display, can be connected.




FIGURE 1.14 The continuous fetch-decode-execute cycle

Another crucial component of the main circuit board is the *system clock*. The clock generates an electronic pulse at regular intervals, which synchronizes the events of the CPU. The rate at which the pulses occur is called the *clock speed*, and it varies depending on the processor. The computer described in Figure 1.8 includes a Pentium 4 processor that runs at a clock speed of 2.8 gigahertz (GHz), or approximately 2.8 billion pulses per second. The speed of the system clock provides a rough measure of how fast the CPU executes instructions. Similar to storage capacities, the speed of processors is constantly increasing with advances in technology.

SELF-REVIEW QUESTIONS *(see answers in Appendix N)*

SR 1.7 How many bytes are in each of the following?

- 3 KB
- 2 MB
- 4 GB

SR 1.8 How many bits are there in each of the following?

- 8 bytes
- 2 KB
- 4 MB

SR 1.9 The music on a CD is created using a sampling rate of 40,000 measurements per second. Each measurement is stored as a number that represents a specific voltage level. Suppose each of these numbers requires two bytes of storage space. How many MB does it take to represent one hour of music?

SR 1.10 What are the two primary hardware components in a computer? How do they interact?

SR 1.11 What is a memory address?

SR 1.12 What does volatile mean? Which memory devices are volatile and which are nonvolatile?

SR 1.13 Select the word from the following list that best matches each of the following phrases:

controller, CPU, main, modem, peripheral, RAM, register, ROM, secondary

- a. Almost all devices in a computer system, other than the CPU and the main memory, are categorized as this.
- b. A device that coordinates the activities of a peripheral device.
- c. Allows information to be sent across a phone line.
- d. This type of memory is usually volatile.
- e. This type of memory is usually nonvolatile.
- f. This term basically is interchangeable with the term “main memory.”
- g. Where the fundamental processing of a computer takes place.

1.3 Networks

KEY CONCEPT

A network consists of two or more computers connected together so that they can exchange information.

A single computer can accomplish a great deal, but connecting several computers together into networks can dramatically increase productivity and facilitate the sharing of information. A *network* is two or more computers connected together so they can exchange information. Using networks has become the normal mode of commercial computer operation. New technologies are emerging every day to capitalize on the connected environments of modern computer systems.

Figure 1.15 shows a simple computer network. One of the devices on the network is a printer, which allows any computer connected to the network to print a document on that printer. One of the computers on the network is designated as a *file server*, which is dedicated to storing programs and data that are needed by many network users. A file server usually has a large amount of secondary memory. When a network has a file server, each individual computer doesn't need its own copy of a program.

Network Connections

If two computers are directly connected, they can communicate in basically the same way that information moves across wires inside a single machine. When




FIGURE 1.15 A simple computer network

connecting two geographically close computers, this solution works well and is called a *point-to-point connection*. However, consider the task of connecting many computers together across large distances. If point-to-point connections are used, every computer is directly connected by a wire to every other computer in the network. A separate wire for each connection is not a workable solution because every time a new computer is added to the network, a new communication line will have to be installed for each computer already in the network. Furthermore, a single computer can handle only a small number of direct connections.

Figure 1.16 shows multiple point-to-point connections. Consider the number of communication lines that would be needed if two or three additional computers were added to the network.

Compare the diagrams in Figure 1.15 and Figure 1.16. All of the computers shown in Figure 1.15 share a single communication line. Each computer on the network has its own *network address*, which uniquely identifies it. These addresses are similar in concept to the addresses in main memory except that they identify individual computers on a network instead of individual memory locations inside a single computer. A message is sent across the line from one




FIGURE 1.16 Point-to-point connections

computer to another by specifying the network address of the computer for which it is intended.

KEY CONCEPT

Sharing a communication line creates delays, but it is cost effective and simplifies adding new computers to the network.

Sharing a communication line is cost effective and makes adding new computers to the network relatively easy. However, a shared line introduces delays. The computers on the network cannot use the communication line at the same time. They have to take turns sending information, which means they have to wait when the line is busy.

One technique to improve network delays is to divide large messages into segments, called *packets*, and then send the individual packets across the network intermixed with pieces of other messages sent by other users. The packets are collected at the destination and reassembled into the original message. This situation is similar to a group of people using a conveyor belt to move a set of boxes from one place to another. If only one person were allowed to use the conveyor belt at a time, and that person had a large number of boxes to move, the others would be waiting a long time before they could use it. By taking turns, each person can put one box on at a time, and they all can get their work done. It's not as fast as having a conveyor belt of your own, but it's not as slow as having to wait until everyone else is finished.

Local-Area Networks and Wide-Area Networks

A *local-area network* (LAN) is designed to span short distances and connect a relatively small number of computers. Usually a LAN connects the machines in only one building or in a single room. LANs are convenient to install and manage and are highly reliable. As computers became increasingly small and versatile, LANs became an inexpensive way to share information throughout an organization. However, having a LAN is like having a telephone system that allows you to call only the people in your own town. We need to be able to share information across longer distances.

KEY CONCEPT

A local-area network (LAN) is an effective way to share information and resources throughout an organization.

A *wide-area network* (WAN) connects two or more LANs, often across long distances. Usually one computer on each LAN is dedicated to handling the communication across a WAN. This technique relieves the other computers in a LAN from having to perform the details of long-distance communication. Figure 1.17 shows several LANs connected into a WAN. The LANs connected by a WAN are often owned by different companies or organizations, and might even be located in different countries.

The impact of networks on computer systems has been dramatic. Computing resources can now be shared among many users, and computer-based communication across the entire world is now possible. In fact, the use of networks is now so pervasive that some computers require network resources in order to operate.




FIGURE 1.17 LANs connected into a WAN

The Internet

Throughout the 1970s, a United States government organization called the Advanced Research Projects Agency (ARPA) funded several projects to explore network technology. One result of these efforts was the ARPANET, a WAN that eventually became known as the Internet. The *Internet* is a network of networks. The term Internet comes from the WAN concept of *internetworking*—connecting many smaller networks together.

From the mid 1980s through the present day, the Internet has grown incredibly. In 1983, there were fewer than 600 computers connected to the Internet. By the year 2000, that number had reached over 10 million. As more and more computers connect to the Internet, the task of keeping up with the larger number of users and heavier traffic has been difficult. New technologies have replaced the ARPANET several times since the initial development, each time providing more capacity and faster processing.

A *protocol* is a set of rules that governs how two things communicate. The software that controls the movement of messages across the Internet must conform to a set of protocols called TCP/IP (pronounced by spelling out the letters, T-C-P-I-P). TCP stands for *Transmission Control Protocol*, and IP stands for *Internet Protocol*. The IP software defines how information is formatted and transferred from the source to the destination. The TCP software handles problems such as

KEY CONCEPT

The Internet is a wide-area network (WAN) that spans the globe.

pieces of information arriving out of their original order or information getting lost, which can happen if too much information converges at one location at the same time.

KEY CONCEPT

Every computer connected to the Internet has an IP address that uniquely identifies it.

Every computer connected to the Internet has an *IP address* that uniquely identifies it among all other computers on the Internet. An example of an IP address is 204.192.116.2. Fortunately, the users of the Internet rarely have to deal with IP addresses. The Internet allows each computer to be given a name. Like IP addresses, the names must be unique. The Internet name of a computer is often referred to as its *Internet address*. Two examples of Internet addresses are spencer.villanova.edu and kant.gestalt-llc.com.

The first part of an Internet address is the local name of a specific computer. The rest of the address is the *domain name*, which indicates the organization to which the computer belongs. For example, villanova.edu is the domain name for the network of computers at Villanova University, and spencer is the name of a particular computer on that campus. Because the domain names are unique, many organizations can have a computer named spencer without confusion. Individual departments might be assigned *subdomains* that are added to the basic domain name to uniquely distinguish their set of computers within the larger organization. For example, the csc.villanova.edu subdomain is devoted to the Department of Computing Sciences at Villanova University.

