Chapter 4

IT Infrastructure: Hardware and Software

LEARNING TRACK 1: HOW COMPUTER HARDWARE AND SOFTWARE WORK

Although managers and business professionals do not need to be computer technology experts, they should have a basic understanding of the role of hardware and software in the organization's information technology (IT) infrastructure so that they can make technology decisions that promote organizational performance and productivity. This chapter surveys the capabilities of computer hardware and computer software and highlights the major issues in the management of the firm's hardware and software assets.

Computer Hardware and Information Technology Infrastructure

Computer hardware provides the underlying physical foundation for the firm's IT infrastructure. Other infrastructure components—software, data, and networks—require computer hardware for their storage or operation.

The Computer System

A contemporary computer system consists of a central processing unit, primary storage, secondary storage, input devices, output devices, and communications devices (see Figure 4-1). The central processing unit manipulates raw data into a more useful form and controls the other parts of the computer system. Primary storage temporarily stores data and program instructions during processing, whereas secondary storage

FIGURE 4-1   Hardware components of a computer system.
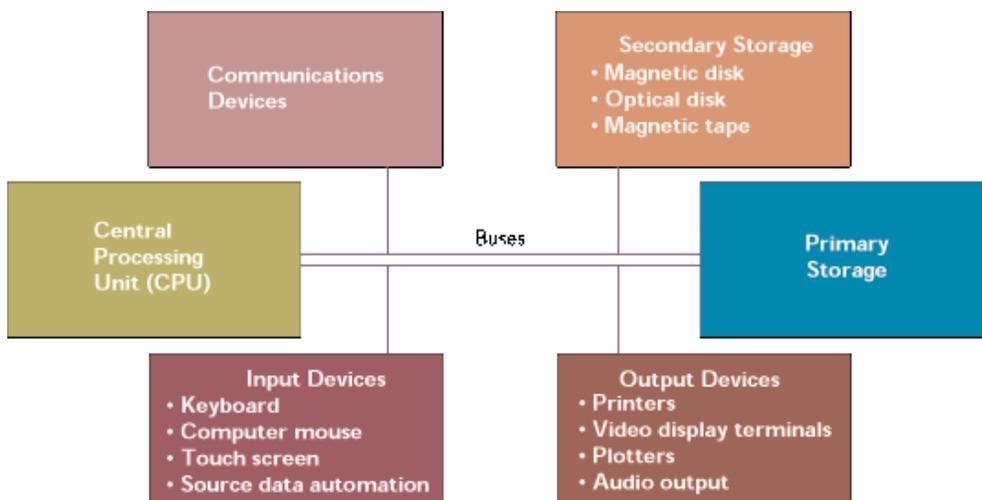
A contemporary computer system can be categorized into six major components. The central processing unit manipulates data and controls the other parts of the computer system; primary storage temporarily stores data and program instructions during processing; secondary storage stores data and instructions when they are not used in processing; input devices convert data and instructions for processing in the computer; output devices present data in a form that people can understand; and communications devices control the passing of information to and from communications networks.

FIGURE 4-2   Bits and bytes.

Bits are represented by either a 0 or 1. A string of eight bits constitutes a byte, which represents a character or number. Illustrated here is a byte representing the letter "A" using the ASCII binary coding standard.
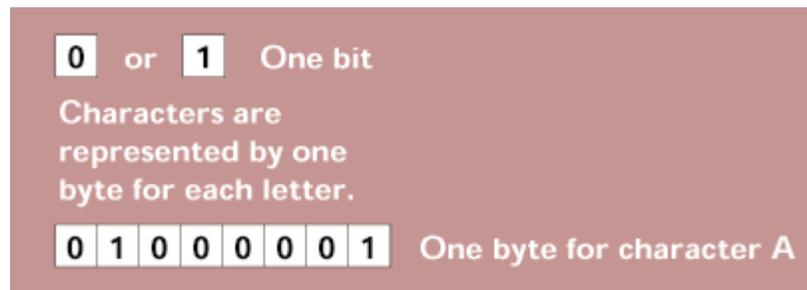
devices (magnetic and optical disks, magnetic tape) store data and programs when they are not being used in processing. Input devices, such as a keyboard or mouse, convert data and instructions into electronic form for input into the computer. Output devices, such as printers and video display terminals, convert electronic data produced by the computer system and display them in a form that people can understand. Communications devices provide connections between the computer and communications networks. Buses are circuitry paths for transmitting data and signals among the parts of the computer system.

In order for information to flow through a computer system and be in a form suitable for processing, all symbols, pictures, or words must be reduced to a string of binary digits. A binary digit is called a bit and represents either a 0 or a 1. In the computer, the presence of an electronic or magnetic signal means one, and its absence signifies zero. Digital computers operate directly with binary digits, either singly or strung together to form bytes. A string of eight bits that the computer stores as a unit is called a byte. Each byte can be used to store a decimal number, a symbol, a character, or part of a picture (see Figure 4-2).

Computers can represent pictures by creating a grid overlay of the picture. Each single point in this grid or matrix is called a pixel (picture element) and consists of a number of bits. The computer then stores this information on each pixel.

The CPU and Primary Storage

The central processing unit (CPU) is the part of the computer system where the manipulation of symbols, numbers, and letters occurs, and it controls the other parts of the computer system (see Figure 4-3). Located near the CPU is primary storage (sometimes called primary memory or main memory), where data and program instructions are stored temporarily during processing. Buses provide pathways for transmitting data and signals between the CPU, primary storage, and the other devices in the computer system. The characteristics of the CPU and primary storage are very important in determining a computer's speed and capabilities.

