Book II

Inside the Box

The 5th Wave   By Rich Tennant

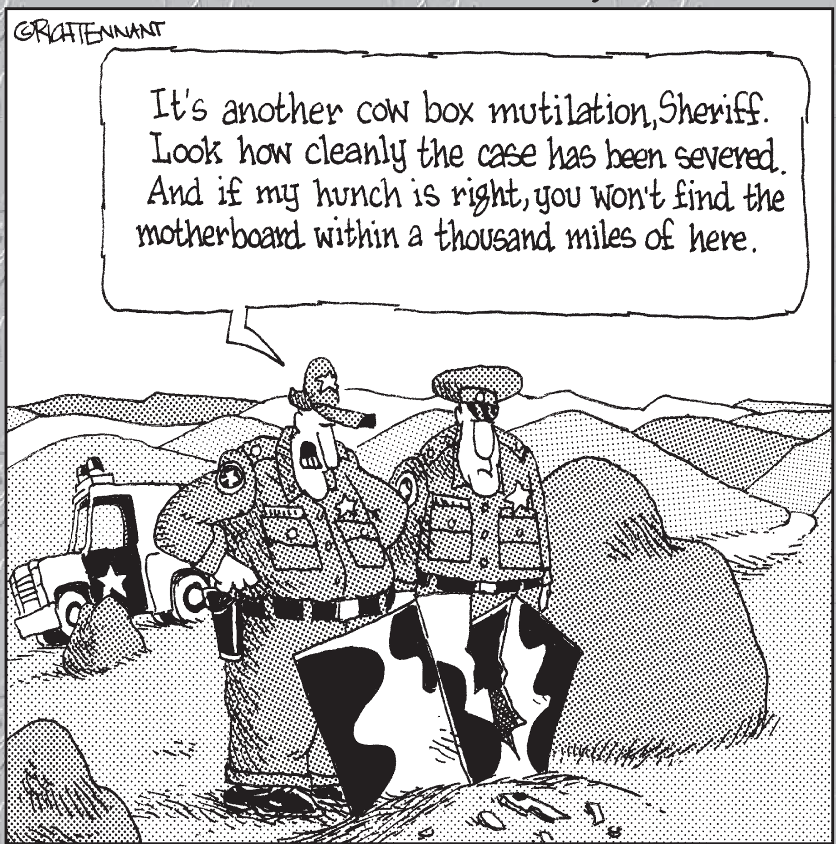

Contents at a Glance

Chapter 1: Knowing Your Motherboard

Exam Objectives

- **Distinguishing motherboard components**
- **Recognizing types of motherboards**
- **Identifying bus architectures**

One of the major replaceable components in your computer is the system board, also known as the motherboard. The *motherboard* is the big green board (that may not be a technical description, but I think that looking inside your system will demonstrate that it is an accurate one) connected to the computer case — it is the motherboard that holds your RAM, processor, and a number of other components in place.

This motherboard is the glue that connects all the other PC components together. For example, you can see how the hard drive is connected to the motherboard by following the IDE ribbon cable from your hard drive to the motherboard. If you do the same with the IDE cable that connects to the floppy drive, you can see that the floppy drive also connects to the motherboard. The memory sockets and the processor socket are likewise located on the motherboard.

All the components that work together to make the computer functional connect to the motherboard. If you take a close look at the motherboard, you can see wires embedded on the board that form little pathways that span the system. Think of these wires as the highway system that data signals use to travel from one location to another.

In this chapter, I introduce you to the different types of components found on the motherboard. After identifying the motherboard components, you find out about the different types of motherboards. Finally, you explore what an expansion bus is and discover the different bus architectures.

Finding Out What's on a Motherboard

When you look at the motherboard inside your computer, you notice that a number of different items connect to it. The memory sockets, the CPU socket, and the BIOS chip are all located on the motherboard. In this section, we identify the different components that are interconnected via the motherboard.

FOR THE EXAM

Remember that the terms *system board* and *motherboard* are interchangeable.

Processor

One of the easiest items to identify on the motherboard is the processor, also known as the *central processing unit* (CPU). The processor is usually the largest chip on the motherboard and is one of the few chips with a heat sink or fan on top of it, as shown in Figure 1-1.

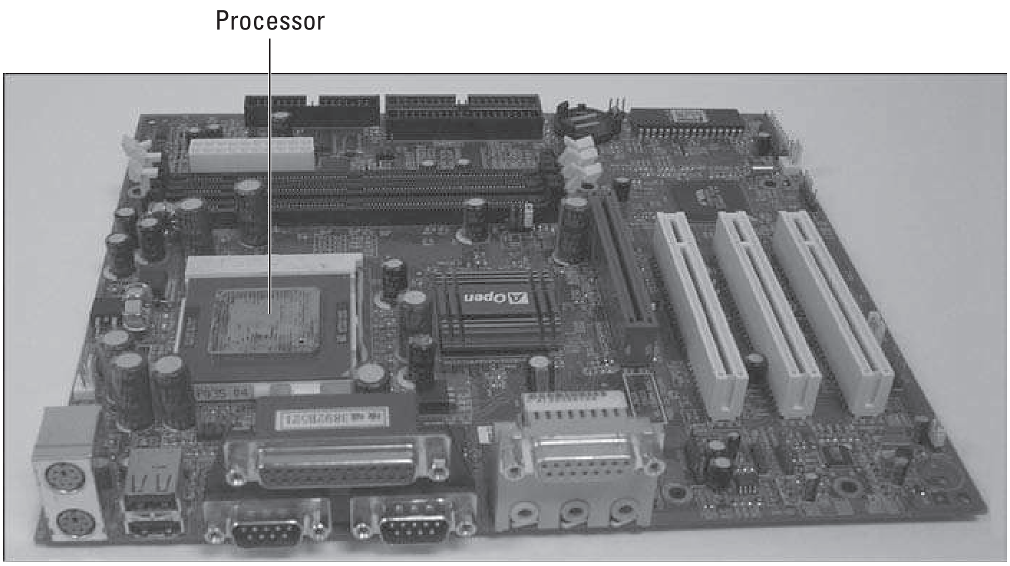

Figure 1-1: Identifying the processor on the motherboard.

The motherboard has a socket that the processor is inserted into. Today, this socket is implemented as a *zero insertion force (ZIF)* socket, which means that the processor chip can be removed or added to the socket with very little effort. ZIF sockets (shown in Figure 1-2) typically have a lever that you pull to pop the processor out of the socket.

When the Pentium II processor was developed, Intel used a different type of packaging, known as the *Single Edge Contact (SEC).* Motherboards had to implement a different type of "socket," known as *slot 1,* to hold this processor. The cartridge would drop into the slot, as shown in Figure 1-3. For more information on processors and sockets, check out Book II, Chapter 2.

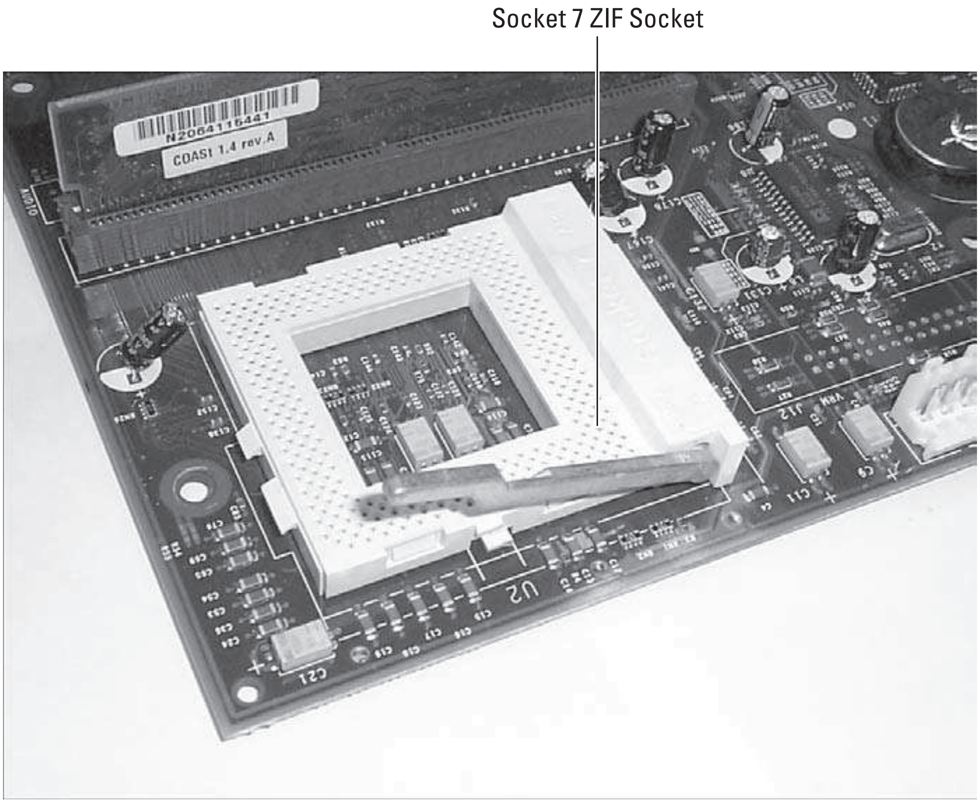

Figure 1-2: Looking at a ZIF socket located on the motherboard.

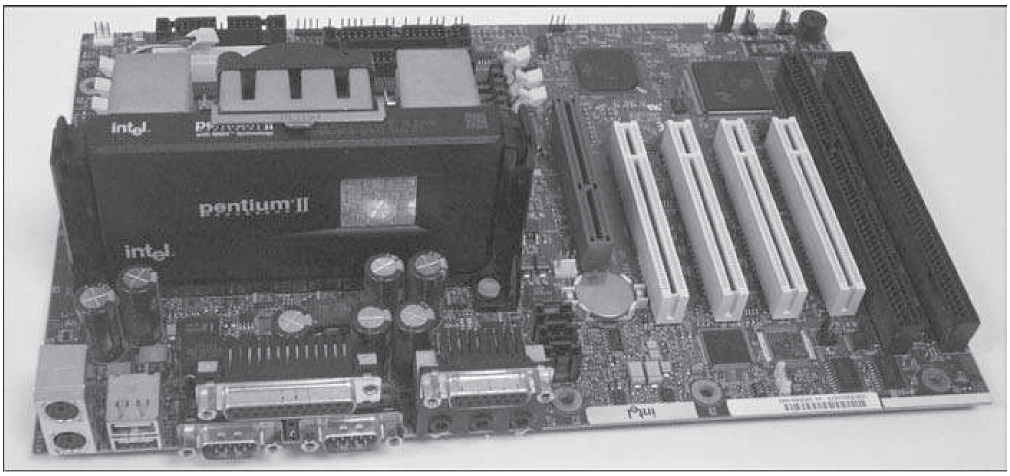

Figure 1-3: Looking at a Pentium II using the SEC packaging.