The last part of each domain name, called a *top-level domain* (TLD), usually indicates the type of organization to which the computer belongs. The TLD edu indicates an educational institution. The TLD com refers to a commercial business. For example, gestalt-llc.com refers to Gestalt, LLC, a company specializing in software technologies. Another common TLD is org, used mostly by nonprofit organizations. Many computers, especially those outside of the United States, use a TLD that denotes the country of origin, such as uk for the United Kingdom or au for Australia. Recently, in response to a diminishing supply of domain names, some new top-level domain names have been created, such as biz, info, and name.

When an Internet address is referenced, it gets translated to its corresponding IP address, which is used from that point on. The software that does this translation is called the *Domain Name System* (DNS). Each organization connected to the Internet operates a *domain server* that maintains a list of all computers at that organization and their IP addresses. It works somewhat like telephone directory assistance in that you provide the name, and the domain server gives back a number. If the local domain server does not have the IP address for the name, it contacts another domain server that does.

The Internet has revolutionized computer processing. Initially, the primary use of interconnected computers was to send electronic mail, but Internet

capabilities continue to improve. One of the most significant uses of the Internet is the World Wide Web.

The World Wide Web

The Internet gives us the capability to exchange information. The *World Wide Web* (also known as WWW or simply the Web) makes the exchange of information easy. Web software provides a common user interface through which many different types of information can be accessed with the click of a mouse.

KEY CONCEPT

The World Wide Web is software that makes sharing information across a network easy.

The Web is based on the concepts of hypertext and hypermedia. The term *hypertext* was first used in 1965 to describe a way to organize information so that the flow of ideas was not constrained to a linear progression. In fact, that concept was entertained as a way to manage large amounts of information as early as the 1940s. Researchers on the Manhattan Project, who were developing the first atomic bomb, envisioned such an approach. The underlying idea is that documents can be linked at various points according to natural relationships so that the reader can jump from one document to another, following the appropriate path for that reader's needs. When other media components are incorporated, such as graphics, sound, animations, and video, the resulting organization is called *hypermedia*.

The terms Internet and World Wide Web are sometimes used interchangeably, but there are important differences between the two. The Internet makes it possible to communicate via computers around the world. The Web makes that communication a straightforward and enjoyable activity. The Web is essentially a distributed information service and is based on a set of software applications. It is not a network. Although it is used effectively with the Internet, it is not inherently bound to it. The Web can be used on a LAN that is not connected to any other network or even on a single machine to display HTML documents.

A *browser* is a software tool that loads and formats Web documents for viewing. *Mosaic*, the first graphical interface browser for the Web, was released in 1993. The designer of a Web document defines to other Web information that might be anywhere on the Internet. Some of the people who developed Mosaic went on to found the Netscape Communications Corporation and create the Netscape Navigator browser. It is currently one of the most popular systems for accessing information on the Web. Microsoft's Internet Explorer is another popular browser.

A computer dedicated to providing access to Web documents is called a *Web server*. Browsers load and interpret documents provided by a Web server. Many such documents are formatted using the *HyperText Markup Language* (HTML). The Java

programming language has an intimate relationship with Web processing because links to Java programs can be embedded in HTML documents and executed through Web browsers. We explore this relationship in more detail in Chapter 2.

Uniform Resource Locators

Information on the Web is found by identifying a *Uniform Resource Locator* (URL). A URL uniquely specifies documents and other information for a browser to obtain and display. The following is an example URL:

KEY CONCEPT

A URL uniquely specifies documents and other information found on the Web for a browser to obtain and display.

`http://www.google.com`

The Web site at this particular URL is a popular *search engine*, which enables you to search the Web for information using particular words or phrases.

A URL contains several pieces of information. The first piece is a protocol, which determines the way the browser transmits and processes information. The second piece is the Internet address of the machine on which the document is stored. The third piece of information is the file name of interest. If no file name is given, as is the case with the Google URL, the Web server usually provides a default page (such as `index.html`).

Let's look at another example URL:

`http://www.gestalt-llc.com/vision.html`

In this URL, the protocol is `http`, which stands for *HyperText Transfer Protocol*. The machine referenced is `www` (a typical reference to a Web server), found at domain `gestalt-llc.com`. Finally, `vision.html` is a file to be transferred to the browser for viewing. Many other forms for URLs exist, but this form is the most common.

SELF-REVIEW QUESTIONS *(see answers in Appendix N)*

- SR 1.14 What is a file server?
- SR 1.15 What is the total number of communication lines needed for a fully connected point-to-point network of five computers? Six computers?
- SR 1.16 Describe a benefit of having computers on a network share a communication line. Describe a cost/drawback of sharing a communication line.
- SR 1.17 What is the origin of the word Internet?
- SR 1.18 The TCP/IP set of protocols describes communication rules for software that uses the Internet. What does TCP stand for? What does IP stand for?

SR 1.19 Explain the parts of the following URLs:

- a. duke.csc.villanova.edu/jss/examples.html
- b. java.sun.com/products/index.html

1.4 The Java Programming Language

Let's now turn our attention to the software that makes a computer system useful. A program is written in a particular *programming language* that uses specific words and symbols to express the problem solution. A programming language defines a set of rules that determines exactly how a programmer can combine the words and symbols of the language into *programming statements*, which are the instructions that are carried out when the program is executed.

Since the inception of computers, many programming languages have been created. We use the Java language in this book to demonstrate various programming concepts and techniques. Although our main goal is to learn these underlying software development concepts, an important side effect will be to become proficient in the development of Java programs.

Java is a relatively new programming language compared to many others. It was developed in the early 1990s by James Gosling at Sun Microsystems. Java was introduced to the public in 1995 and has gained tremendous popularity since.

Java has undergone various changes since its creation. The most recent Java technology is generally referred to as the *Java 2 Platform*, which is organized into three major groups:

- Java 2 Platform, Standard Edition (J2SE)
- Java 2 Platform, Enterprise Edition (J2EE)
- Java 2 Platform, Micro Edition (J2ME)

This book focuses on the Standard Edition, which, as the name implies, is the mainstream version of the language and associated tools. Furthermore, this book is based on the most recent version of the Standard Edition, which is J2SE 5.0.

Some parts of early Java technologies have been *deprecated*, which means they are considered old-fashioned and should not be used. When it is important, we point out deprecated elements and discuss their preferred alternatives.

One reason Java got some initial attention was because it was the first programming language to deliberately embrace the concept of writing programs that can be executed using the Web. The original hype about Java's Web capabilities initially obscured the far more important features that make it a useful general-purpose programming language.

KEY CONCEPT

This book focuses on the principles of object-oriented programming.

Java is an *object-oriented programming language*. Objects are the fundamental elements that make up a program. The principles of object-oriented software development are the cornerstone of this book. We explore object-oriented programming concepts later in this chapter and throughout the rest of the book.

The Java language is accompanied by a library of extra software that we can use when developing programs. This software, referred to as the Java *standard class library*, provides the ability to create graphics, communicate over networks, and interact with databases, among many other features. The standard library that supports Java programming is huge and quite versatile. Although we won't be able to cover all aspects of the library, we will explore many of them.

Java is used in commercial environments all over the world. It is one of the fastest growing programming technologies of all time. So not only is it a good language in which to learn programming concepts, it is also a practical language that will serve you well in the future.

A Java Program

Let's look at a simple but complete Java program. The program in Listing 1.1 prints two sentences to the screen. This particular program prints a quote by Abraham Lincoln. The output is shown below the program listing.

All Java applications have a similar basic structure. Despite its small size and simple purpose, this program contains several important features. Let's carefully dissect it and examine its pieces.

KEY CONCEPT

Comments do not affect a program's processing; instead, they serve to facilitate human comprehension.