THE ARITHMETIC-LOGIC UNIT AND CONTROL UNIT

Figure 4-3 also shows that the CPU consists of an arithmetic-logic unit and a control unit. The arithmetic-logic unit (ALU) performs the computer's principal logical and arithmetic operations. It adds, subtracts, multiplies, and divides, determining whether a number is positive, negative, or zero. In addition to performing arithmetic functions, an ALU must be able to determine when one quantity is greater than or less than another and when two quantities are equal. The ALU can perform logic operations on letters as well as numbers.

FIGURE 4-3   The CPU and primary storage.
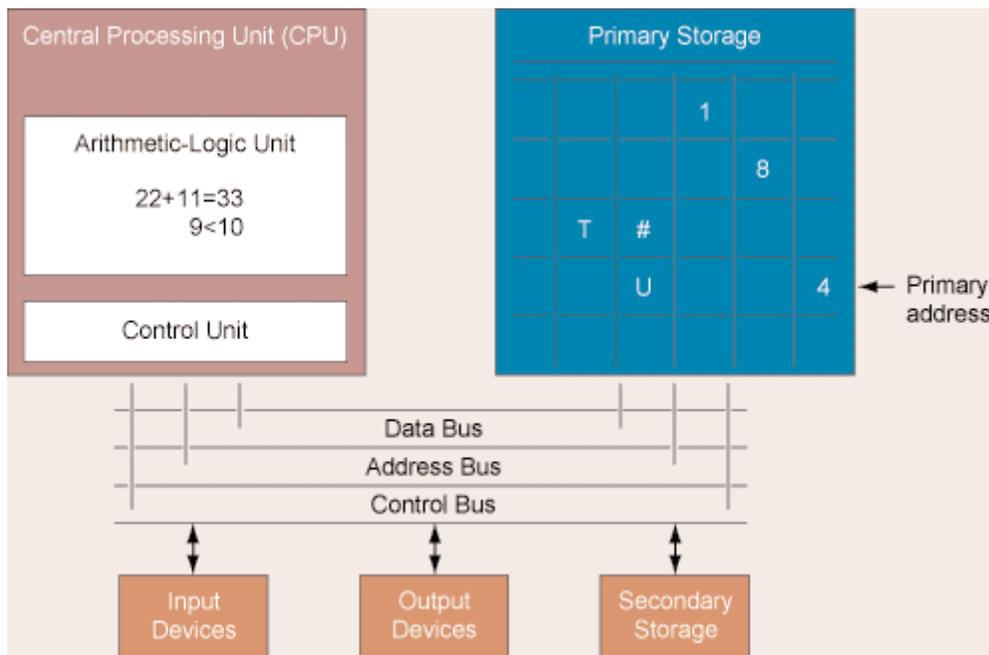
The CPU contains an arithmetic-logic unit and a control unit. Data and instructions are stored in unique addresses in primary storage that the CPU can access during processing. The data bus, address bus, and control bus transmit signals between the central processing unit, primary storage, and other devices in the computer system.

The control unit coordinates and controls the other parts of the computer system. It reads a stored program, one instruction at a time, and directs other components of the computer system to perform the program's required tasks. The series of operations required to process a single machine instruction is called the machine cycle. Older computers and PCs have machine cycle times measured in microseconds (millionths of a second). More powerful machines have machine cycle times measured in nanoseconds (billionths of a second) or picoseconds (trillionths of a second). Another measure of machine cycle time is by MIPS, or millions of instructions per second.

PRIMARY STORAGE

Primary storage has three functions. It stores all or part of the software program that is being executed. Primary storage also stores the operating system programs that manage the operation of the computer. Finally, the primary storage area holds data that the program is using. Internal primary storage is often called RAM, or random access memory. It is called RAM because it can directly access any randomly chosen location in the same amount of time.

Primary memory is divided into storage locations called bytes. Each location contains a set of eight binary switches or devices, each of which can store one bit of information. The set of eight bits found in each storage location is sufficient to store one letter, one digit, or one special symbol (such as $). Each byte has a unique address, similar to a mailbox, indicating where it is located in RAM. The computer can remember where the data in all of the bytes are located simply by keeping track of these addresses.

Computer storage capacity is measured in bytes. Table 4-1 lists computer storage capacity measurements. One thousand bytes (actually 1,024 storage positions) is called a kilobyte. One million bytes is called a megabyte, one billion bytes is called a gigabyte, and one trillion bytes is called a terabyte.

Primary storage is composed of semiconductors, which are integrated circuits made by printing thousands and even millions of tiny transistors on small silicon chips.

TABLE 4-1   Computer storage capacity

Byte	String of eight bits
Kilobyte	1,000 bytes*
Megabyte	1,000,000 bytes
Gigabyte	1,000,000,000 bytes
Terabyte	1,000,000,000,000 bytes

Actually 1024 storage positions.

There are several different kinds of semiconductor memory used in primary storage. RAM is used for short-term storage of data or program instructions. RAM is volatile: Its contents will be lost when the computer's electric supply is disrupted by a power outage or when the computer is turned off. ROM, or read-only memory, can only be read from; it cannot be written to. ROM chips come from the manufacturer with programs already burned in, or stored. ROM is used in general-purpose computers to store important or frequently used programs.

Computer Processing

The processing capability of the CPU plays a large role in determining the amount of work that a computer system can accomplish.