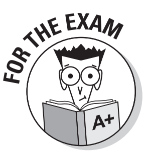

Remember that classic Pentium chips are inserted into socket 5 or socket 7, whereas Pentium II processors are inserted into slot 1. With newer Pentium processors, such as Pentium III and Pentium 4, Intel has moved away from the SEC. The Pentium III is placed in Socket 370 while the Pentium 4 is placed is Socket 423 or Socket 478.

SIMM/DIMM sockets

When you look at a motherboard, one of the first items that should stand out is the processor; the next things you will usually notice are the memory slots that are used to install RAM into the system.

There are typically two types of sockets to install memory: *Single Inline Memory Module (SIMM)* sockets and *Dual Inline Memory Module (DIMM)* sockets. Original Pentium systems typically have either four 72-pin SIMM sockets or two 168-pin DIMM sockets to install memory, while newer motherboards today have up to four DIMM sockets and no SIMM sockets. There are no rules as to how many SIMM or DIMM sockets a motherboard manufacturer may use, as you can see with Figure 1-4. Figure 1-4 shows a motherboard with four 72-pin SIMM sockets *and* two DIMM sockets used to hold memory. SIMMs have been phased out and are only available on older motherboards.

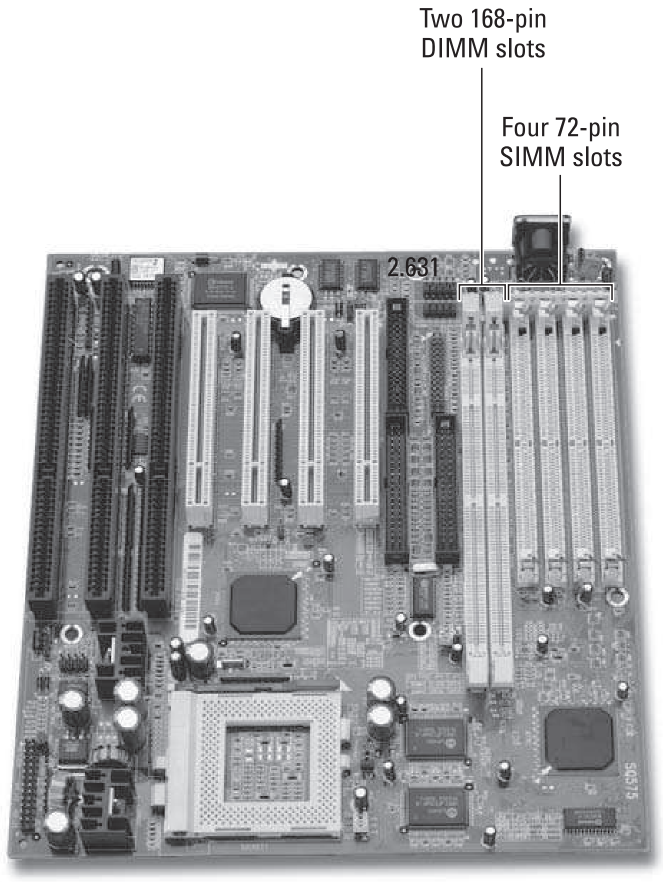

Figure 1-4: Identifying SIMM and DIMM memory slots on the motherboard.

When installing SIMMs in Pentium motherboards, you have to install them in pairs, but when installing DIMMs, you can install them individually. The reason for the difference is that when installing memory, you must fill a memory bank, which is the size of the processor's data path. That is, if you install 72-pin (32-bit) SIMMs onto a Pentium (64-bit) motherboard, then you have to install two modules to fill the 64-bit data path of the processor. DIMMs are 64-bit memory modules, the same number of bits as the data path of the CPU, which is why you're able to install only one at a time. For more information on memory banks and installing memory, check out Book II, Chapter 3.

Cache memory

Cache memory increases performance by storing frequently used program code or data that can be later accessed by the processor. Cache memory is much faster memory than normal RAM and, as a result, is more expensive. The system stores data accessed from RAM in cache memory when the data is accessed the first time, making subsequent requests to the same data faster because the data is accessed from cache (which is faster than RAM) for subsequent calls.

All processors today have integrated cache memory, which is known as *level 1 cache. Integrated cache* is cache memory that is built into the processor, while nonintegrated cache — known as *external cache* — is built outside the processor, typically on the motherboard. The types of cache are as follows:

- **L1 (level-1) cache:** Cache that is integrated within the processor.
- **L2 (level-2) cache:** Cache that is located outside the processor, usually on the motherboard.

Older motherboards implemented cache memory as rows of DIP (dual inline package) chips placed directly on the motherboard. This area was sometimes even labeled "cache." Unfortunately, you can't expect a motherboard to be well-labeled; if you find labels (in English), consider it an added bonus! For more information on chip packages, check out Book II, Chapter 2.

Other systems have implemented the cache as a memory module, so you may see an empty slot on the motherboard that looks like a SIMM slot, but it will really hold a cache module. A lot of times, this will be labeled as "cache" on the motherboard. Figure 1-5 shows L2 cache on an older motherboard.

FOR THE EXAM

Remember, L2 cache is usually located on the motherboard near the processor. That way, data travels over a shorter distance from cache to processor — increasing overall system performance. Also, today's processors implement both L1 and L2 cache in the casing of the processor. For more information on cache memory, refer to Book II, Chapter 3.

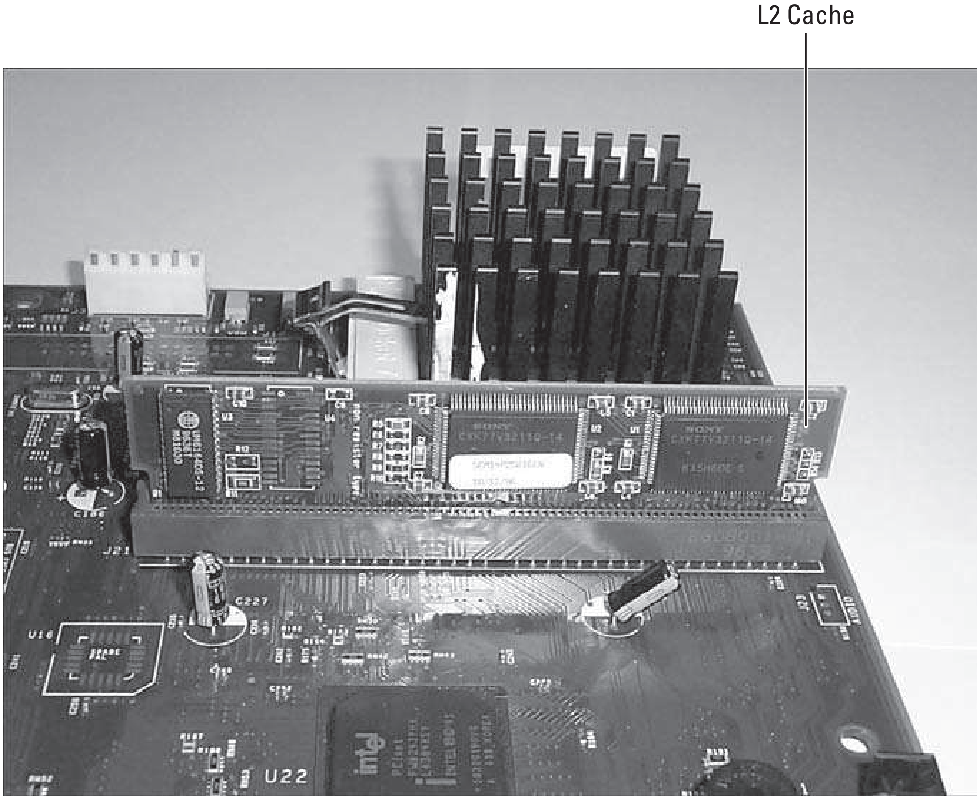

Figure 1-5: Looking at L2 cache located on the motherboard.

Motherboard chipset

Each hardware component in the system has circuitry that is responsible for managing a specific hardware part. This circuitry is known as the *controller* for that specific piece of hardware. For example, access to memory is controlled by the *memory controller*, the hard disk is managed by the *hard disk controller*, and the keyboard is managed by the *keyboard controller*.

The combination of computer chips that hold the logic for these controllers is known as the *motherboard chipset.* Together, the computer chips make up the chipset control communication from the CPU to each of the hardware devices in the system.

Two chips that make up a big part of a motherboard's chipset are the North Bridge and the South Bridge. The *North Bridge* chip is responsible for communication from the CPU to memory and the *advanced graphics port* (AGP) device (more on AGP later in this chapter). The *South Bridge* chip is responsible for communication between the CPU and other devices, such as PCI, ISA, and IDE devices. These two chips contain the bulk of the logic that allows a CPU to communicate with other hardware. Figure 1-6 displays the relationship between the processor and the North Bridge and South Bridge chips.

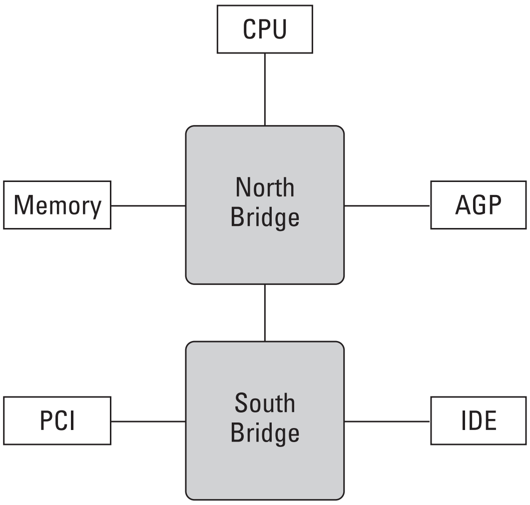

Figure 1-6: Looking at the relationship between the CPU and the motherboard chipset.

Locating the North Bridge and South Bridge chips on a motherboard can sometimes be challenging. The North Bridge chip is typically the second largest chip (after the processor) and typically contains a heat sink or fan on top of the chip to keep it cool. The North Bridge chip is typically located between the processor and the AGP slot, while the South Bridge is normally located farther from the processor — usually beside the PCI slots, as shown in Figure 1-7. Notice in the figure that the North Bridge chip has the words AOPEN on it, while the South Bridge is the chip above the PCI slots.