The first few lines of the program are comments, which start with the `//` symbols and continue to the end of the line. Comments don't affect what the program does but are included to make the program easier to understand by humans. Programmers can and should include comments as needed throughout a program to clearly identify the purpose of the program and describe any special processing. Any written comments or documents, including a user's guide and technical references, are called *documentation*. Comments included in a program are called *inline documentation*.

The rest of the program is a *class definition*. This class is called `Lincoln`, though we could have named it just about anything we wished. The class definition runs from the first opening brace (`{`) to the final closing brace (`}`) on the last line of the program. All Java programs are defined using class definitions.

Inside the class definition are some more comments describing the purpose of the `main` method, which is defined directly below the comments. A *method* is a group of programming statements that is given a name. In this case, the name of



Video Note
Overview of program elements.

LISTING 1.1

```
/**
 * Lincoln.java      Author: Lewis/Loftus
 *
 * Demonstrates the basic structure of a Java application.
 */

public class Lincoln
{
    //-----
    // Prints a presidential quote.
    //-----
    public static void main (String[] args)
    {
        System.out.println ("A quote by Abraham Lincoln:");

        System.out.println ("Whatever you are, be a good one.");
    }
}
```

OUTPUT

```
A quote by Abraham Lincoln:
Whatever you are, be a good one.
```

the method is `main` and it contains only two programming statements. Like a class definition, a method is also delimited by braces.

All Java applications have a `main` method, which is where processing begins. Each programming statement in the `main` method is executed, one at a time in order, until the end of the method is reached. Then the program ends, or *terminates*. The `main` method definition in a Java program is always preceded by the words `public`, `static`, and `void`, which we examine later in the text. The use of `String` and `args` does not come into play in this particular program. We describe these later also.

The two lines of code in the `main` method invoke another method called `println` (pronounced print line). We *invoke*, or *call*, a method when we want it to execute. The `println` method prints the specified characters to the screen. The characters to be printed are represented as a *character string*, enclosed in double quote characters (`"`). When the program is executed, it calls the `println` method to print the first statement, calls it again to print the second statement, and then, because that is the last line in the `main` method, the program terminates.

The code executed when the `println` method is invoked is not defined in this program. The `println` method is part of the `System.out` object, which is part of the Java standard class library. It's not technically part of the Java language, but is always available for use in any Java program. We explore the `println` method in more detail in Chapter 2.

Comments

Let's examine comments in more detail. Comments are the only language feature that allows programmers to compose and communicate their thoughts independent of the code. Comments should provide insight into the programmer's original intent. A program is often used for many years, and often many modifications are made to it over time. The original programmer often will not remember the details of a particular program when, at some point in the future, modifications are required. Furthermore, the original programmer is not always available to make the changes; thus, someone completely unfamiliar with the program will need to understand it. Good documentation is therefore essential.

As far as the Java programming language is concerned, the content of comments can be any text whatsoever. Comments are ignored by the computer; they do not affect how the program executes.

The comments in the `Lincoln` program represent one of two types of comments allowed in Java. The comments in `Lincoln` take the following form:

```
// This is a comment.
```

This type of comment begins with a double slash (`//`) and continues to the end of the line. You cannot have any characters between the two slashes. The computer ignores any text after the double slash to the end of the line. A comment can follow code on the same line to document that particular line, as in the following example:

```
System.out.println ("Monthly Report"); // always use this title
```

The second form a Java comment may have is the following:

```
/* This is another comment. */
```

This comment type does not use the end of a line to indicate the end of the comment. Anything between the initiating slash-asterisk (`/*`) and the terminating asterisk-slash (`*/`) is part of the comment, including the invisible *newline* character that represents the end of a line. Therefore, this type of comment can extend over multiple lines. No space can be between the slash and the asterisk.

If there is a second asterisk following the `/*` at the beginning of a comment, the content of the comment can be used to automatically generate external documentation about your program by using a tool called *javadoc*. More information about javadoc is given in Appendix I.

The two basic comment types can be used to create various documentation styles, such as:

```
// This is a comment on a single line.

//-----
// Some comments such as those above methods or classes
// deserve to be blocked off to focus special attention
// on a particular aspect of your code. Note that each of
// these lines is technically a separate comment.
//-----

/*
   This is one comment
   that spans several lines.
*/
```

Programmers often concentrate so much on writing code that they focus too little on documentation. You should develop good commenting practices and follow them habitually. Comments should be well written, often in complete sentences. They should not belabor the obvious but should provide appropriate insight into the intent of the code. The following examples are *not* good comments:

```
System.out.println ("hello"); // prints hello
System.out.println ("test"); // change this later
```

The first comment paraphrases the obvious purpose of the line and does not add any value to the statement. It is better to have no comment than a useless one. The second comment is ambiguous. What should be changed later? When is later? Why should it be changed?

KEY CONCEPT

Inline documentation should provide insight into your code. It should not be ambiguous or belabor the obvious.

Identifiers and Reserved Words

The various words used when writing programs are called *identifiers*. The identifiers in the `Lincoln` program are `class`, `Lincoln`, `public`, `static`, `void`, `main`, `String`, `args`, `System`, `out`, and `println`. These fall into three categories:

- words that we make up when writing a program (`Lincoln` and `args`)
- words that another programmer chose (`String`, `System`, `out`, `println`, and `main`)
- words that are reserved for special purposes in the language (`class`, `public`, `static`, and `void`)

While writing the program, we simply chose to name the class `Lincoln`, but we could have used one of many other possibilities. For example, we could have called it `Quote`, or `Abe`, or `GoodOne`. The identifier `args` (which is short for arguments) is often used in the way we use it in `Lincoln`, but we could have used just about any other identifier in its place.

The identifiers `String`, `System`, `out`, and `println` were chosen by other programmers. These words are not part of the Java language. They are part of the Java standard library of predefined code, a set of classes and methods that someone has already written for us. The authors of that code chose the identifiers in that code—we're just making use of them.

Reserved words are identifiers that have a special meaning in a programming language and can only be used in predefined ways. A reserved word cannot be used for any other purpose, such as naming a class or method. In the `Lincoln` program, the reserved words used are `class`, `public`, `static`, and `void`. Throughout the book, we show Java reserved words in blue type. Figure 1.18 lists all of the Java reserved words in alphabetical order. The words marked with an asterisk are reserved for possible future use in later versions of the language but currently have no meaning in Java.


An identifier that we make up for use in a program can be composed of any combination of letters, digits, the underscore character (`_`), and the dollar sign (`$`), but it cannot begin with a digit. Identifiers may be of any length. Therefore, `total`, `label17`, `nextStockItem`, `NUM_BOXES`, and `$amount` are all valid identifiers, but `4th_word` and `coin#value` are not valid.

Both uppercase and lowercase letters can be used in an identifier, and the difference is important. Java is *case sensitive*, which means that two identifier names that differ only in the case of their letters are considered to be different identifiers.

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

FIGURE 1.18 Java reserved words

Identifier



An identifier is a letter followed by zero or more letters and digits. A Java Letter includes the 26 English alphabetic characters in both uppercase and lowercase, the \$ and _ (underscore) characters, as well as alphabetic characters from other languages. A Java Digit includes the digits 0 through 9.

Examples:

```
total
MAX_HEIGHT
num1
Keyboard
System
```

Therefore, `total`, `Total`, `toTaL`, and `TOTAL` are all different identifiers. As you can imagine, it is not a good idea to use multiple identifiers that differ only in their case, because they can be easily confused.

Although the Java language doesn't require it, using a consistent case format for each kind of identifier makes your identifiers easier to understand. There are various Java conventions regarding identifiers that should be followed, though technically they don't have to be. For example, we use *title case* (uppercase for the first letter of each word) for class names. Throughout the text, we describe the preferred case style for each type of identifier when it is first encountered.