MICROPROCESSORS AND PROCESSING POWER

Contemporary CPUs use semiconductor chips called microprocessors, which integrate all of the memory, logic, and control circuits for an entire CPU onto a single chip. The speed and performance of a computer's microprocessors help determine a computer's processing power and are based on the microprocessor's word length, cycle speed, and data bus width. Word length refers to the number of bits that the computer can process at one time. A 32-bit chip can process 32 bits, or 4 bytes, of data in a single machine cycle. A 64-bit chip can process 64 bits or 8 bytes in a single cycle. The larger the word length, the greater the computer's speed.

A second factor affecting chip speed is cycle speed. Every event in a computer must be sequenced so that one step logically follows another. The control unit sets a beat to the chip. This beat is established by an internal clock and is measured in megahertz (abbreviated MHz, which stands for millions of cycles per second). The Intel 8088 chip, for instance, originally had a clock speed of 4.47 megahertz, whereas the Intel Pentium III chip has a clock speed that ranges from 450 to more than 900 megahertz and the Pentium 4 chip clock speed can reach up to 2 gigahertz.

A third factor affecting speed is the data bus width. The data bus acts as a highway between the CPU, primary storage, and other devices, determining how much data can be moved at one time. The 8088 chip used in the original IBM personal computer, for example, had a 16-bit word length but only an 8-bit data bus width. This meant that data were processed within the CPU chip itself in 16-bit chunks but could only be moved 8 bits at a time between the CPU, primary storage, and external devices. On the other hand, Intel's Itanium chip has both a 64-bit word length and a 64-bit data bus width. To have a computer execute more instructions per second and work through programs or handle users expeditiously, it is necessary to increase the processor's word length, the data bus width, or the cycle speed—or all three.

Microprocessors can be made faster by using reduced instruction set computing (RISC) in their design. Conventional chips, based on complex instruction set computing, have several hundred or more instructions hard-wired into their circuitry, and they

FIGURE 4-4   Sequential and parallel processing.

SEQUENTIAL PROCESSING

PARALLEL PROCESSING
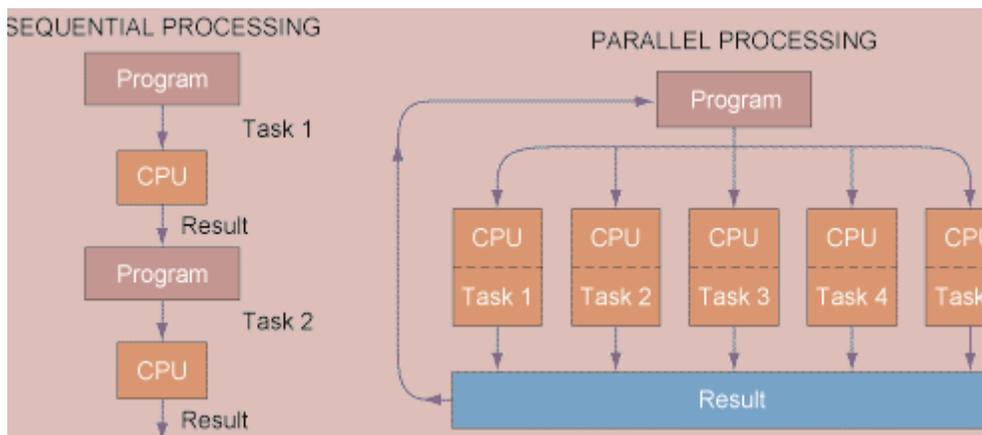
During sequential processing, each task is assigned to one CPU that processes one instruction at a time. In parallel processing, multiple tasks are assigned to multiple processing units to expedite the result.

may take several clock cycles to execute a single instruction. If the little-used instructions are eliminated, the remaining instructions can execute much faster. RISC processors have only the most frequently used instructions embedded in them. A RISC CPU can execute most instructions in a single machine cycle and sometimes multiple instructions at the same time. RISC is often used in scientific and workstation computing.

PARALLEL PROCESSING

Processing can also be sped up by linking several processors to work simultaneously on the same task. Figure 4-4 compares parallel processing to serial processing used in conventional computers. In parallel processing, multiple processing units (CPUs) break down a problem into smaller parts and work on it simultaneously. Getting a group of processors to attack the same problem at once requires both rethinking the problems and special software that can divide problems among different processors in the most efficient way possible, providing the needed data, and reassembling the many subtasks to reach an appropriate solution.

Massively parallel computers have huge networks of hundreds or even thousands of processor chips interwoven in complex and flexible ways to attack large computing problems. As opposed to parallel processing, where small numbers of powerful but expensive specialized chips are linked together, massively parallel machines link hundreds or even thousands of inexpensive, commonly used chips to break problems into many small pieces and solve them.

Storage, Input and Output Technology

The capabilities of computer systems depend not only on the speed and capacity of the CPU but also on the speed, capacity, and design of storage, input and output technology. Storage, input and output devices are called peripheral devices because they are outside the main computer system unit.

Secondary Storage Technology

The rise of digital firms has made storage a strategic technology. Although electronic commerce and electronic business are reducing manual processes, data of all types must be stored electronically and available whenever needed. Most of the information

Secondary storage devices such as floppy disks, optical disks, and hard disks are used to store large quantities of data outside the CPU and primary storage. They provide direct access to data for easy retrieval.
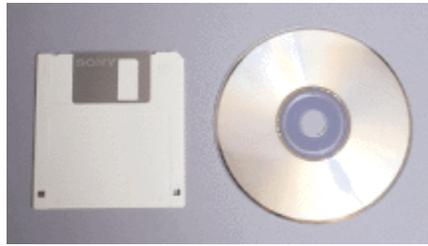
used by a computer application is stored on secondary storage devices located outside of the primary storage area. Secondary storage is used for relatively long term storage of data outside the CPU. Secondary storage is nonvolatile and retains data even when the computer is turned off. The most important secondary storage technologies are magnetic disk, optical disk, and magnetic tape.