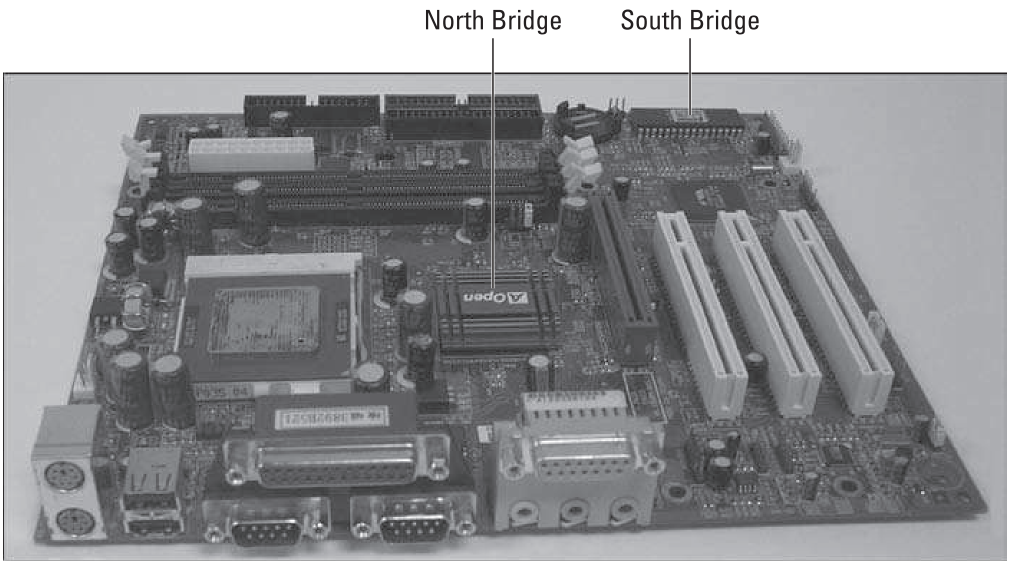

Figure 1-7: Identifying the North Bridge and South Bridge chips on an AOPEN motherboard.

BIOS chip

The *Basic Input-Output System* (*BIOS*) is the low-level program code that allows all the system devices to communicate with one another. This low-level program code is stored in the BIOS chip on the motherboard.

Locating the BIOS chip on the motherboard is easy; it is usually rectangular and generally features a label with the manufacturer's name and the year the chip was manufactured. Some of the popular manufacturers are AMI, AWARD, and Phoenix.

The BIOS chip is a *Read-Only Memory* (ROM) chip, which means that you can read information from the chip, but you can't write to the chip under normal circumstances. Today's implementations of BIOS chips are *EEPROM (Electrically Erasable Programmable ROM),* which means that you can get special software from the manufacturer of the BIOS to write to the chip.

Why would you want to erase the BIOS? Suppose, for example, that your BIOS is programmed to support a hard disk up to 2GB in size, but that you want to install a new, larger hard disk instead. What can you do about it? You can contact the BIOS manufacturer and get an update for your BIOS chip, which is usually a software program (in the past, you generally had to install a new chip). Running the software program writes new instructions to the BIOS to make it aware that there are hard disks bigger than 2GB and provides instructions for dealing with them. But before new instructions can be written, the old instructions need to be erased.

The BIOS chip also contains code that controls the boot process for your system. It contains code that will perform a *power-on self-test (POST),* which means that the computer goes through a number of tests, checking itself out and making sure that it is okay. After it has made it past the POST, the BIOS then locates a bootable partition and calls on the master boot record, which loads an operating system. Figure 1-8 shows a BIOS chip on a motherboard. For more information on the system BIOS, refer to Book II, Chapter 4.

Battery

The computer keeps track of its inventory in what is known as the *Complementary Metal-Oxide Semiconductor* (*CMOS*). CMOS holds a listing of system components, such as the size of the hard disk, the amount of RAM, and the resources (IRQs and I/O addresses) used by the serial and parallel ports.

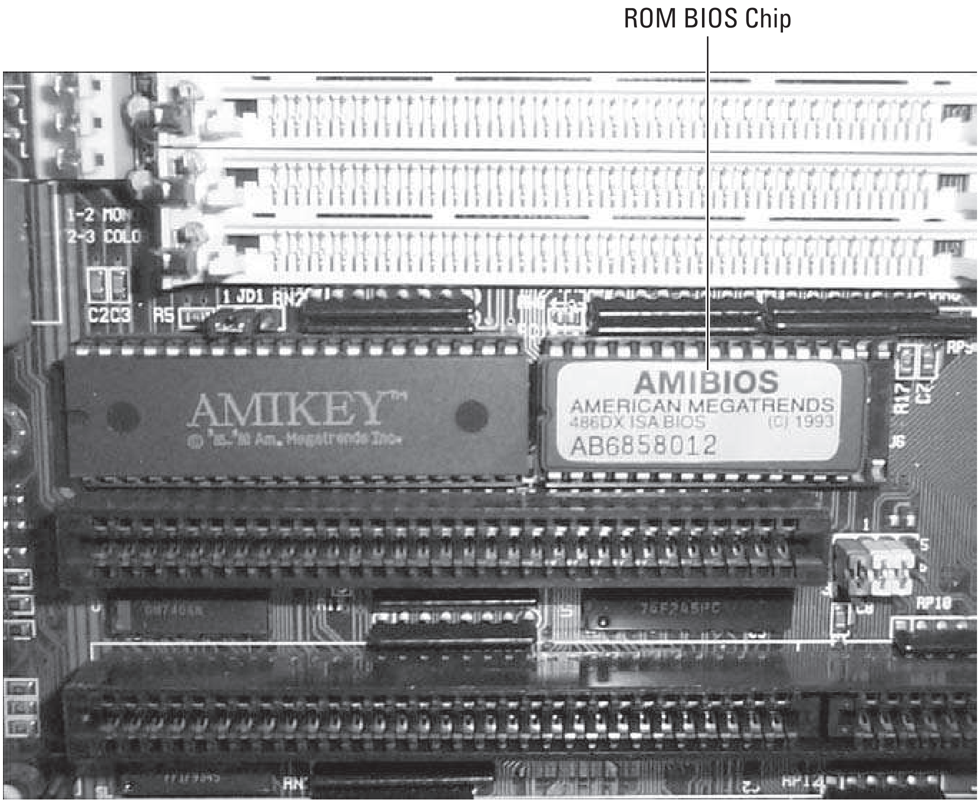

Figure 1-8: Looking at a BIOS chip located on the motherboard.

This inventory list is stored in what is known as CMOS RAM, which is a bit of a problem because RAM loses its content when the power is shut off. You don't want the computer to forget that it has a hard disk or forget how much RAM it has installed. To prevent this sort of problem, a small watch-like battery on the motherboard maintains enough energy that CMOS RAM doesn't lose its charge. If CMOS RAM loses its charge, it results in the CMOS content being lost. Figure 1-9 identifies a battery on the motherboard. For more information on CMOS, check out Book II, Chapter 4.

Expansion slots

Most motherboards have one or more *expansion slots,* which serve the purpose of adding functionality to the computer. For example, assume that your computer doesn't have sound capability — you can install a sound card into the expansion slot to add that capability to your system.

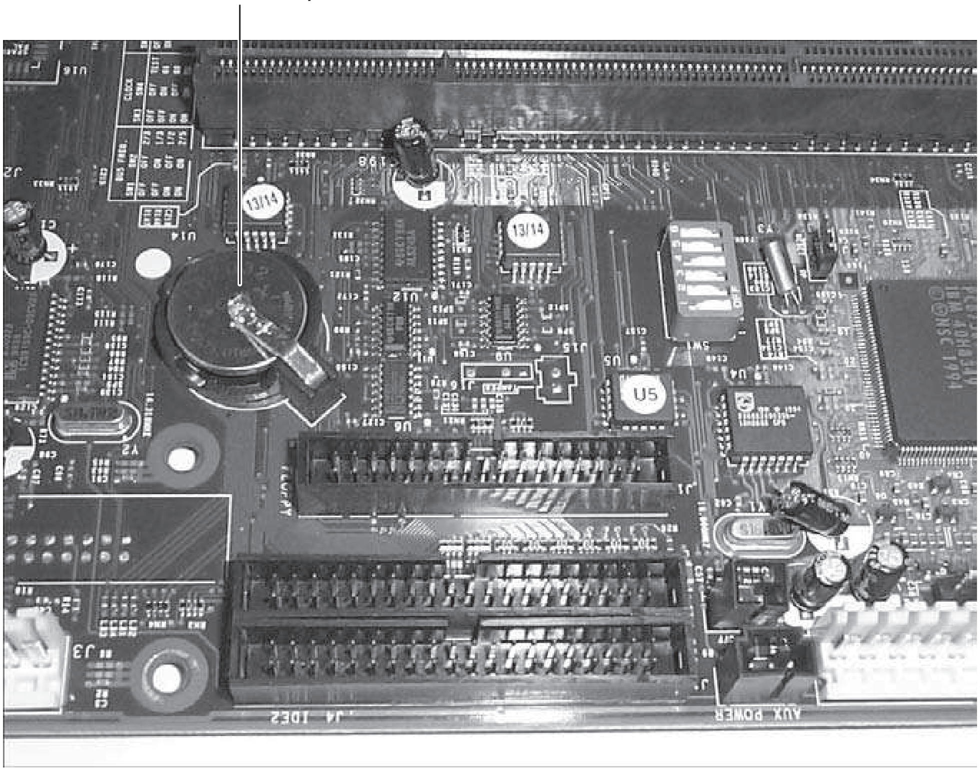

Figure 1-9: Identifying the battery on the motherboard that is used to maintain a charge to CMOS RAM.

Expansion slots come in different varieties, and it is extremely important to understand the benefits of each type. We discuss these issues later in the chapter, in the section titled "Understanding Bus Architectures." For now, I just want you to be able to identify the expansion slots on the motherboard.

If you look at the motherboard, you can see a number of expansion slots. There are probably some white, narrow PCI slots on the board, as well as a tan-colored AGP slot (used for video cards). You may also see some larger black slots; these are ISA slots used by older devices. Most motherboards today do not have ISA slots, or may only have one. Figure 1-10 displays ISA, PCI, and AGP expansion slots used to add expansion cards to the system. For more information on expansion slots, refer to the "Understanding Bus Architectures" section, later in this chapter.

Ports and connectors

There are a number of *ports* on the back of the motherboard that connect the keyboard, mouse, printer, and other devices to the system. This section identifies those ports. Figure 1-11 displays a number of built-in *input/output* (I/O) ports on the back of an ATX motherboard.

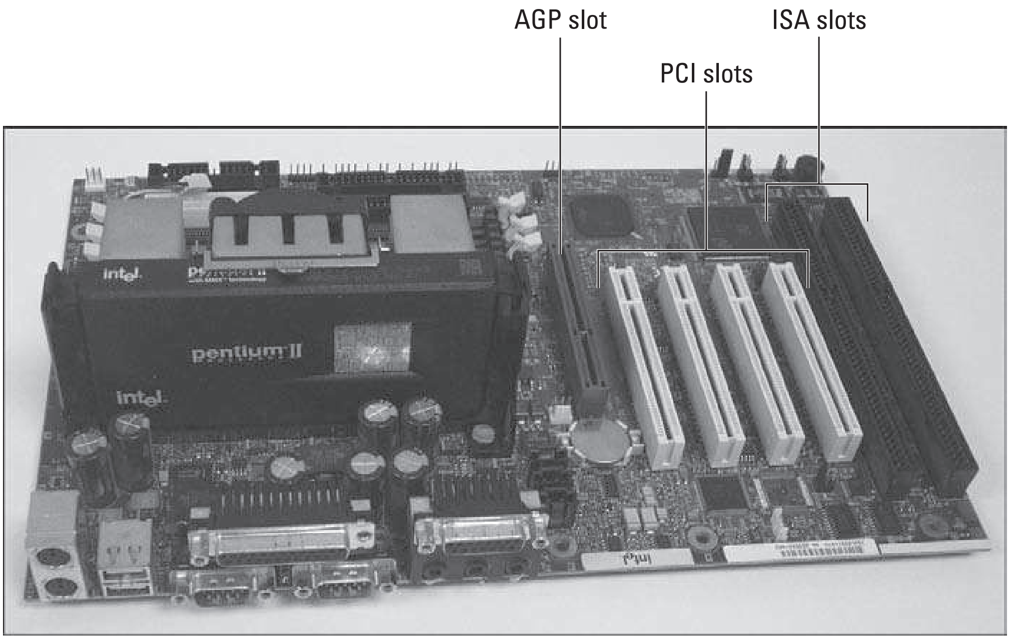

Figure 1-10: Identifying expansion slots such as AGP, PCI, and ISA on a motherboard.