While an identifier can be of any length, you should choose your names carefully. They should be descriptive but not verbose. You should avoid meaningless names such as `a` or `x`. An exception to this rule can be made if the short name is actually descriptive, such as using `x` and `y` to represent (x, y) coordinates on a two-dimensional grid. Likewise, you should not use unnecessarily long names, such as the identifier `theCurrentItemBeingProcessed`. The name `currentItem` would serve just as well. As you might imagine, the use of identifiers that are verbose is a much less prevalent problem than the use of names that are not descriptive.

You should always strive to make your programs as readable as possible. Therefore, you should always be careful when abbreviating words. You might think

KEY CONCEPT

Java is case sensitive. The uppercase and lowercase versions of a letter are distinct.

KEY CONCEPT

Identifier names should be descriptive and readable.

`curStVal` is a good name to represent the current stock value, but another person trying to understand the code may have trouble figuring out what you meant. It might not even be clear to you two months after writing it.

A *name* in Java is a series of identifiers separated by the dot (period) character. The name `System.out` is the way we designate the object through which we invoked the `println` method. Names appear quite regularly in Java programs.

White Space

All Java programs use *white space* to separate the words and symbols used in a program. White space consists of blanks, tabs, and newline characters. The phrase white space refers to the fact that, on a white sheet of paper with black printing, the space between the words and symbols is white. The way a programmer uses white space is important because it can be used to emphasize parts of the code and can make a program easier to read.

KEY CONCEPT

Appropriate use of white space makes a program easier to read and understand.

Except when it's used to separate words, the computer ignores white space. It does not affect the execution of a program. This fact gives programmers a great deal of flexibility in how they format a program. The lines of a program should be divided in logical places and certain lines should be indented and aligned so that the program's underlying structure is clear.

Because white space is ignored, we can write a program in many different ways. For example, taking white space to one extreme, we could put as many words as possible on each line. The code in Listing 1.2, the `Lincoln2` program, is formatted quite differently from `Lincoln` but prints the same message.

LISTING 1.2

```
//*****
//  Lincoln2.java      Author: Lewis/Loftus
//
//  Demonstrates a poorly formatted, though valid, program.
//*****

Lincoln2{  main(String[]args){
System.out.println("A quote by Abraham Lincoln:");
System.out.println("Whatever you are, be a good one.");}}
```

OUTPUT

```
A quote by Abraham Lincoln:
Whatever you are, be a good one.
```

Taking white space to the other extreme, we could write almost every word and symbol on a different line with varying amounts of spaces, such as `Lincoln3`, shown in Listing 1.3.

All three versions of `Lincoln` are technically valid and will execute in the same way, but they are radically different from a reader's point of view. Both of the latter examples show poor style and make the program difficult to understand. You may be asked to adhere to particular guidelines when you write your programs. A software development company often has a programming style policy that it requires its programmers to follow. In any case, you should adopt and consistently use a set of style guidelines that increase the readability of your code.

KEY CONCEPT

You should adhere to a set of guidelines that establish the way you format and document your programs.

LISTING 1.3

```

//*****
//  Lincoln3.java      Author: Lewis/Loftus
//
//  Demonstrates another valid program that is poorly formatted.
//*****

        public      class
    Lincoln3
    {
        public
    static
    void
    main
    (
String
    String
        []
        args
        )
    {
    System.out.println
    ("A quote by Abraham Lincoln:"
    );
        System.out.println
        (
            "Whatever you are, be a good one."
        )
    ;
    }
    }

```

OUTPUT

```

A quote by Abraham Lincoln:
Whatever you are, be a good one.

```

SELF-REVIEW QUESTIONS *(see answers in Appendix N)*

SR 1.20 When was the Java programming language developed? By whom? When was it introduced to the public?

SR 1.21 Where does processing begin in a Java application?

SR 1.22 What do you predict would be the result of the following line in a Java program?

```
System.out.println("Hello"); // prints hello
```

SR 1.23 What do you predict would be the result of the following line in a Java program?

```
// prints hello System.out.println("Hello");
```

SR 1.24 Which of the following are not valid Java identifiers? Why?

- a. RESULT
- b. result
- c. 12345
- d. x12345y
- e. black&white
- f. answer_7

SR 1.25 Suppose a program requires an identifier to represent the sum of the test scores of a class of students. For each of the following names, state whether or not each is a good name to use for the identifier. Explain your answers.

- a. x
- b. scoreSum
- c. sumOfTheTestScoresOfTheStudents
- d. smTstScr

SR 1.26 What is white space? How does it affect program execution? How does it affect program readability?

1.5 Program Development

The process of getting a program running involves various activities. The program has to be written in the appropriate programming language, such as Java. That program has to be translated into a form that the computer can execute. Errors can occur at various stages of this process and must be fixed. Various software

tools can be used to help with all parts of the development process as well. Let's explore these issues in more detail.

Programming Language Levels

Suppose a particular person is giving travel directions to a friend. That person might explain those directions in any one of several languages, such as English, Russian, or Italian. The directions are the same no matter which language is used to explain them, but the manner in which the directions are expressed is different. The friend must be able to understand the language being used in order to follow the directions.

Similarly, a problem can be solved by writing a program in one of many programming languages, such as Java, Ada, C, C++, C#, Pascal, and Smalltalk. The purpose of the program is essentially the same no matter which language is used, but the particular statements used to express the instructions, and the overall organization of those instructions, vary with each language. A computer must be able to understand the instructions in order to carry them out.

Programming languages can be categorized into the following four groups. These groups basically reflect the historical development of computer languages.

- machine language
- assembly language
- high-level languages
- fourth-generation languages

In order for a program to run on a computer, it must be expressed in that computer's *machine language*. Each type of CPU has its own language. For that reason, we can't run a program specifically written for a Sun Workstation, with its Sparc processor, on a Dell PC, with its Intel processor.

Each machine language instruction can accomplish only a simple task. For example, a single machine language instruction might copy a value into a register or compare a value to zero. It might take four separate machine language instructions to add two numbers together and to store the result. However, a computer can do millions of these instructions in a second, and therefore many simple commands can be executed quickly to accomplish complex tasks.

Machine language code is expressed as a series of binary digits and is extremely difficult for humans to read and write. Originally, programs were entered into the computer by using switches or some similarly tedious method. Early programmers found these techniques to be time consuming and error prone.

KEY CONCEPT

All programs must be translated to a particular CPU's machine language in order to be executed.

These problems gave rise to the use of *assembly language*, which replaced binary digits with *mnemonics*, short English-like words that represent commands or data. It is much easier for programmers to deal with words than with binary digits. However, an assembly language program cannot be executed directly on a computer. It must first be translated into machine language.

Generally, each assembly language instruction corresponds to an equivalent machine language instruction. Therefore, similar to machine language, each assembly language instruction accomplishes only a simple operation. Although assembly language is an improvement over machine code from a programmer's perspective, it is still tedious to use. Both assembly language and machine language are considered *low-level languages*.

KEY CONCEPT

High-level languages allow a programmer to ignore the underlying details of machine language.

Today, most programmers use a *high-level language* to write software. A high-level language is expressed in English-like phrases, and thus is easier for programmers to read and write. A single high-level language programming statement can accomplish the equivalent of many—perhaps hundreds—of machine language instructions. The term high-level refers to the fact that the programming statements are expressed in a way that is far removed from the machine language that is ultimately executed. Java is a high-level language, as are Ada, C++, Smalltalk, and many others.

Figure 1.19 shows equivalent expressions in a high-level language, assembly language, and machine language. The expressions add two numbers together. The assembly language and machine language in this example are specific to a Sparc processor.

The high-level language expression in Figure 1.19 is readable and intuitive for programmers. It is similar to an algebraic expression. The equivalent assembly

High-Level Language	Assembly Language	Machine Language
a + b	ld [%fp-20], %o0	...
	ld [%fp-24], %o1	1101 0000 0000 0111
	add %o0, %o1, %o0	1011 1111 1110 1000
		1101 0010 0000 0111
		1011 1111 1110 1000
		1001 0000 0000 0000
		...