MAGNETIC DISK

The most widely used secondary storage medium today is magnetic disk. There are two kinds of magnetic disks: floppy disks (used in PCs) and hard disks (used on large commercial disk drives and PCs). Large mainframe or midrange computer systems have multiple hard disk drives because they require immense disk storage capacity in the gigabyte and terabyte range. PCs also use floppy disks, which are removable and portable, with storage of up to 2.8 megabytes and a much slower access rate than hard disks. Removable disk drives are popular backup storage alternatives for PC systems. Magnetic disks on both large and small computers permit direct access to individual records so that data stored on the disk can be directly accessed regardless of the order in which the data were originally recorded. Disk technology is useful for systems requiring rapid and direct access to data.

Disk drive performance can be further enhanced by using a disk technology called RAID (Redundant Array of Inexpensive Disks). RAID devices package more than a hundred disk drives, a controller chip, and specialized software into a single large unit. Traditional disk drives deliver data from the disk drive along a single path, but RAID delivers data over multiple paths simultaneously, improving disk access time and relia-bility. For most RAID systems, data on a failed disk can be restored automatically without the computer system having to be shut down.

OPTICAL DISKS

Optical disks, also called compact disks or laser optical disks, user laser technology to store massive quantities of data in a highly compact form. They are available for both PCs and large computers. The most common optical disk system used with PCs is called CD-ROM (compact disk read-only memory). A 4.75-inch compact disk for PCs can store up to 660 megabytes, nearly 300 times more than a high-density floppy disk. Optical disks are most appropriate for applications where enormous quantities of unchanging data must be stored compactly for easy retrieval or for applications com-bining text, sound, and images.

CD-ROM is read-only storage. No new data can be written to it; it can only be read. WORM (write once/read many) and CD-R (compact disk-recordable) optical disk sys-tems allow users to record data only once on an optical disk. Once written, the data cannot be erased but can be read indefinitely. CD-RW (CD-ReWritable) technology has been developed to allow users to create rewritable optical disks for applications requiring large volumes of storage where the information is only occasionally updated.
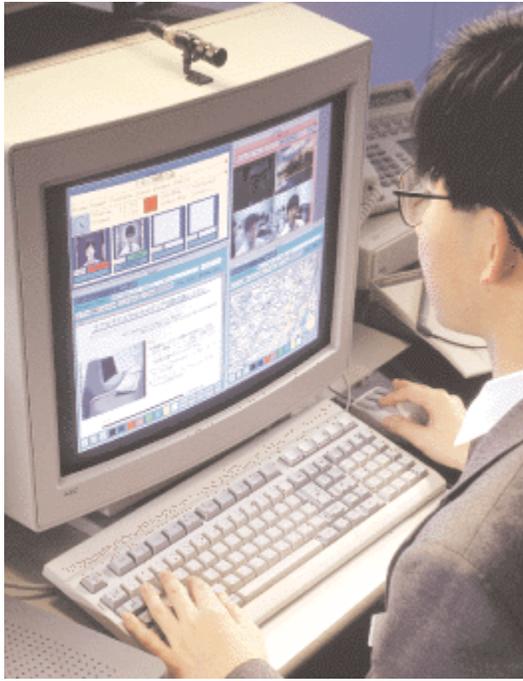
Multimedia combines text, graphics, sound, and video into a computer-based experience that permits two-way communication. Many organizations use this technology for interactive training.

Digital video disks (DVDs), also called digital versatile disks, are optical disks the same size as CD-ROMs but of even higher capacity. They can hold a minimum of 4.7 gigabytes of data, enough to store a full-length, high-quality motion picture. DVDs are initially being used to store movies and multimedia applications using large amounts of video and graphics, but they may replace CD-ROMs because they can store large amounts of digitized text, graphics, audio, and video data. Once read-only, writable and re-writable DVD drives and media are now available.

MAGNETIC TAPE

Magnetic tape is an older storage technology that still is employed for secondary storage of large quantities of data that are needed rapidly but not instantly. Magnetic tape is very inexpensive and relatively stable. However, it stores data sequentially and is relatively slow compared to the speed of other secondary storage media. In order to find an individual record stored on magnetic tape, such as an employment record, the tape must be read from the beginning up to the location of the desired record.

NEW STORAGE ALTERNATIVES: STORAGE AREA NETWORKS (SANS)

To meet the escalating demand for data-intensive multimedia, Web, and other services, the amount of data that companies need to store is increasing from 75 to 150 percent every year. Companies are turning to new kinds of storage infrastructures to deal with their mushrooming storage requirements and their difficulties managing large volumes of data.

Storage area networks (SANs) can provide a solution for companies with the need to share information across applications and computing platforms. A storage area network (SAN) is a high-speed network dedicated to storage that connects different kinds of storage devices, such as tape libraries and disk arrays. The network moves data among pools of servers and storage devices, creating an enterprise-wide infrastructure for data storage. The SAN creates a large central pool of storage that can be shared by multiple servers so that users can rapidly share data across the SAN. Every user in a company can access data from any server in the organization. Figure 4-5 illustrates

FIGURE 4-5   A storage area network (SAN).