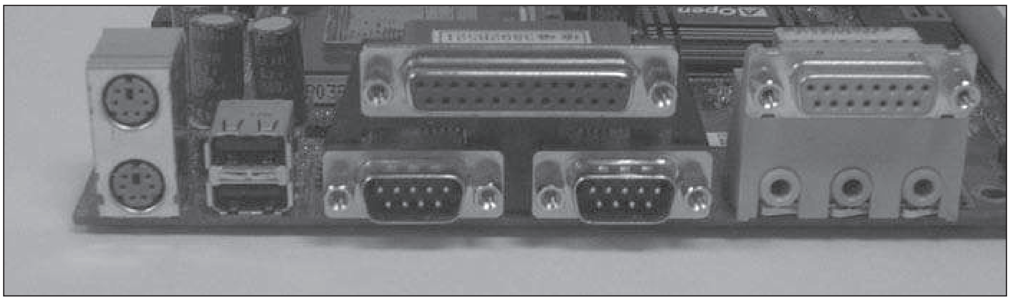

Figure 1-11: Looking at the built-in ports on the back of an ATX motherboard.

Serial ports

Most motherboards have serial ports integrated directly into the board. The *serial ports* are also known as *communications (COM) ports.* The reason that they are called *serial ports* is because they send data in a series — a single bit at a time. If eight bits of data are being delivered to a device connected to the COM ports, then the system sends the eight bits of data, one bit at a time, in single file. Typically, there are two COM ports — COM1 and COM2 — on each system.

The official standard that governs serial communication is known as *RS-232,* and you may see serial ports referred to as *RS-232 ports.*

You usually connect an external modem or a serial mouse to a serial port. Each of these devices is used for communication; a modem allows your computer to talk to another computer across phone lines, while a serial mouse allows you to communicate with the system. Figure 1-12 shows two serial ports connected to a motherboard.

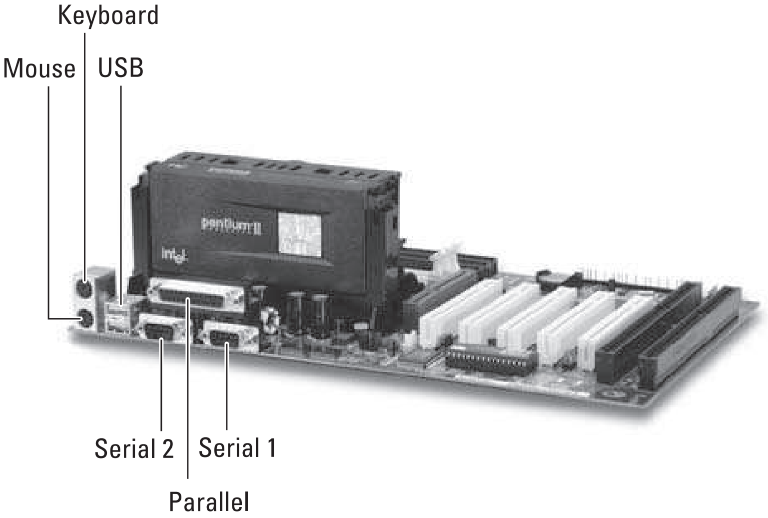

Figure 1-12: Identifying the integrated ports on the back of an ATX motherboard.

Serial ports on the back of the motherboard are one of two types:

- **DB9-male** is a serial port with 9 pins.
- **DB25-male** is a serial port with 25 pins.

Parallel port

Another type of connector that you will have on the back of the motherboard is the parallel port. The *parallel port* is also known as the *printer port,* or *LPT1.* The parallel port gets its name by being able to send information eight bits at a time. Whereas serial ports send only one bit at a time in single file, parallel ports can send eight bits in one operation — side-by-side rather than single file. Refer to Figure 1-12 to see a parallel port connected to a motherboard.

The parallel port, which is known as DB25-female, has 25 pins and is located on the back of the motherboard. Looking back at Figure 1-12, you can see the parallel port located above the two serial ports.

It is important not to confuse the serial port with the parallel port. The serial port is a male port (meaning that the port has a number of pins in it), whereas the parallel port is a female port (meaning that it does not contain the pins but the pin holes).

You connect the parallel port to a printer by using a parallel cable that has a different type of connector at each end. On one end of the cable is a DB25 connector that attaches to the parallel port on the back of the computer. The other end of the cable (the end that connects to the printer) has a 36-pin Centronics connector.

Remember, a standard printer cable has a different type of connector on each end. One end has a DB25-male connector with 25 pins, while the other end has a 36-pin Centronics connector.

Video adapter

In the past, a motherboard came with a built-in *video adapter,* sometimes called a *video card* or *video controller.* The video adapter is responsible for converting digital data from the processor and preparing the information to be displayed on the screen. Figure 1-13 displays a video adapter port — you can identify it by the three rows of five pins. The video port is a 15-pin female port.

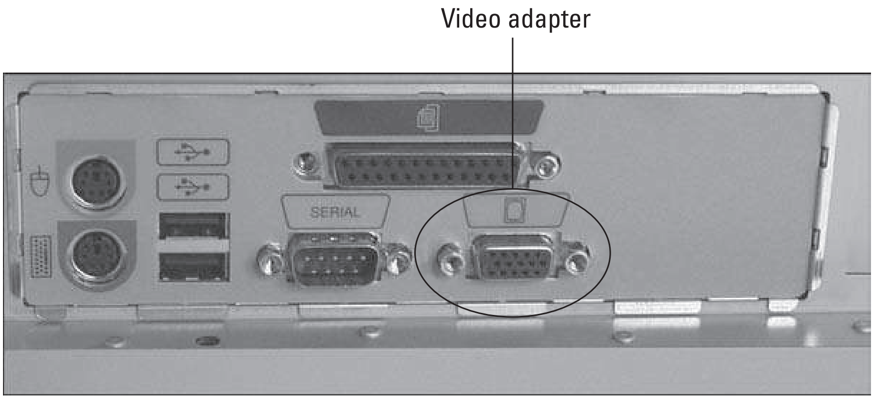

Figure 1-13: Identifying the video adapter port on the system.

Many systems today use the ATX motherboard form factor and, as a result, have an AGP slot to hold the video adapter. This means that the video adapter is not integrated into the motherboard like it was in the past.

Figure 1-14 shows how information flows from the computer system to the monitor. The following steps refer to the numbers in Figure 1-14.

1. **The video adapter is responsible for receiving digital data from the processor, which instructs the video adapter on how the images are to be drawn on the screen.**
2. **The video adapter stores the information about drawing the images in its memory and starts converting the information into analog data that the monitor can understand.**
3. **The data is sent in analog format from the video adapter to the monitor.**

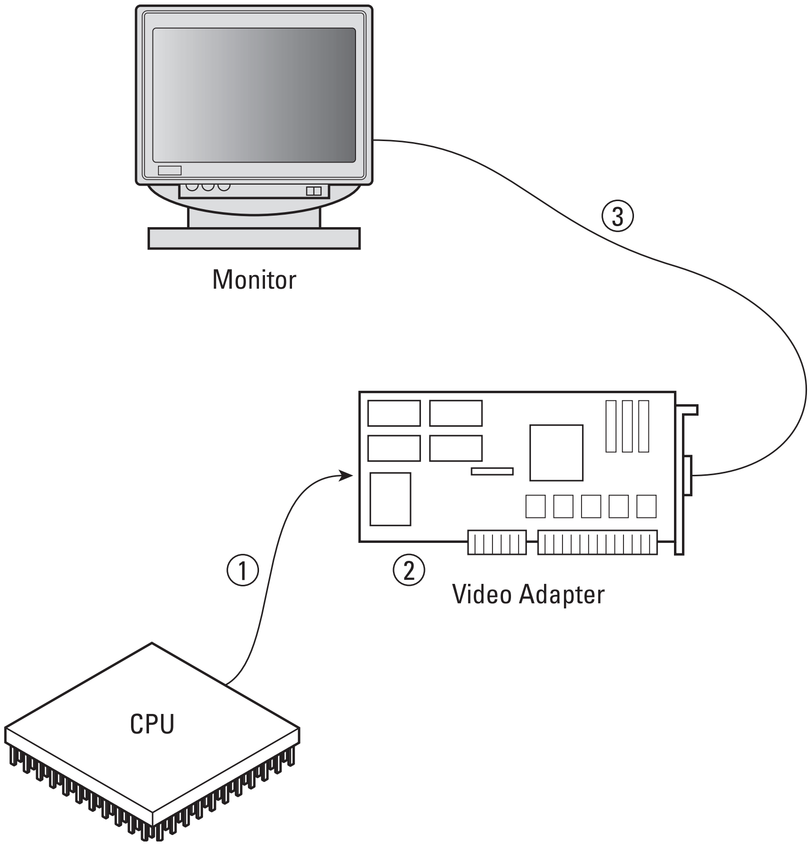

Figure 1-14: Looking at how information flows from the processor to the display.

Keyboard/mouse connectors

The mouse and keyboard connectors on motherboards today are most likely PS/2 style connectors or USB connectors. Let's focus on PS/2 connectors for a now. A PS/2 connector is a small circular six-pin connector. In Figure 1-15, you can see the keyboard and mouse connectors on the left side of the diagram.

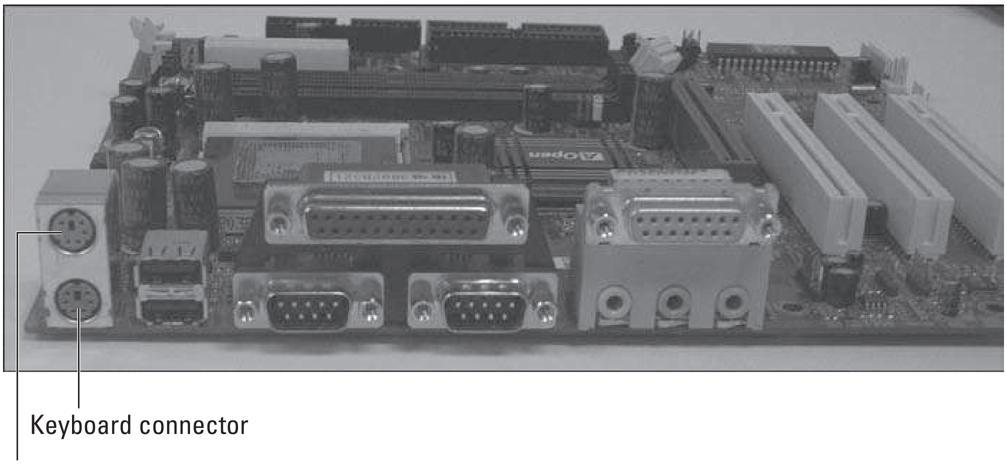

Figure 1-15: Identifying the PS/2 connections used for a keyboard and mouse.