FIGURE 1.19 A high-level expression and its assembly language and machine language equivalent

language code is somewhat readable, but it is more verbose and less intuitive. The machine language is basically unreadable and much longer. In fact, only a small portion of the binary machine code to add two numbers together is shown in Figure 1.19. The complete machine language code for this particular expression is over 400 bits long.

A high-level language insulates programmers from needing to know the underlying machine language for the processor on which they are working. But high-level language code must be translated into machine language in order to be executed.

Some programming languages are considered to operate at an even higher level than high-level languages. They might include special facilities for automatic report generation or interaction with a database. These languages are called *fourth-generation languages*, or simply 4GLs, because they followed the first three generations of computer programming: machine, assembly, and high-level.

Editors, Compilers, and Interpreters

Several special-purpose programs are needed to help with the process of developing new programs. They are sometimes called software tools because they are used to build programs. Examples of basic software tools include an editor, a compiler, and an interpreter.

Initially, you use an *editor* as you type a program into a computer and store it in a file. There are many different editors with many different features. You should become familiar with the editor you will use regularly because it can dramatically affect the speed at which you enter and modify your programs.

Figure 1.20 shows a very basic view of the program development process. After editing and saving your program, you attempt to translate it from high-level code into a form that can be executed. That translation may result in errors, in which case you return to the editor to make changes to the code to fix the problems. Once the translation occurs successfully, you can execute the program and evaluate the results. If the results are not what you want, or if you want to enhance your existing program, you again return to the editor to make changes.

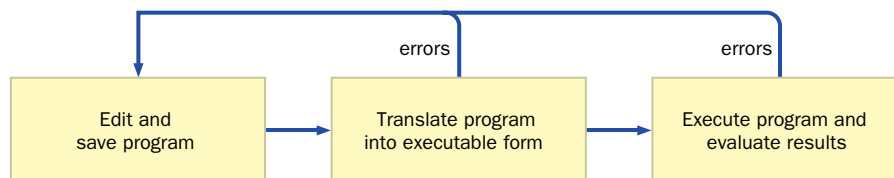


FIGURE 1.20 Editing and running a program

The translation of source code into (ultimately) machine language for a particular type of CPU can occur in a variety of ways. A *compiler* is a program that translates code in one language to equivalent code in another language. The original code is called *source code*, and the language into which it is translated is called the *target language*. For many traditional compilers, the source code is translated directly into a particular machine language. In that case, the translation process occurs once (for a given version of the program), and the resulting executable program can be run whenever needed.

An *interpreter* is similar to a compiler but has an important difference. An interpreter interweaves the translation and execution activities. A small part of the source code, such as one statement, is translated and executed. Then another statement is translated and executed, and so on. One advantage of this technique is that it eliminates the need for a separate compilation phase. However, the program generally runs more slowly because the translation process occurs during each execution.

KEY CONCEPT

A Java compiler translates Java source code into Java bytecode, a low-level, architecture-neutral representation of the program.

The process generally used to translate and execute Java programs combines the use of a compiler and an interpreter. This process is pictured in Figure 1.21. The Java compiler translates Java source code into Java *bytecode*, which is a representation of the program in a low-level form similar to machine language code. The Java interpreter reads Java bytecode and executes it on a specific machine. Another compiler could translate the bytecode into a particular machine language for efficient execution on that machine.

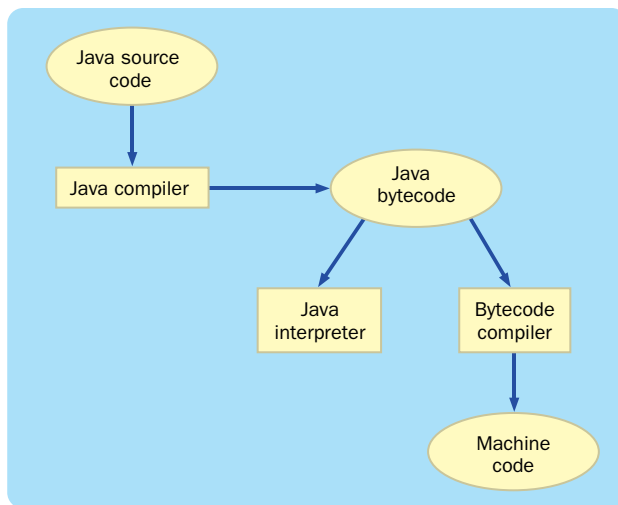


FIGURE 1.21 The Java translation and execution process

The difference between Java bytecode and true machine language code is that Java bytecode is not tied to any particular processor type. This approach has the distinct advantage of making Java *architecture neutral*, and therefore easily portable from one machine type to another. The only restriction is that there must be a Java interpreter or a bytecode compiler for each processor type on which the Java bytecode is to be executed.

Since the compilation process translates the high-level Java source code into a low-level representation, the interpretation process is more efficient than interpreting high-level code directly. Executing a program by interpreting its bytecode is still slower than executing machine code directly, but it is fast enough for most applications. Note that for efficiency, Java bytecode could be compiled into machine code.

Development Environments

A software *development environment* is the set of tools used to create, test, and modify a program. Some development environments are available for free while others, which may have advanced features, must be purchased. Some environments are referred to as *integrated development environments* (IDEs) because they integrate various tools into one software program.

Any development environment will contain certain key tools, such as a Java compiler and interpreter. Some will include a *debugger*, which helps you find errors in a program. Other tools that may be included are documentation generators, archiving tools, and tools that help you visualize your program structure.

Sun Microsystems, the creator of the Java programming language, provides the Java *Software Development Kit* (SDK), which is sometimes referred to simply as the *Java Development Kit* (JDK). The SDK can be downloaded free of charge for various hardware platforms from Sun's Java Web site, java.sun.com, and is also included on the CD that accompanies this book.

The SDK tools are not an integrated environment. The commands for compilation and interpretation are executed on the command line. That is, the SDK does not have a GUI. It also does not include an editor, although any editor that can save a document as simple text can be used.

Sun also has a Java IDE called NetBeans (www.netbeans.org) that incorporates the development tools of the SDK into one convenient GUI-based program. IBM promotes a similar IDE called Eclipse (www.eclipse.org). Both NetBeans and Eclipse are *open source* projects, meaning that they are developed by a wide collection of programmers and are available for free.



Video Note
Comparison of Java IDEs.

KEY CONCEPT

Many different development environments exist to help you create and modify Java programs.

A research group at Auburn University has developed jGRASP, a free Java IDE that is included on the CD that accompanies this book. It can also be downloaded from www.jgrasp.com. In addition to fundamental development tools, jGRASP contains tools that graphically display program elements.

Various other Java development environments are available. A Web search will unveil dozens of them. The choice of which development environment to use is important. The more you know about the capabilities of your environment, the more productive you can be during program development.

Syntax and Semantics

Each programming language has its own unique *syntax*. The syntax rules of a language dictate exactly how the vocabulary elements of the language can be combined to form statements. These rules must be followed in order to create a program. We've already discussed several Java syntax rules. For instance, the fact that an identifier cannot begin with a digit is a syntax rule. The fact that braces are used to begin and end classes and methods is also a syntax rule. Appendix L formally defines the basic syntax rules for the Java programming language, and specific rules are highlighted throughout the text.

During compilation, all syntax rules are checked. If a program is not syntactically correct, the compiler will issue error messages and will not produce byte-code. Java has a similar syntax to the programming languages C and C++, and therefore the look and feel of the code is familiar to people with a background in those languages.

The *semantics* of a statement in a programming language define what will happen when that statement is executed. Programming languages are generally unambiguous, which means the semantics of a program are well defined. That is, there is one and only one interpretation for each statement. On the other hand, the *natural languages* that humans use to communicate, such as English and Italian, are full of ambiguities. A sentence can often have two or more different meanings. For example, consider the following sentence:

Time flies like an arrow.