The SAN stores data on many different types of storage devices, providing data to the enterprise. The SAN supports communication between any server and the storage unit as well as between different storage devices in the network.
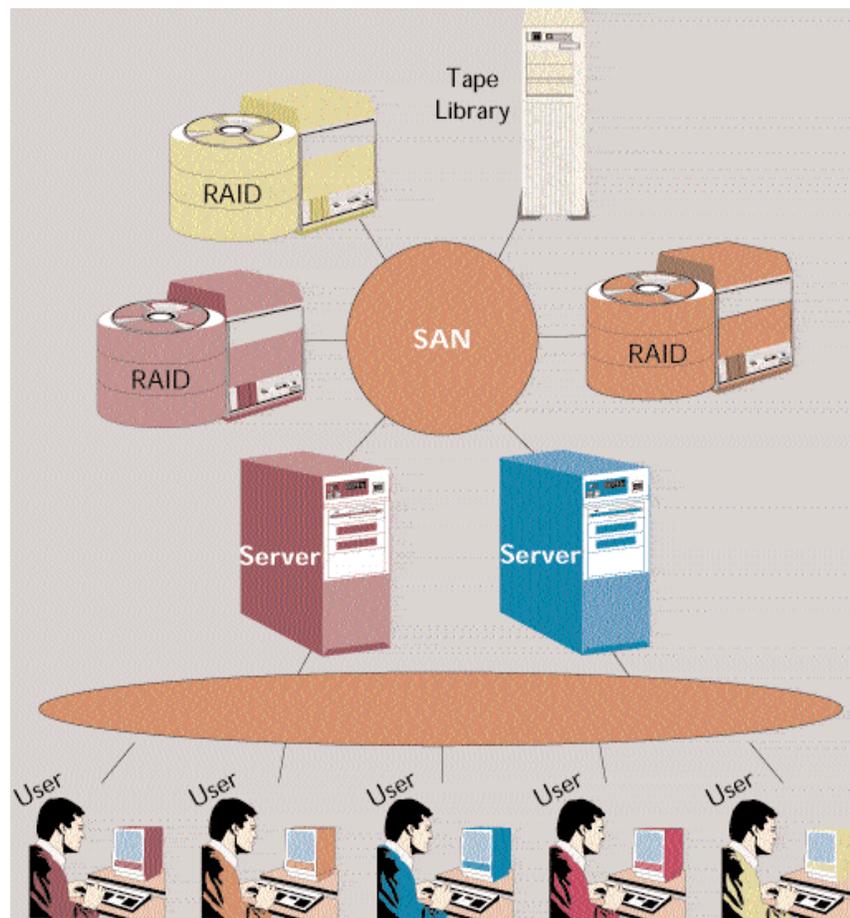
how a SAN works. The SAN storage devices are located on their own network and connected using a high-transmission technology such as Fibre Channel. SANs can be expensive and difficult to manage, but they are very useful for companies that can benefit from consolidating their storage resources and providing rapid data access to widely distributed users.

Input and Output Devices

Human beings interact with computer systems largely through input and output devices. Input devices gather data and convert them into electronic form for use by the computer, whereas output devices display data after they have been processed. Table 4-2 describes the principal input devices, and Table 4-3 describes the major output devices.

The principal input devices consist of keyboards, pointing devices (such as the computer mouse and touch screens), and source data automation technologies (optical and magnetic ink character recognition, pen-based input, digital scanners, audio input, and sensors), which capture data in computer-readable form at the time and place they are created.

The principal output devices are cathode ray tube terminals (CRTs), sometimes called video display terminals or VDTs, and printers.

TABLE 4-2   Principal Input Devices

Input Device	Description
Keyboard	Principal method of data entry for text and numerical data.
Computer mouse	Handheld device with point-and-click capabilities that is usually connected to the computer by a cable. The computer user can move the mouse around on a desktop to control the cursor's position on a computer display screen, pushing a button to select a command. Trackballs and touch pads often are used in place of the mouse as pointing devices on laptop PCs.
Touch screen	Allows user to enter limited amounts of data by touching the surface of a sensitized video display monitor with a finger or a pointer. Often found in information kiosks in retail stores, restaurants, and shopping malls.
Optical character recognition	Devices that can translate specially designed marks, characters, and codes into digital form. The most widely used optical code is the *bar code*, which is used in point-of-sale systems in supermarkets and retail stores. The codes can include time, date, and location data in addition to identification data.
Magnetic ink character recognition (MICR)	Used primarily in check processing for the banking industry. Characters on the bottom of a check identify the bank, checking account, and check number and are preprinted using a special magnetic ink. A MICR reader translates these characters into digital form for the computer.
Pen-based input	Handwriting-recognition devices such as pen-based tablets, notebooks, and notepads convert the motion made by an electronic stylus pressing on a touch-sensitive screen into digital format.
Digital scanner	Translates images such as pictures or documents into digital form and are an essential component of image-processing systems.
Audio input	Voice input devices that convert spoken words into digital form for processing by the computer. Microphones and tape cassette players can serve as input devices for music and other sounds.
Sensors	Devices that collect data directly from the environment for input into a computer system. For instance, today's farmers can use sensors to monitor the moisture of the soil in their fields to help them with irrigation.

TABLE 4-3   Principal Output Devices

| Output Device	Description
Cathode ray tube (CRT)	Electronic gun that shoots a beam of electrons illuminating tiny points on a display screen. Laptop computers use flat panel displays, which are less bulky than CRT monitors.
Printers	Devices that produce a printed hard copy of information ouput. They include impact printers (such as dot matrix printers) and nonimpact printers (such as laser, inkjet, and thermal transfer printers).
Audio output	Voice output devices that convert digital output data back into intelligible speech. Other audio output, such as music, can be delivered by speakers connected to the computer.