Older motherboards may have a *DIN keyboard connector,* also known as an *AT connector,* which you can see on AT and Baby AT motherboards. These systems did not have any other ports on the back of the system, so you would need to insert an I/O card for other ports (such as serial and parallel ports).

Sound

Most motherboards today have built-in sound capabilities, allowing you to connect speakers and a microphone to the computer. Figure 1-16 shows the integrated sound ports on a motherboard. There are three different ports on the integrated sound card:

- **Line-in:** The line-in port is normally blue and allows you to connect many audio sources to the system.
- **Line-out:** The speaker port is normally green and allows you to connect speakers to the computer.
- **Microphone:** The MIC-in port is red and allows you to connect a microphone to the system for recording purposes.

Network interface card and modem

A number of systems today have built-in network support via an integrated *network interface card* (NIC), or network card for short. These systems may have a built-in modem as well. The built-in network card has an RJ45 port on the back of the system that looks like an oversized telephone jack, as shown in Figure 1-17.

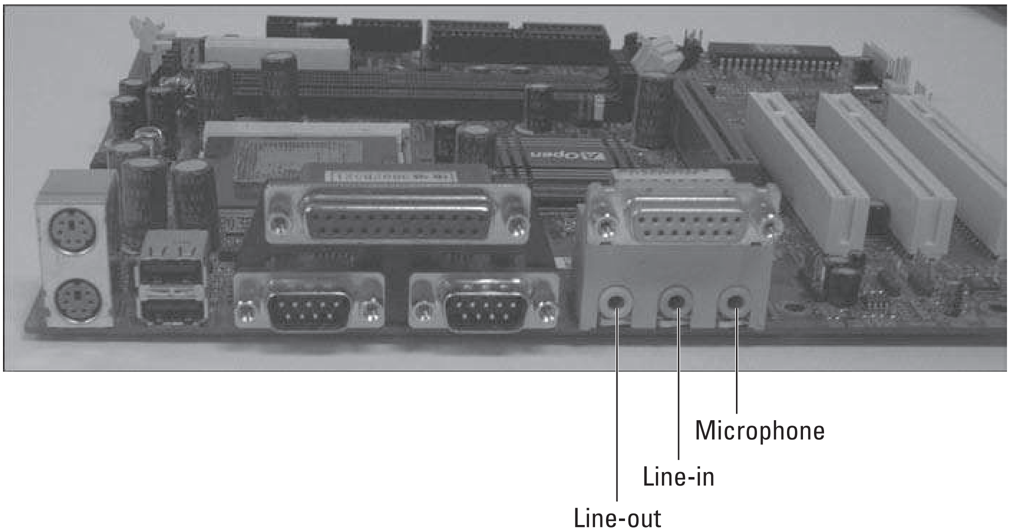

Figure 1-16: Identifying the sound ports on the system.

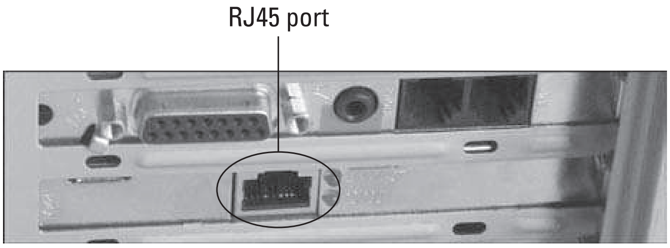

Figure 1-17: Identifying the RJ45 port.

USB ports

Universal Serial Bus (USB) is a high-speed serial technology that transfers data at 12 Mbps (USB 1.0) and 480 Mbps (USB 2.0). One of the major benefits of USB is the fact that all USB devices use the same type of connector, so you won't have to guess which ports to connect the mouse, keyboard, or scanner to. If they are all USB devices, they connect to the same type of port!

TIP

The newest USB standard, called USB 2.0, has a transfer rate of 480 Mbps, which is much faster than the USB 1.0 standard. Be sure that you also have USB 2.0 drivers installed to leverage the performance benefits of your USB 2.0 devices.

USB devices also support *daisy chaining.* For example, you can connect Device A to the back of the computer and then connect Device B to Device A, and so on. You can connect up to 127 devices to a system using USB. Figure 1-18 identifies the USB ports on the back of an ATX motherboard.

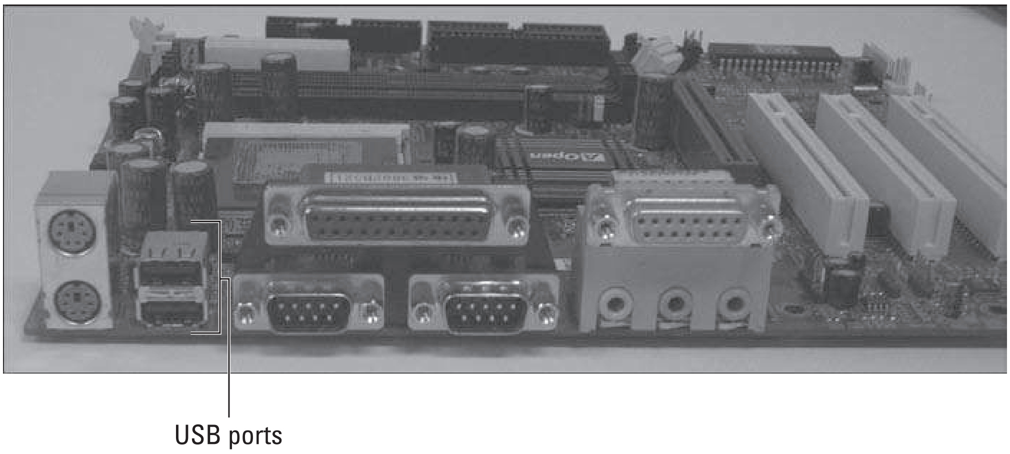

Figure 1-18: Identifying the USB ports on the back of the motherboard.

A USB device that connects to the computer and then has other devices connected to it is considered a *hub device.* If you don't have a USB device that can act as a hub device, you can purchase a specific USB hub that allows you to chain four or more other devices off of it. With a USB hub, you can easily increase the number of USB ports your system has by connecting USB devices to the hub and connecting the hub to the back of the computer. Figure 1-19 shows a USB hub.

Figure 1-19: A USB hub with four ports.

FireWire (IEEE 1394)

When USB 1.0 was introduced, it ran at 12 Mbps. This was a fairly good speed for most types of devices but was a little too slow when it came to multimedia devices, such as digital video cameras. Typically, these types of devices use a *FireWire* connection, which has a transfer rate of up to 400 Mbps and supports 63 devices in a chain. This is a huge jump compared to the USB 1.0 standard. The official standard that defines FireWire is known as the IEEE 1394 — be sure to remember that for the exam!

Just as USB had a second version with a faster transfer rate so does FireWire. The second version of FireWire is defined as the IEEE 1394b standard and transfers data at 800 Mbps! This second version of FireWire is also known as FireWire 800. Figure 1-20 shows a digital video camera being plugged into a FireWire port.

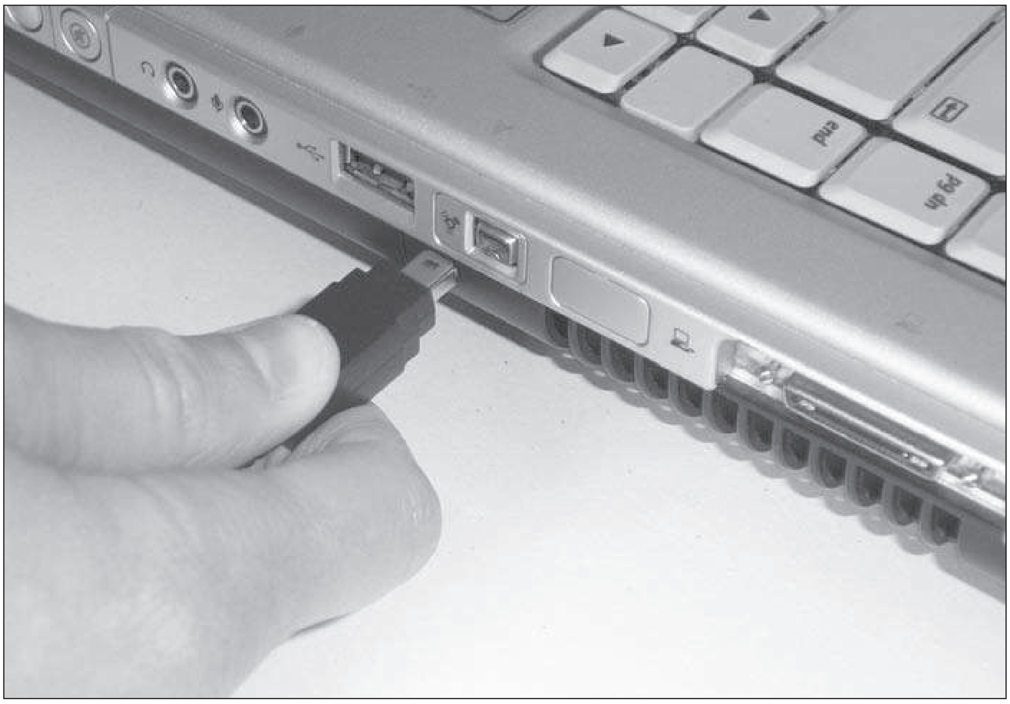

Figure 1-20: A digital video camera being connected to a system by the FireWire port.

FOR THE EXAM

For the exam, remember that the original version of FireWire runs at 400 Mbps and is known as IEEE 1394. The second version of FireWire, FireWire 800, is also known as IEEE 1394b and runs at 800 Mbps. Also, FireWire supports 63 devices in a daisy chain.

For more information on common ports and connectors such as keyboard, mouse, serial, parallel, USB, and FireWire refer to Book III, Chapter 1.

Power connectors

All the devices connected to the motherboard need to get power from somewhere, so the power supply is connected to the motherboard, which supplies power to the board and its components. The following sections discuss power connectors on older and newer motherboards.

Older motherboard power connectors

Figure 1-21 shows power connectors on an older motherboard. There are power cables coming from the power supply to connect to the motherboard with very unique connectors on the end. These power connectors coming from the power supply that connect to the motherboard may be labeled as P1 and P2, or on some systems, P8 and P9.