KEY CONCEPT

Syntax rules dictate the form of a program. Semantics dictate the meaning of the program statements.

The average human is likely to interpret this sentence as a general observation: that time moves quickly in the same way that an arrow moves quickly. However, if we interpret the word *time* as a verb (as in “run the 50-yard dash and I’ll time you”) and the word *flies* as a noun (the plural of fly), the interpretation changes completely. We know that arrows don’t time things, so we wouldn’t normally interpret the sentence that way, but it is a valid interpretation of the words in the sentence. A computer

would have a difficult time trying to determine which meaning is intended. Moreover, this sentence could describe the preferences of an unusual insect known as a “time fly,” which might be found near an archery range. After all, fruit flies like a banana.

The point is that one specific English sentence can have multiple valid meanings. A computer language cannot allow such ambiguities to exist. If a programming language instruction could have two different meanings, a computer would not be able to determine which one should be carried out.

Errors

Several different kinds of problems can occur in software, particularly during program development. The term computer error is often misused and varies in meaning depending on the situation. From a user’s point of view, anything that goes awry when interacting with a machine can be called a computer error. For example, suppose you charged a \$23 item to your credit card, but when you received the bill, the item was listed at \$230. After you have the problem fixed, the credit card company apologizes for the “computer error.” Did the computer arbitrarily add a zero to the end of the number, or did it perhaps multiply the value by 10? Of course not. A computer follows the commands we give it and operates on the data we provide. If our programs are wrong or our data inaccurate, then we cannot expect the results to be correct. A common phrase used to describe this situation is “garbage in, garbage out.”

You will encounter three kinds of errors as you develop programs:

- compile-time error
- run-time error
- logical error

The compiler checks to make sure you are using the correct syntax. If you have any statements that do not conform to the syntactic rules of the language, the compiler will produce a *syntax error*. The compiler also tries to find other problems, such as the use of incompatible types of data. The syntax might be technically correct, but you may be attempting to do something that the language doesn’t semantically allow. Any error identified by the compiler is called a *compile-time error*. If a compile-time error occurs, an executable version of the program is not created.

The second kind of problem occurs during program execution. It is called a *run-time error* and causes the program to terminate abnormally. For example, if

KEY CONCEPT

The programmer is responsible for the accuracy and reliability of a program.

KEY CONCEPT

A Java program must be syntactically correct or the compiler will not produce bytecode.

**Video Note**

Examples of various error types.

we attempt to divide by zero, the program will “crash” and halt execution at that point. Because the requested operation is undefined, the system simply abandons its attempt to continue processing your program. The best programs are *robust*; that is, they avoid as many run-time errors as possible. For example, the program code could guard against the possibility of dividing by zero and handle the situation appropriately if it arises. In Java, many run-time problems are called *exceptions* that can be caught and dealt with accordingly.

The third kind of software problem is a *logical error*. In this case, the software compiles and executes without complaint, but it produces incorrect results. For example, a logical error occurs when a value is calculated incorrectly or when a graphical button does not appear in the correct place. A programmer must test the program thoroughly, comparing the expected results to those that actually occur. When defects are found, they must be traced back to the source of the problem in the code and corrected. The process of finding and correcting defects in a program is called *debugging*. Logical errors can manifest themselves in many ways, and the actual root cause might be difficult to discover.

SELF-REVIEW QUESTIONS *(see answers in Appendix N)*

- SR 1.27 We all know that computers are used to perform complex jobs. In this section, you learned that a computer’s instructions can do only simple tasks. Explain this apparent contradiction.
- SR 1.28 What is the relationship between a high-level language and a machine language?
- SR 1.29 What is Java bytecode?
- SR 1.30 Select the word from the following list that best matches each of the following phrases:
- assembly, compiler, high-level, IDE, interpreter, Java, low-level, machine
- a. A program written in this type of language can run directly on a computer.
 - b. Generally, each language instruction in this type of language corresponds to an equivalent machine language instruction.
 - c. Most programmers write their programs using this type of language.
 - d. Java is an example of this type of language.
 - e. This type of program translates code in one language to code in another language.
 - f. This type of program interweaves the translation of code and the execution of the code.

- SR 1.31 What do we mean by the syntax and semantics of a programming language?
- SR 1.32 Categorize each of the following situations as a compile-time error, run-time error, or logical error.
- Misspelling a Java reserved word.
 - Calculating the average of an empty list of numbers by dividing the sum of the numbers on the list (which is zero) by the size of the list (which is also zero).
 - Outputting a student's high test grade when their average test grade should have been output.

1.6 Object-Oriented Programming

As we stated earlier in this chapter, Java is an object-oriented (OO) language. As the name implies, an *object* is a fundamental entity in a Java program. This book is focused on the idea of developing software by defining objects that interact with each other.

The principles of object-oriented software development have been around for many years, essentially as long as high-level programming languages have been used. The programming language Simula, developed in the 1960s, had many characteristics that define the modern OO approach to software development. In the 1980s and 1990s, object-oriented programming became wildly popular, due in large part to the development of programming languages like C++ and Java. It is now the dominant approach used in commercial software development.

One of the most attractive characteristics of the object-oriented approach is the fact that objects can be used quite effectively to represent real-world entities. We can use a software object to represent an employee in a company, for instance. We'd create one object per employee, each with behaviors and characteristics that we need to represent. In this way, object-oriented programming allows us to map our programs to the real situations that the programs represent. That is, the object-oriented approach makes it easier to solve problems, which is the point of writing a program in the first place.

Let's discuss the general issues related to problem solving, then explore the specific characteristics of the object-oriented approach that helps us solve those problems.

KEY CONCEPT

Object-oriented programming helps us solve problems, which is the purpose of writing a program.

Problem Solving

In general, problem solving consists of multiple steps:

1. Understanding the problem.
2. Designing a solution.
3. Considering alternatives to the solution and refining the solution.
4. Implementing the solution.
5. Testing the solution and fixing any problems that exist.

Although this approach applies to any kind of problem solving, it works particularly well when developing software. These steps aren't purely linear. That is, some of the activities will overlap others. But at some point, all of these steps should be carefully addressed.

The first step, understanding the problem, may sound obvious, but a lack of attention to this step has been the cause of many misguided software development efforts. If we attempt to solve a problem we don't completely understand, we often end up solving the wrong problem or at least going off on improper tangents. Each problem has a *problem domain*, the real-world issues that are key to our solution. For example, if we are going to write a program to score a bowling match, then the problem domain includes the rules of bowling. To develop a good solution, we must thoroughly understand the problem domain.

The key to designing a problem solution is breaking it down into manageable pieces. A solution to any problem can rarely be expressed as one big task. Instead, it is a series of small cooperating tasks that interact to perform a larger task. When developing software, we don't write one big program. We design separate pieces that are responsible for certain parts of the solution, then integrate them with the other parts.

Our first inclination toward a solution may not be the best one. We must always consider alternatives and refine the solution as necessary. The earlier we consider alternatives, the easier it is to modify our approach.

Implementing the solution is the act of taking the design and putting it in a usable form. When developing a software solution to a problem, the implementation stage is the process of actually writing the program. Too often programming is thought of as writing code. But in most cases, the act of designing the program should be far more interesting and creative than the process of implementing the design in a particular programming language.

At many points in the development process, we should test our solution to find any errors that exist so that we can fix them. Testing cannot guarantee that there aren't still problems yet to be discovered, but it can raise our confidence that we have a viable solution.

KEY CONCEPT

Program design involves breaking a solution down into manageable pieces.

Throughout this text we explore techniques that allow us to design and implement elegant programs. Although we will often get immersed in these details, we should never forget that our primary goal is to solve problems.