Batch and On-Line Input and Processing

The manner in which data are input into the computer affects how the data can be processed. Information systems collect and process information in one of two ways: through batch or through on-line processing. In batch processing, transactions, such as orders or payroll time cards, are accumulated and stored in a group or batch until the time when, because of some reporting cycle, it is efficient or necessary to process them. Batch processing is found primarily in older systems where users need only occasional reports. In on-line processing, the user enters transactions into a device (such as a data entry keyboard or bar code reader) that is directly connected to the computer system. The transactions usually are processed immediately. Most processing today is on-line processing.

Figure 4-6 compares batch and on-line processing. Batch systems often use tape as a storage medium, whereas on-line processing systems use disk storage, which permits immediate access to specific items. In batch systems, transactions are accumulated in a transaction file, which contains all the transactions for a particular time period. Periodically, this file is used to update a master file, which contains permanent information on entities. (An example is a payroll master file with employee earnings and deduction data. It is updated with weekly time-card transactions.) Adding the transaction data to the existing master file creates a new master file. In on-line processing,

FIGURE 4-6 A comparison of batch and on-line processing.

In batch processing, transactions are accumulated and stored in a group. Because batches are processed at regular intervals, such as daily, weekly, or monthly, information in the system will not always be up to date. In on-line processing, transactions are input immediately and usually processed immediately. Information in the system is generally up to date. A typical on-line application is an airline reservation system.
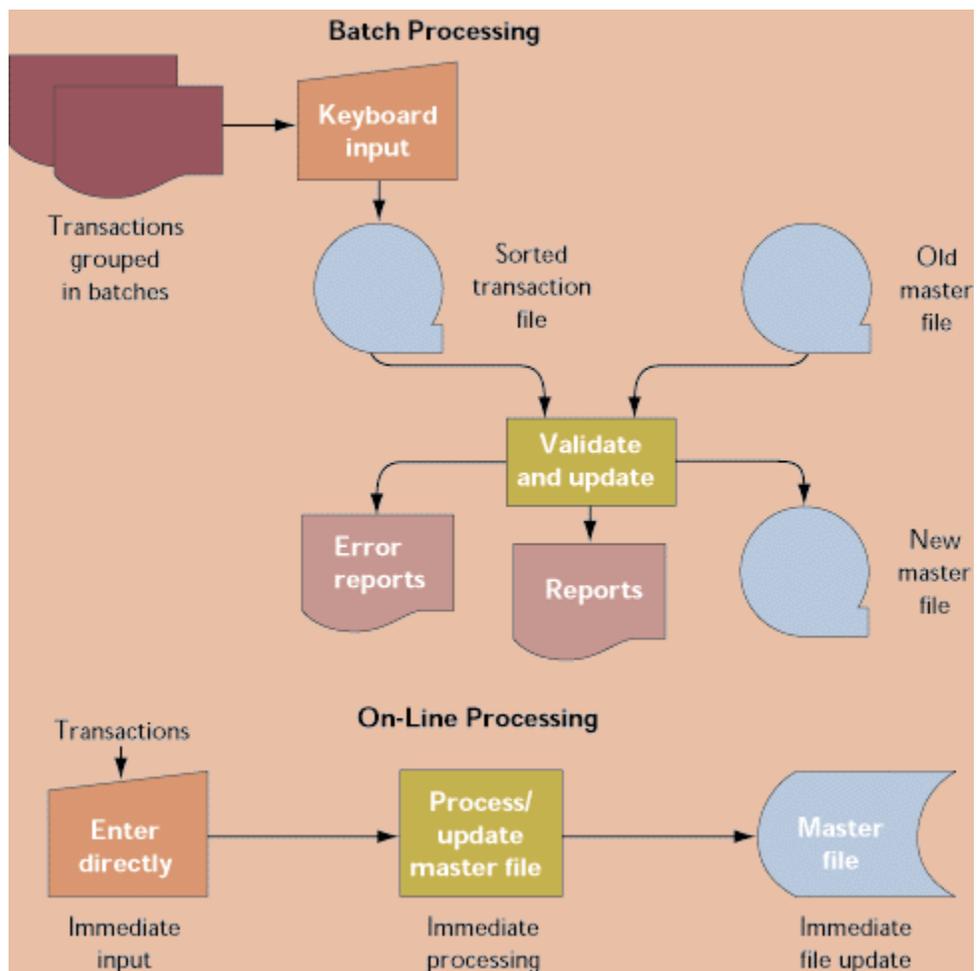
transactions are entered into the system immediately using a keyboard, pointing device, or source data automation, and the system usually responds immediately. The master file is updated continually.

Interactive Multimedia

The processing, input, output, and storage technologies we have just described can be used to create multimedia applications that integrate sound and full-motion video, or animation with graphics and text into a computer-based application. Multimedia is becoming the foundation of new consumer products and services, such as electronic books and newspapers, electronic classroom-presentation technologies, full-motion videoconferencing, imaging, graphics design tools, and video and voice mail. PCs today come with built-in multimedia capabilities, including high-resolution color monitors, CD-ROM drives or DVD drives to store video, audio, and graphic data, and stereo speakers for amplifying audio output.

Interactive Web pages replete with graphics, sound, animations, and full-motion video have made multimedia popular on the Internet. For example, visitors to the CNN Interactive Web site can access news stories from CNN, photos, on-air transcripts, video clips, and audio clips. The video and audio clips are made available using streaming technology, which allows audio and video data to be processed as a steady and continuous stream as they are downloaded from the Web.