Motherboard power connector

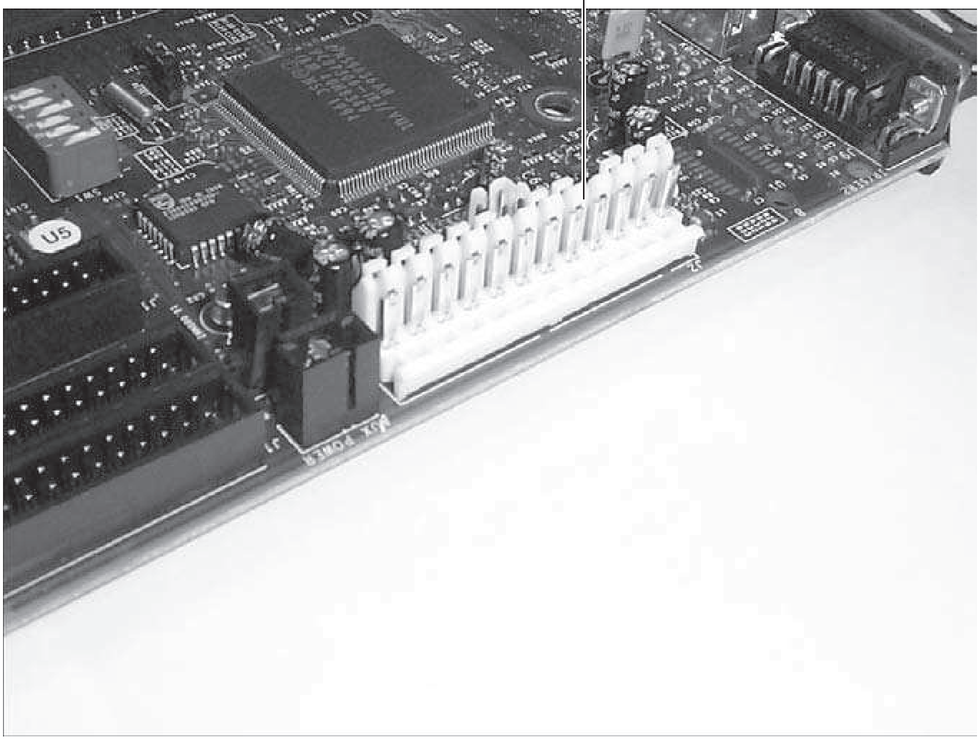

Figure 1-21: Looking at the power connectors on an older mother-board.

You have to be extremely careful to make sure that the connectors on the cable coming from the power supply to the motherboard are inserted properly, or you could damage the motherboard. Often, these connectors are *keyed* (meaning that they can go in only one way) so that you cannot put both of the connectors in the wrong way. These older power connectors supplied power in 5 volts and 12 volts.

ATX power connectors

Newer ATX motherboards use a different power connector than the one shown in the preceding section. The ATX power connector supplies 3.3 volts, 5 volts, and 12 volts. The ATX power connector, shown in Figure 1-22, is typically labeled as P1.

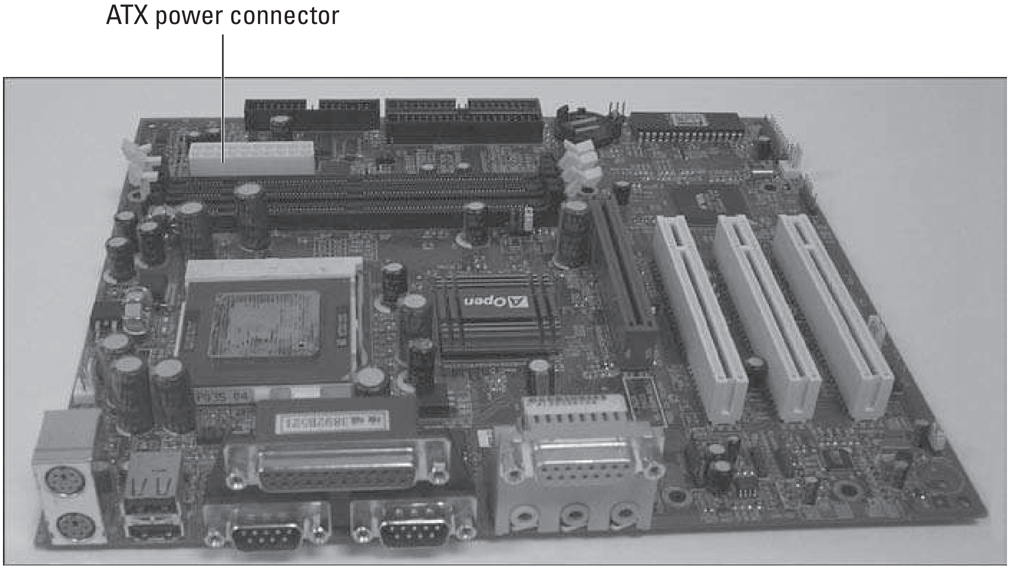

Figure 1-22: The ATX power connector on an ATX motherboard.

Some systems, like ones that use the Pentium 4 boards, use an additional power connector, known as the P4 connector, which supplies an additional 12 volts to the ATX board. Figure 1-23 displays the P4 power connector.

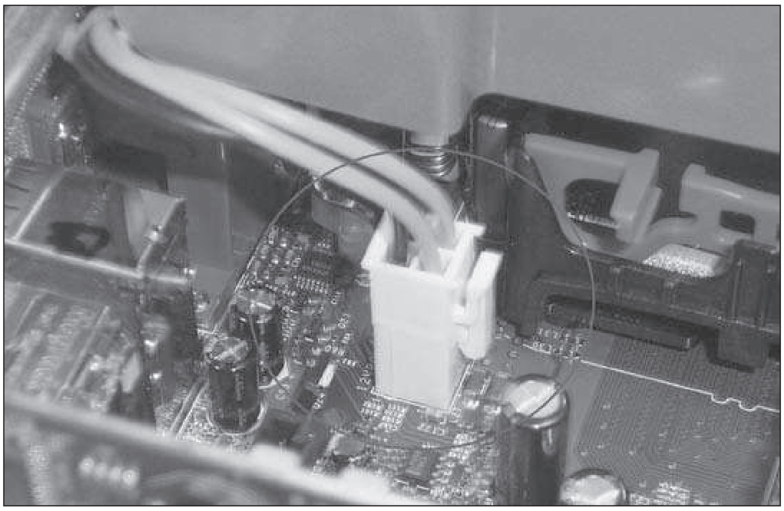

Figure 1-23: The P4 power connector on an ATX motherboard.

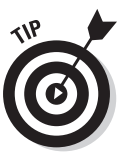

For more information on power supplies and their connectors, check out Book II, Chapter 6.

Drive connectors

You need to be able to identify the different types of connectors that link hard drives to your system. As you may already be aware, the hard drives are used to store information permanently on the computer but in order to access that information the drives have a physical connection to the system via the motherboard.

There are four major types of drives in systems today, *IDE* (*Integrated Drive Electronics*) drives, *SATA* (*Serial Advanced Technology Attachment*) drives, SCSI drives, and floppy drives. Each type of drive has its own type of connection on the motherboard. Before you purchase a hard disk to add to the system, you need to be aware of what types of drives your motherboard supports.

IDE connections

IDE drives have been around since the 1980s, and although the technology has improved from a performance perspective, IDE drives connect to the system in the same way they always have. If your motherboard supports IDE, you will have two IDE connectors that are made up of 40 pins each, as shown in Figure 1-24.

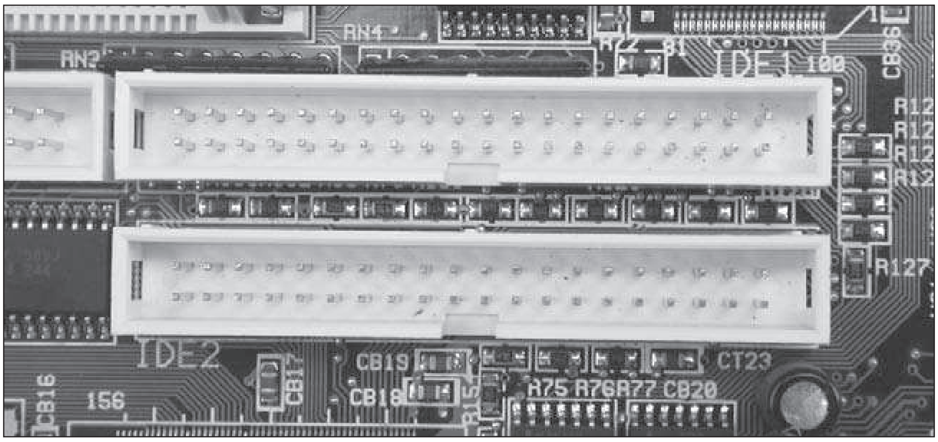

Figure 1-24: IDE connectors on the motherboard.

You will connect the drive to the connector on the motherboard by using a 40-wire or 80-wire *IDE ribbon cable.* This ribbon cable typically has two connectors on it — one end connects to the drive, while the other end connects to the motherboard. You can also find IDE ribbon cables with three connectors that allow you to connect two drives to each IDE connector on the motherboard. This means that you can have up to four IDE devices on a system. Figure 1-25 shows an IDE ribbon cable connector.

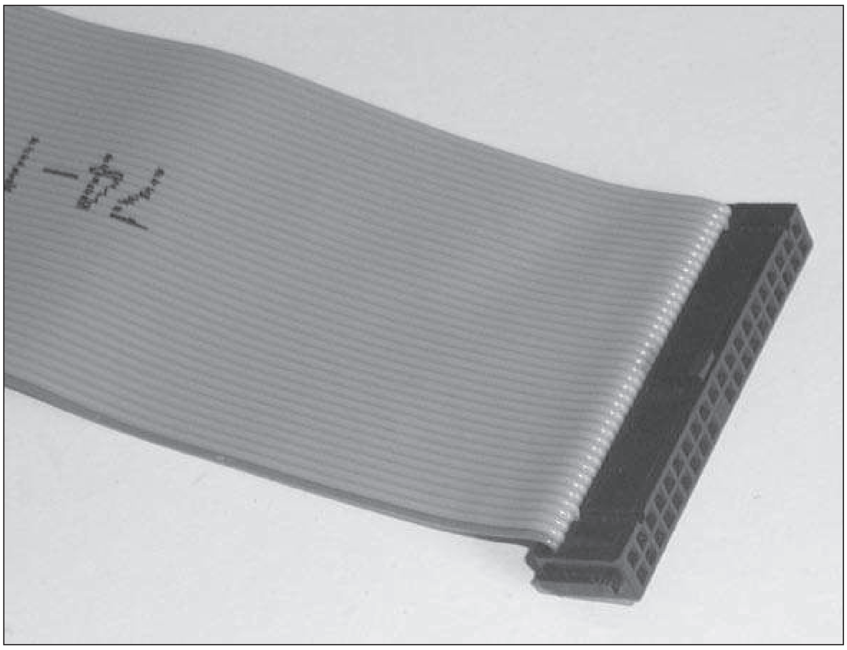

Figure 1-25: A 40-wire IDE ribbon cable.