Object-Oriented Software Principles

Object-oriented programming ultimately requires a solid understanding of the following terms:

- object
- attribute
- method
- class
- encapsulation
- inheritance
- polymorphism

In addition to these terms, there are many associated concepts that allow us to tailor our solutions in innumerable ways. This book is designed to help you evolve your understanding of these concepts gradually and naturally. This section provides an overview of these ideas at a high level to establish some terminology and provide the big picture.

We mentioned earlier that an *object* is a fundamental element in a program. A software object often represents a real object in our problem domain, such as a bank account. Every object has a *state* and a set of *behaviors*. By “state” we mean state of being—fundamental characteristics that currently define the object. For example, part of a bank account’s state is its current balance. The behaviors of an object are the activities associated with the object. Behaviors associated with a bank account probably include the ability to make deposits and withdrawals.

In addition to objects, a Java program also manages primitive data. *Primitive data* includes fundamental values such as numbers and characters. Objects usually represent more interesting or complex entities.

An object’s *attributes* are the values it stores internally, which may be represented as primitive data or as other objects. For example, a bank account object may store a floating point number (a primitive value) that represents the balance of the account. It may contain other attributes, such as the name of the account owner. Collectively, the values of an object’s attributes define its current state.

As mentioned earlier in this chapter, a *method* is a group of programming statements that is given a name. When a method is

KEY CONCEPT

Each object has a state, defined by its attributes, and a set of behaviors, defined by its methods.

invoked, its statements are executed. A set of methods is associated with an object. The methods of an object define its potential behaviors. To define the ability to make a deposit into a bank account, we define a method containing programming statements that will update the account balance accordingly.

An object is defined by a *class*. A class is the model or blueprint from which an object is created. Consider the blueprint created by an architect when designing a house. The blueprint defines the important characteristics of the house—its walls, windows, doors, electrical outlets, and so on. Once the blueprint is created, several houses can be built using it, as depicted in Figure 1.22.

In one sense, the houses built from the blueprint are different. They are in different locations, have different addresses, contain different furniture, and are inhabited by different people. Yet in many ways they are the “same” house. The layout of the rooms and other crucial characteristics are the same in each. To create a different house, we would need a different blueprint.

A class is a blueprint of an object. It establishes the kind of data an object of that type will hold and defines the methods that represent the behavior of such objects. However, a class is not an object any more than a blueprint is a house. In general, a class contains no space to store data. Each object has space for its own data, which is why each object can have its own state.

Once a class has been defined, multiple objects can be created from that class. For example, once we define a class to represent the concept of a bank account,

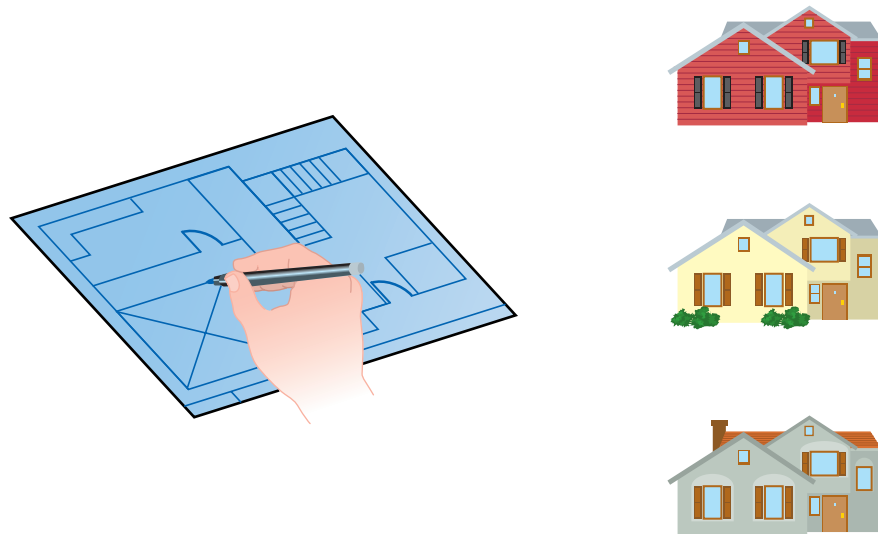


FIGURE 1.22 A house blueprint and three houses created from it

we can create multiple objects that represent specific, individual bank accounts. Each bank account object would keep track of its own balance.

An object should be *encapsulated*, which means it protects and manages its own information. That is, an object should be self-governing. The only changes made to the state of the object should be accomplished by that object's methods. We should design objects so that other objects cannot "reach in" and change their states.

Classes can be created from other classes by using *inheritance*. That is, the definition of one class can be based on another class that already exists. Inheritance is a form of *software reuse*, capitalizing on the similarities between various kinds of classes that we may want to create. One class can be used to derive several new classes. Derived classes can then be used to derive even more classes. This creates a hierarchy of classes, where the attributes and methods defined in one class are inherited by its children, which in turn pass them on to their children, and so on. For example, we might create a hierarchy of classes that represent various types of accounts. Common characteristics are defined in high-level classes, and specific differences are defined in derived classes.

Polymorphism is the idea that we can refer to multiple types of related objects over time in consistent ways. It gives us the ability to design powerful and elegant solutions to problems that deal with multiple objects.

Some of the core object-oriented concepts are depicted in Figure 1.23. We don't expect you to understand these ideas fully at this point. Most of this book is designed to flesh out these ideas. This overview is intended only to set the stage.

KEY CONCEPT

A class is a blueprint of an object. Multiple objects can be created from one class definition.

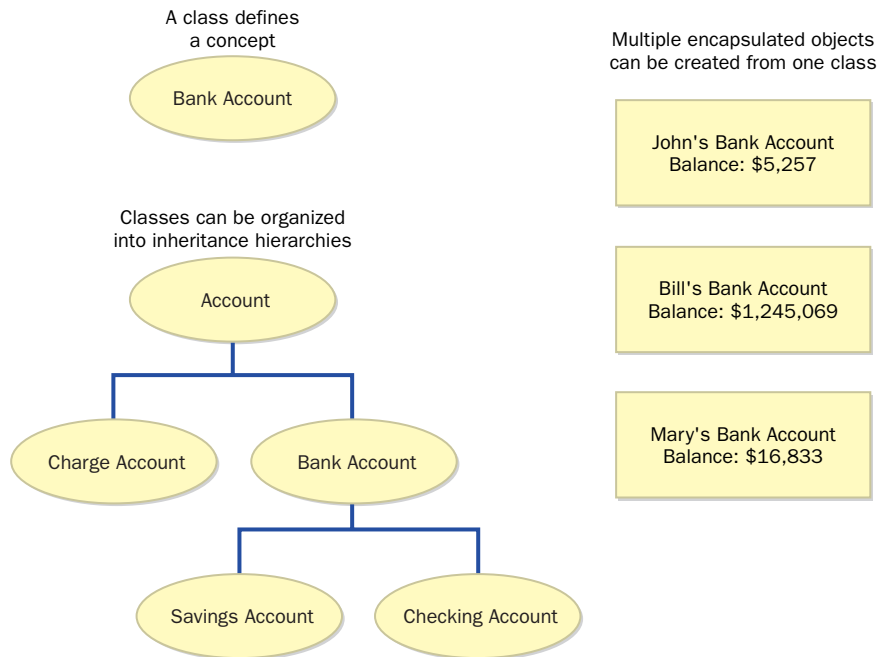


FIGURE 1.23 Various aspects of object-oriented software

SELF-REVIEW QUESTIONS (see answers in Appendix N)

- SR 1.33** List the five general steps required to solve a problem.
- SR 1.34** Why is it important to consider more than one approach to solving a problem? Why is it important to consider alternatives early in the process of solving a problem?
- SR 1.35** What are the primary concepts that support object-oriented programming?