Multimedia Web sites are also being used to sell digital products, such as digitized music clips. A compression standard known as MP3, also called MPEG3, which stands for Motion Picture Experts Group, audio layer 3, can compress audio files down to one-tenth or one-twelfth of their original size with virtually no loss in quality. Visitors to Web sites such as MP3.com can download MP3 music clips over the Internet and play them on their own computers.

System Software and PC Operating Systems

System software coordinates the various parts of the computer system and mediates between application software and computer hardware. The system software that manages and controls the computer's activities is called the operating system. Other system software consists of computer language translation programs that convert programming languages into machine language that can be understood by the computer and utility programs that perform common processing tasks.

FUNCTIONS OF THE OPERATING SYSTEM

The operating system is the computer system's chief manager. The operating system allocates and assigns system resources, schedules the use of computer resources and computer jobs, and monitors computer system activities. The operating system provides locations in primary memory for data and programs, and controls the input and output devices, such as printers, terminals, and telecommunication links. The operating system also coordinates the scheduling of work in various areas of the computer so that different parts of different jobs can be worked on at the same time. Finally, the operating system keeps track of each computer job and may also keep track of who is using the system, of what programs have been run, and of any unauthorized attempts to access the system.

MULTIPROGRAMMING

A series of specialized operating system capabilities enables the computer to handle many different tasks and users at the same time. Multiprogramming permits multiple programs to share a computer system's resources at any one time through concurrent

FIGURE 4-7   Single-program execution versus multiprogramming.

In multiprogramming, the computer can be used much more efficiently because a number of programs can be executing concurrently. Several complete programs are loaded into memory. The first program executes until an input/output event is read in the program. The CPU then moves to the second program until an input/output statement occurs. At this point, the CPU switches to the execution of the third program, and so forth, until eventually all three programs have been executed.
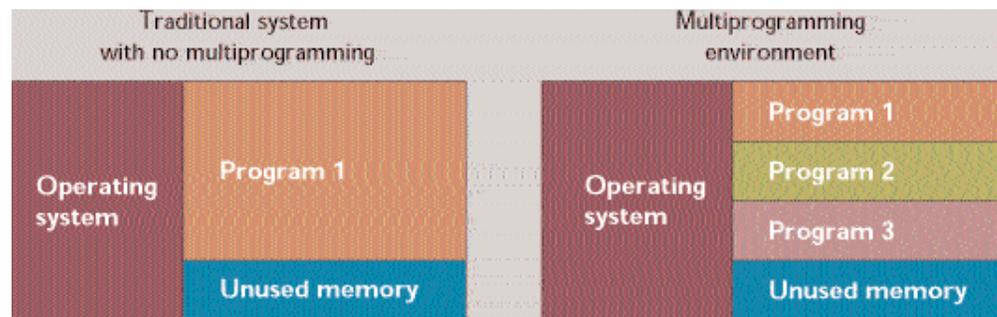

use of a CPU. Only one program is actually using the CPU at any given moment, but the input/output needs of other programs can be serviced at the same time. Two or more programs are active at the same time, but they do not use the same computer resources simultaneously. With multiprogramming, a group of programs takes turns using the processor. Figure 4-7 shows how three programs in a multiprogramming environment can be stored in primary storage. Multiprogramming on single-user operating systems such as those in older personal computers is called multitasking.

VIRTUAL STORAGE

Virtual storage handles programs more efficiently because the computer divides the programs into small fixed- or variable-length portions, storing only a small portion of the program in primary memory at one time. Only a few statements of a program actually execute at any given moment. This permits a very large number of programs to reside in primary memory, because only a tiny portion of each program is actually located there (see Figure 4-8), using this resource more efficiently. All other program pages are stored on a peripheral disk unit until they are ready for execution.

FIGURE 4-8   Virtual storage.

In virtual storage, programs are broken down into small sections that are read into memory only when needed. The rest of the program is stored on disk until it is required. In this way, very large programs can be executed by small machines, or a large number of programs can be executed concurrently by a single machine.
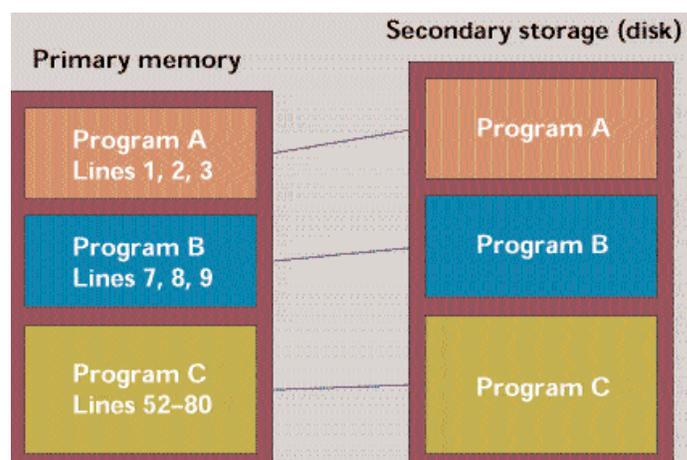
TIME SHARING

Time sharing is an operating system capability that allows many users to share computer processing resources simultaneously. It differs from multiprogramming in that the CPU spends a fixed amount of time on one program before moving on to another. Thousands of users are each allocated a tiny slice of computer time, when each is free to perform any required operations; at the end of this period, another user is given another tiny time slice of the CPU. This arrangement permits many users to be connected to a CPU simultaneously, with each receiving only a tiny amount of CPU time.