When connecting the IDE ribbon cable to the drive and motherboard, the colored wire on the ribbon cable connects to pin 1 on the connector. This is known as the *pin-1 rule.* Pin 1 is normally labeled on the motherboard and drive. If it isn't labeled, see whether the manufacturer has labeled pin 40 — if so, pin 1 is at the other end!

TECHNICAL STUFF

IDE controllers is a popular term used in the computer industry for the IDE connectors. Although in theory these are not controllers, it is a term used in the industry to describe the IDE connections on the motherboard. The *actual* IDE controller is the circuitry located on the circuit board on the drive itself — it is responsible for controlling the flow of information to and from the drive.

SATA connections

Limitations of the IDE architecture have kept its data transfer rate around 150 MBps. As drives become more powerful, a new standard is needed. The first new standard to replace IDE is known as *SATA* and is now becoming popular in desktop computers. SATA can reach transfer rates of up to 600 MBps! This is quite a bit (450 megabytes, to be exact) faster than the 150 MBps currently offered by high-end IDE drives.

SATA uses it's own unique four-wire cable to connect to the motherboard. Figure 1-26 shows a SATA cable connected to the SATA connector on the motherboard. Notice that the cable is quite a bit thinner than the IDE ribbon cable; this allows for better airflow in the system and improves overall temperature control of the computer.

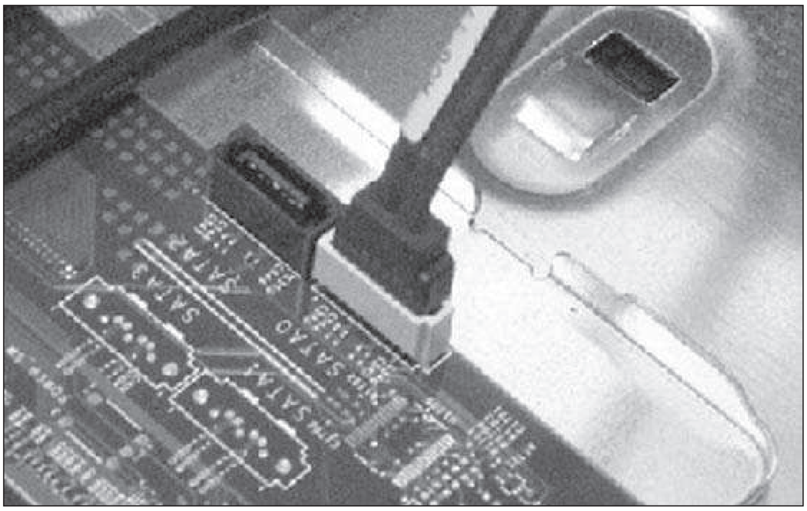

Figure 1-26: A SATA cable connected to the motherboard.

Unlike IDE drives, you cannot connect more than one SATA drive to a connector. For example, if your motherboard has two SATA connectors then you can only connect two SATA drives to the system unless you purchase a SATA card that has additional connectors.

SCSI controller

Some high-end machines, particularly those designed for use as servers, may have a controller on the motherboard with 50 pins on it. This is the footprint of a *SCSI (Small Computer System Interface) controller*. Because SCSI devices outperform IDE devices, SCSI controllers are extremely popular for servers (which have greater hard disk access and storage needs than regular desktop computers). To connect a SCSI drive to the 50-pin SCSI connector on the system, you use a 50-wire ribbon cable.

Remember, IDE uses a 40-pin connector that a 40/80-wire ribbon cable connects to, and an internal SCSI connector has 50 pins that connect to a 50-wire ribbon cable.

Lab 1-1 and Lab 1-2 will help you identify the major motherboard components on the motherboard. Lab 1-1 and Lab 1-2 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Floppy disk connectors

Located very close to the IDE connectors on the motherboard, you should see a smaller *floppy drive connector* which contains 34 pins instead of the 40 pins found with the hard drive IDE connectors. The floppy drive connector on the motherboard is used to connect the floppy drive to the motherboard using a 34-wire ribbon cable.

When connecting the floppy drive to the system, you will notice that the wires on one end of the ribbon cable are twisted. This twisted end must be connected to the floppy drive. The opposite, untwisted end connects to the motherboard. Also note that one wire of ribbon cable is colored, usually red, which indicates wire one. Like the IDE drives, you need to connect wire one to pin 1 on the motherboard and on the floppy drive.

To find out more information about IDE, SATA, SCSI, and floppy drives, check out Book II, Chapter 5.

Jumpers and DIP switches

A jumper is a set of pins that have a plastic cap enclosed over them to create an electrical connection. The plastic cap contains a piece of metal that makes contact with the pins and creates the electrical circuit. The circuit that is created enables a feature on the motherboard. Most motherboards (and expansion cards) use *jumpers* to implement different settings. Figure 1-27 displays a jumper on an expansion card.

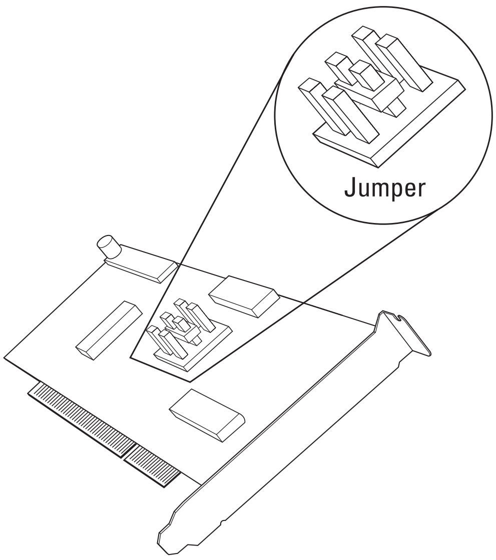

Figure 1-27: Identifying a jumper.

Notice in the figure that the jumper has three sets of pins that the cap may be placed over. The idea behind the three sets of pins is that each set of pins would enable a different setting. For example, looking back at Figure 1-27, the three different sets of pins may be used to assign three different IRQs to the card. You choose which IRQ is assigned to the card by setting the jumper over a set of pins. Keep in mind that you no longer assign IRQs with jumpers, it was something that was done years ago (for more information on IRQs, check out Book III, Chapter 4).

Today you will find jumpers on motherboards, hard drives, CDROM drives, and DVD drives. There are many different features that can be enabled or disabled on a motherboard using jumpers. For example, there usually is a jumper on the motherboard that is used to clear the CMOS password of a system, to change the voltage supplied to the processor socket, or to change the speed of the motherboard. In order to know what jumper to set you need to check the documentation for the motherboard.

Another popular component of a motherboard or expansion cards in the past that was used to enable or disable different features is the *dual inline package (DIP) switch*. A DIP switch (seen in Figure 1-28) is a set of switches that can be turned on or turned off to enable functionality on the board. In order to know what to set for on/off combinations you would need to consult the documentation for the board.

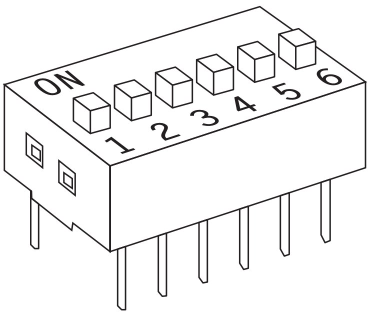

Figure 1-28: Looking at a DIP switch.

Identifying the Types of Motherboards

Now that you understand some of the major components of the motherboard (system board), it is important to mention the different motherboard form factors. A *motherboard form factor* just describes the dimensions of the motherboard and the layout of the motherboard components.

It is important to understand the different motherboard form factors because you can't just take *any* motherboard and place it in a computer case. You must put a full AT motherboard in a full AT case, a Baby AT board in a Baby AT case, and an ATX board in an ATX case. Figure 1-29 shows the three major types of motherboards and gives you an idea of size and shape differences between the three types.

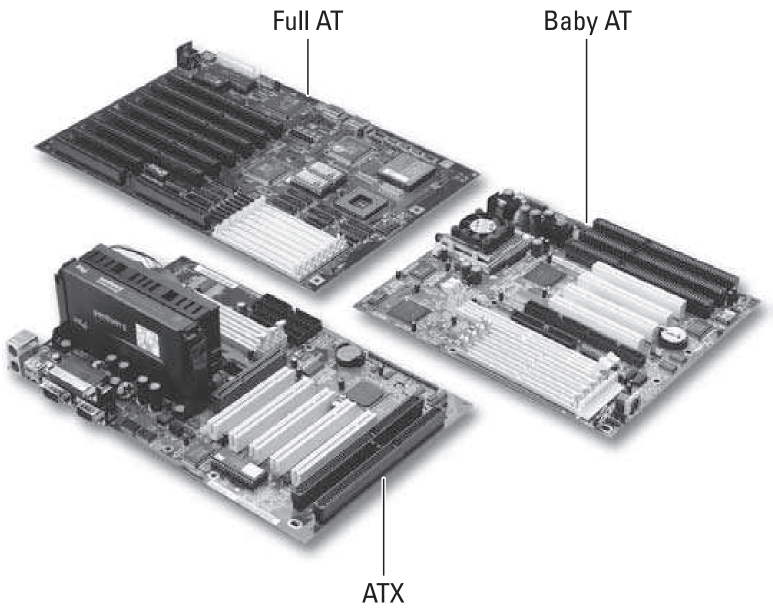

Figure 1-29: Looking at different motherboard form factors.

Full AT

The first type of motherboard that we want to talk about is the *full AT* motherboard. The full AT motherboard is 12 inches wide and 11 inches long and is easily recognized by the fact that it only has a keyboard connector on the back of the motherboard — it contains no other I/O ports.

The full AT suffers from a problem with accessing some of the items on the motherboard because the drive bays hang over the motherboard. This situation makes installation and troubleshooting of the components on the motherboard very difficult.

Another problem with the layout of the full AT board is that the expansion cards, once inserted into the systems, cover the processor. This situation leads to cooling problems because ventilation is insufficient to keep the chip from overheating. Figure 1-30 displays a full AT motherboard being installed in a full AT case.