Summary of Key Concepts

- A computer system consists of hardware and software that work in concert to help us solve problems.
- The CPU reads the program instructions from main memory, executing them one at a time until the program ends.
- The operating system provides a user interface and manages computer resources.
- As far as the user is concerned, the interface *is* the program.
- Digital computers store information by breaking it into pieces and representing each piece as a number.
- Binary is used to store information in a computer because the devices that store and manipulate binary data are inexpensive and reliable.
- There are exactly 2^N combinations of N bits. Therefore, N bits can represent up to 2^N unique items.
- The core of a computer is made up of main memory, which stores programs and data, and the CPU, which executes program instructions one at a time.
- An address is a unique number associated with each memory location.
- Main memory is volatile, meaning the stored information is maintained only as long as electric power is supplied.
- The surface of a CD has both smooth areas and small pits. A pit represents a binary 1 and a smooth area represents a binary 0.
- A rewritable CD simulates the pits and smooth areas of a regular CD by using a coating that can be made amorphous or crystalline as needed.
- The fetch-decode-execute cycle forms the foundation of computer processing.
- A network consists of two or more computers connected together so that they can exchange information.
- Sharing a communication line creates delays, but it is cost effective and simplifies adding new computers to the network.
- A local-area network (LAN) is an effective way to share information and resources throughout an organization.
- The Internet is a wide-area network (WAN) that spans the globe.
- Every computer connected to the Internet has an IP address that uniquely identifies it.

- The World Wide Web is software that makes sharing information across a network easy.
- A URL uniquely specifies documents and other information found on the Web for a browser to obtain and display.
- This book focuses on the principles of object-oriented programming.
- Comments do not affect a program's processing; instead, they serve to facilitate human comprehension.
- Inline documentation should provide insight into your code. It should not be ambiguous or belabor the obvious.
- Java is case sensitive. The uppercase and lowercase versions of a letter are distinct.
- Identifier names should be descriptive and readable.
- Appropriate use of white space makes a program easier to read and understand.
- You should adhere to a set of guidelines that establish the way you format and document your programs.
- All programs must be translated to a particular CPU's machine language in order to be executed.
- High-level languages allow a programmer to ignore the underlying details of machine language.
- A Java compiler translates Java source code into Java bytecode, a low-level, architecture-neutral representation of the program.
- Many different development environments exist to help you create and modify Java programs.
- Syntax rules dictate the form of a program. Semantics dictate the meaning of the program statements.
- The programmer is responsible for the accuracy and reliability of a program.
- A Java program must be syntactically correct or the compiler will not produce bytecode.
- Object-oriented programming helps us solve problems, which is the purpose of writing a program.
- Program design involves breaking a solution down into manageable pieces.
- Each object has a state, defined by its attributes, and a set of behaviors, defined by its methods.
- A class is a blueprint of an object. Multiple objects can be created from one class definition.

Exercises

- EX 1.1 Describe the hardware components of your personal computer or of a computer in a lab to which you have access. Include the processor type and speed, storage capacities of main and secondary memory, and types of I/O devices. Explain how you determined your answers.
- EX 1.2 Why do we use the binary number system to store information on a computer?
- EX 1.3 How many unique items can be represented with each of the following?
- 1 bit
 - 3 bits
 - 6 bits
 - 8 bits
 - 10 bits
 - 16 bits
- EX 1.4 If a picture is made up of 128 possible colors, how many bits would be needed to store each pixel of the picture? Why?
- EX 1.5 If a language uses 240 unique letters and symbols, how many bits would be needed to store each character of a document? Why?
- EX 1.6 How many bits are there in each of the following? How many bytes are there in each?
- 12 KB
 - 5 MB
 - 3 GB
 - 2 TB
- EX 1.7 Explain the difference between random access memory (RAM) and read-only memory (ROM).
- EX 1.8 A disk is a random-access device but it is not RAM (random access memory). Explain.
- EX 1.9 Determine how your computer, or a computer in a lab to which you have access, is connected to others across a network. Is it linked to the Internet? Draw a diagram to show the basic connections in your environment.
- EX 1.10 Explain the differences between a local-area network (LAN) and a wide-area network (WAN). What is the relationship between them?

- EX 1.11 What is the total number of communication lines needed for a fully connected point-to-point network of eight computers? Nine computers? Ten computers? What is a general formula for determining this result?
- EX 1.12 Explain the difference between the Internet and the World Wide Web.
- EX 1.13 List and explain the parts of the URLs for:
- your school
 - the Computer Science department of your school
 - your instructor's Web page
- EX 1.14 Use a Web browser to access information through the Web about the following topics. For each one, explain the process you used to find the information and record the specific URLs you found.
- the Philadelphia Phillies baseball team
 - wine production in California
 - the subway systems in two major cities
 - vacation opportunities in the Caribbean
- EX 1.15 Give examples of the two types of Java comments and explain the differences between them.
- EX 1.16 Which of the following are not valid Java identifiers? Why?
- `Factorial`
 - `anExtremelyLongIdentifierIfYouAskMe`
 - `2ndLevel`
 - `level2`
 - `MAX_SIZE`
 - `highest$`
 - `hook&ladder`
- EX 1.17 Why are the following valid Java identifiers not considered good identifiers?
- `q`
 - `totVal`
 - `theNextValueInTheList`
- EX 1.18 Java is case sensitive. What does that mean?
- EX 1.19 What do we mean when we say that the English language is ambiguous? Give two examples of English ambiguity (other

than the example used in this chapter) and explain the ambiguity. Why is ambiguity a problem for programming languages?

- EX 1.20 Categorize each of the following situations as a compile-time error, run-time error, or logical error.
- multiplying two numbers when you meant to add them
 - dividing by zero
 - forgetting a semicolon at the end of a programming statement
 - spelling a word wrong in the output
 - producing inaccurate results
 - typing a { when you should have typed (

Programming Projects

- PP 1.1 Enter, compile, and run the following application:

```
public class Test
{
    public static void main (String[] args)
    {
        System.out.println ("An Emergency Broadcast");
    }
}
```

- PP 1.2 Introduce the following errors, one at a time, to the program from PP 1.1. Record any error messages that the compiler produces. Fix the previous error each time before you introduce a new one. If no error messages are produced, explain why. Try to predict what will happen before you make each change.

- change `Test` to `test`
- change `Emergency` to `emergency`
- remove the first quotation mark in the string
- remove the last quotation mark in the string
- change `main` to `man`
- change `println` to `bogus`
- remove the semicolon at the end of the `println` statement
- remove the last brace in the program

- PP 1.3 Write an application that prints, on separate lines, your name, your birthday, your hobbies, your favorite book, and your favorite movie. Label each piece of information in the output.



Video Note

Developing a solution of PP 1.2.

- PP 1.4 Write an application that prints the phrase `Knowledge is Power`:
- on one line
 - on three lines, one word per line, with the words centered relative to each other
 - inside a box made up of the characters `=` and `|`

PP 1.5 Write an application that prints a list of four or five web sites that you enjoy. Print both the site name and the URL.

PP 1.6 Write an application that prints the first few verses of a song (your choice). Label the chorus.

PP 1.7 Write an application that prints the following diamond shape. Don't print any unneeded characters. (That is, don't make any character string longer than it has to be.)

```

*
***
*****
*****
*****
*****
*****
***
*

```

PP 1.8 Write an application that displays your initials in large block letters. Make each large letter out of the corresponding regular character. For example:

```

JJJJJJJJJJJJJJ  AAAAAAAAAA  LLLL
JJJJJJJJJJJJJJ  AAAAAAAAAAAA LLLL
                JJJJ   AAA   AAA  LLLL
                JJJJ   AAA   AAA  LLLL
                JJJJ   AAAAAAAAAA LLLL
J                JJJJ   AAAAAAAAAA LLLL
JJ               JJJJ   AAA   AAA  LLLL
                JJJJJJJJJJ  AAA   AAA  LLLLLLLLLLLLLLLL
                JJJJJJJJJ  AAA   AAA  LLLLLLLLLLLLLLLL

```