MULTIPROCESSING

Multiprocessing is an operating system capability that links together two or more CPUs to work in parallel in a single computer system. The operating system can assign multiple CPUs to execute different instructions from the same program or from different programs simultaneously, dividing the work between the CPUs. Whereas multiprogramming uses concurrent processing with one CPU, multiprocessing uses simultaneous processing with multiple CPUs.

LANGUAGE TRANSLATION AND UTILITY SOFTWARE

System software includes special language translator programs that translate high-level language programs written in programming languages such as COBOL, FORTRAN, or C into machine language that the computer can execute. The program in the high-level language before translation into machine language is called source code. A compiler translates source code into machine code called object code, which is linked to other object code modules and then executed by the computer. Some programming languages, such as BASIC, do not use a compiler but an interpreter, which translates each source code statement one at a time into machine code and executes it.

System software includes utility programs for routine, repetitive tasks, such as copying, clearing primary storage, computing a square root, or sorting. Utility programs can be shared by all users of a computer system and can be used in many different information system applications when requested.

PC OPERATING SYSTEMS AND GRAPHICAL USER INTERFACES

Like any other software, PC software is based on specific operating systems and computer hardware. Software written for one PC operating system generally cannot run on another.

When a user interacts with a computer, including a PC, the interaction is controlled by an operating system. A user communicates with an operating system through the user interface of that operating system. Contemporary PC operating systems use a graphical user interface, often called a GUI, which makes extensive use of icons, buttons, bars, and boxes to perform tasks. It has become the dominant model for the user interface of PC operating systems and for many types of application software.

Application Software and Programming Languages

Application software is primarily concerned with accomplishing the tasks of end users. Many different languages and software tools can be used to develop application software. Managers should understand which software tools and programming languages are appropriate for their organization's objectives.

PROGRAMMING LANGUAGES

The first generation of computer languages consisted of machine language, which required the programmer to write all program instructions in the 0s and 1s of binary

FIGURE 4-9   Assembly language.

This sample assembly language command adds the contents of register 3 to register 5 and stores the result in register 5. (A register is a temporary storage location in the CPU for small amounts of data or instructions).

AR 5, 3

code and to specify storage locations for every instruction and item of data used. Programming in machine language was a very slow, labor-intensive process. As computer hardware improved and processing speed and memory size increased, programming languages became progressively easier for humans to understand and use. From the mid-1950s to the mid-1970s, high-level programming languages emerged, allowing programs to be written with regular words using sentence-like statements. We now briefly describe the most important high-level languages.

ASSEMBLY LANGUAGE

Assembly language is the next level of programming language up from machine language and is considered a "second-generation" language. Like machine language, assembly language (Figure 4-9) is designed for a specific machine and specific microprocessors. Assembly language makes use of certain mnemonics (e.g., load, sum) to represent machine language instructions and storage locations. Although assembly language gives programmers great control, it is difficult and costly to write and learn. Assembly language is used primarily today in system software.

THIRD-GENERATION LANGUAGES: FORTRAN, COBOL, BASIC, PASCAL, AND C

Third-generation languages specify instructions as brief statements that are more like natural languages than assembly language. All are less efficient in the use of computer resources than earlier languages, they are easier to write and understand and have made it possible to create software for business and scientific problems. Important third-generation languages include FORTRAN, COBOL, BASIC, and C.

FORTRAN FORTRAN (FORmula TRANslator) (Figure 4-10) was developed in 1956 to provide an easy way of writing scientific and engineering applications. FORTRAN is especially useful in processing numeric data. Some business applications can be written in FORTRAN, and contemporary versions provide sophisticated structures for controlling program logic.

COBOL COBOL (COmmon Business Oriented Language) (Figure 4-11) was developed in the early 1960s by a committee representing both government and industry. Rear Admiral Grace M. Hopper was a key committee member who played a major role in COBOL development. COBOL was designed with business administration in mind, for processing large data files with alphanumeric characters (mixed alphabetic

FIGURE 4-10 FORTRAN.

This sample FORTRAN program code is part of a program to compute sales figures for a particular item.

READ (5,100) ID, QUANT, PRICE
TOTAL = QUANT * PRICE

FIGURE 4-11 COBOL.

MULTIPLY QUANT-SOLD BY UNIT-PRICE GIVING SALES-TOTAL.

This sample COBOL program code is part of a routine to compute total sales figures for a particular item.

and numeric data) and for performing repetitive tasks such as payroll. It is poor at complex, mathematical calculations. Also, there are many versions of COBOL, and not all are compatible with each other.

BASIC and Pascal BASIC and Pascal are used primarily in education to teach programming. BASIC (Beginners All-purpose Symbolic Instruction Code) was developed in 1964 by John Kemeny and Thomas Kurtz to teach students at Dartmouth College how to use computers. BASIC is easy to use but does few computer processing tasks well, even though it does them all. Different versions of BASIC exist.

Named after Blaise Pascal, the seventeenth-century mathematician and philosopher, Pascal was developed by the Swiss computer science professor Niklaus Wirth of Zurich in the late 1960s. Pascal is used primarily in computer science courses to teach sound programming practices.

C and C++ C is a powerful and efficient language developed at AT&T's Bell Labs in the early 1970s. It combines machine portability with tight control and efficient use of computer resources, and it can work on a variety of different computers. It is used primarily by professional programmers to create operating systems and application software, especially for PCs.

C++ is a newer version of C that is object-oriented. It has all the capabilities of C plus additional features for working with software objects. C++ is used for developing application software.