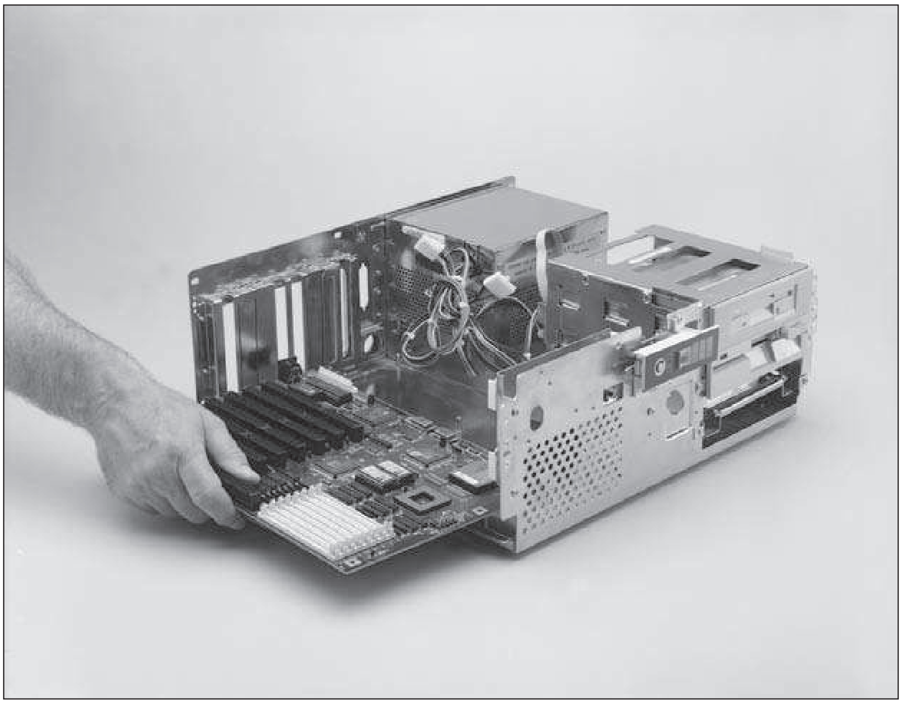

Figure 1-30: Looking at a full AT motherboard.

Baby AT

The *Baby AT* motherboard form factor had been one of the most popular motherboard types until recent years. The Baby AT board is 8.5 inches wide and 10 inches long. This motherboard can be easily recognized because it usually has a DIN keyboard connector in the top-right corner of the board. This keyboard connector is the only I/O connector on the back of the motherboard.

The Baby AT board is about two-thirds the size of the full AT board and typically incorporates a socket 7 ZIF (zero insertion force) slot for classic Pentium processors. The Baby AT board usually has a mixture of ISA/EISA and PCI slots located on the motherboard and includes a Plug and Play BIOS. Figure 1-31 shows a Baby AT motherboard and identifies the popular components.

Take a minute to consider some of the key components on the Baby AT motherboard. You can see the socket 7 ZIF slot at the bottom of the motherboard where the processor is to be installed. Also notice the SIMM and DIMM sockets on the right side of the motherboard, which house the system memory. To the left of the SIMM and DIMM slots, you can see the primary and secondary EIDE connectors (sometimes called controllers) for connecting the hard drives to the board. To the left of the EIDE controllers, notice the types of expansion slots that are used: There are four PCI slots and three EISA slots. Above the PCI slots, you can also see a silver circle, which is the CMOS battery.

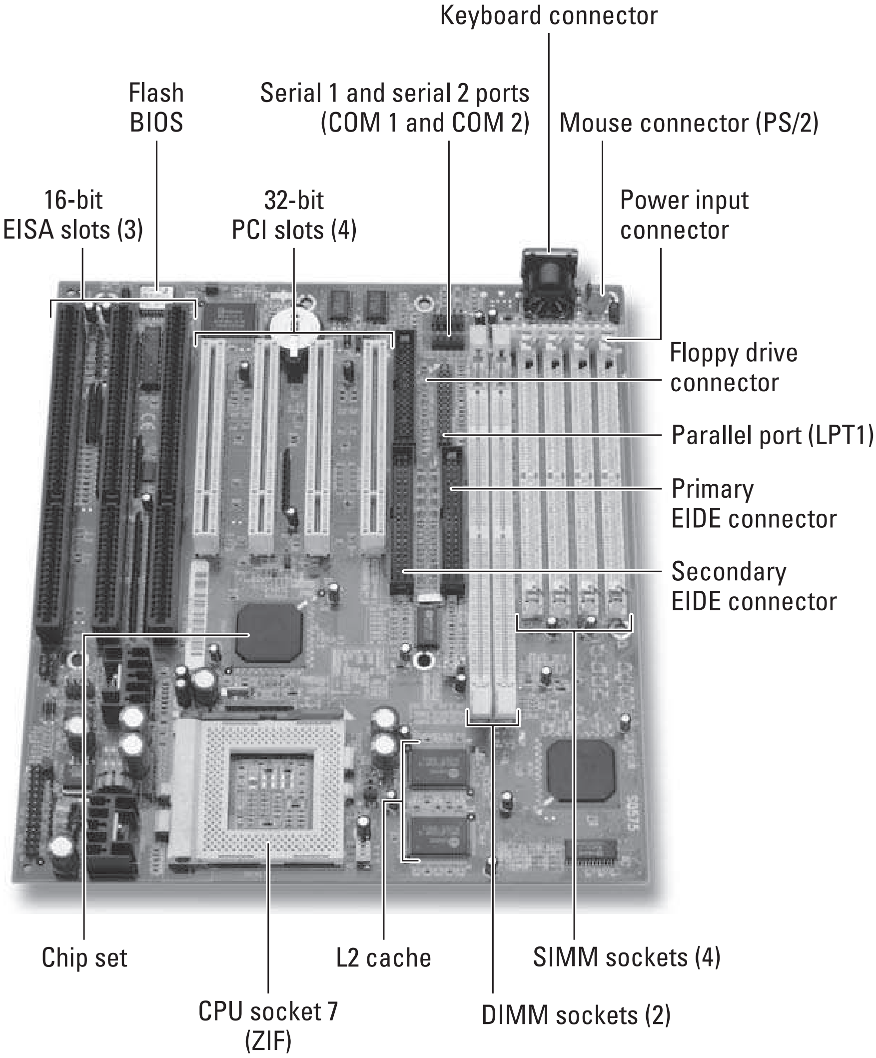

Figure 1-31: Identifying components on a Baby AT motherboard.

LPX/NLX

In an effort to allow computers to take up much less space, a slimline desktop system was designed with a smaller motherboard. After the era of the Baby AT came the *LPX (low profile extended),* which was then replaced by *NLX (New Low-profile eXtended)* motherboard. Both motherboard types served the same purpose — to create low-profile computers.

The NLX motherboard is identifiable by the I/O ports along the back of the motherboard. This motherboard is unlike the full and Baby AT because they incorporated only the keyboard connector. The NLX provides a keyboard and mouse connector, serial and parallel ports, and a video connector.

The NLX form factor is 9 inches wide by 13.6 inches long and uses a riser card to house the bus architectures. The riser card typically connects to the side of the motherboard and is then secured along the side of the case. Figure 1-32 shows an NLX motherboard with a riser card. (We cover bus architectures in the section, "Understanding Bus Architectures," later in this chapter.)

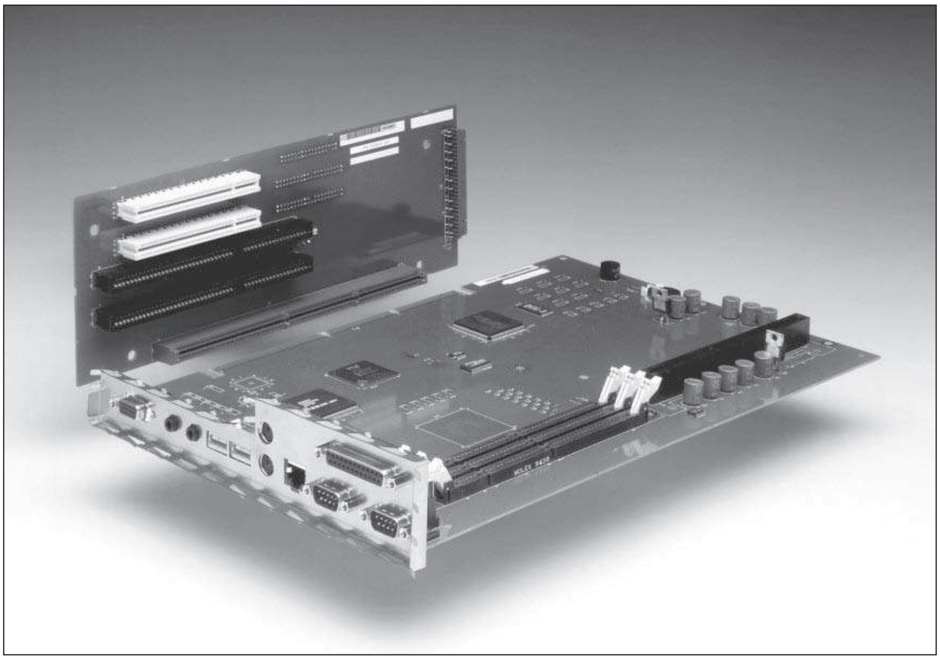

Figure 1-32: An NLX form factor motherboard.

ATX

In 1995, Intel wanted a motherboard that would support the Pentium II processor and the new AGP slot, so the ATX form factor was built (shown in Figure 1-33). The ATX board is 7.5 inches wide and 12 inches long and has most of the I/O ports integrated directly into the board, including USB ports.

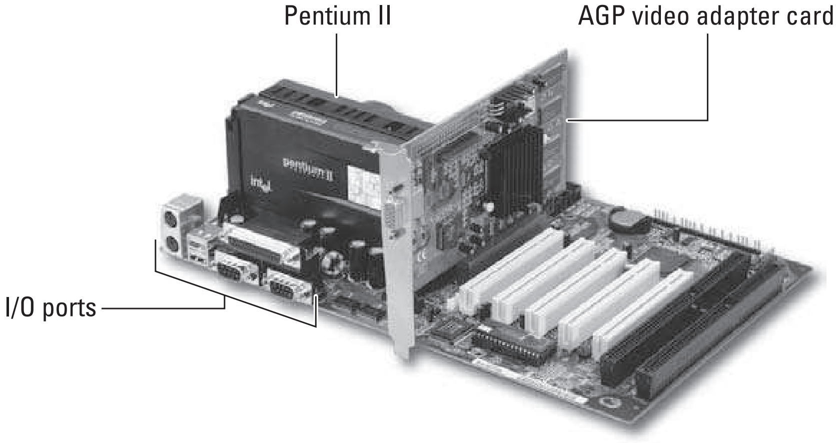

Figure 1-33: The position of the I/O ports on an ATX motherboard.

Remember, the ATX motherboard incorporates the I/O ports and includes an AGP slot for high-performance video cards. Figure 1-33 displays the ports on the back of the ATX motherboard — they are clustered in the left corner of the board and do not spread across the length of the board like they do with the NLX form factor.

The ATX board introduced a 100 MHz system bus and has been increased to speeds of 533 MHz and more. The ATX motherboard has one AGP slot for the video card, which means that the built-in I/O ports on the back of the board do not have a built-in video card like the NLX. The ATX board also has *soft power support,* which allows software developers to create software that controls the startup and shutdown of the system.

The ATX form factor rotated the Baby AT components by 90 degrees so that any cards inserted into the bus architectures would not cover the processor and prevent proper cooling. Figure 1-34 shows an ATX motherboard.

Figure 1-34 also highlights some of the common components on the ATX board. Notice, for instance, slot 1, where a Pentium II chip can be inserted. Newer versions of the ATX motherboard use a ZIF socket to house the processor. Notice also, in the top-right corner, the BIOS chip with a white label on top of it. At the top of the figure, you can identify the EISA and PCI slots, and located in the center of the board is an AGP slot. The hard drive controllers are located on the left side beside the three slots that hold the DIMM memory.

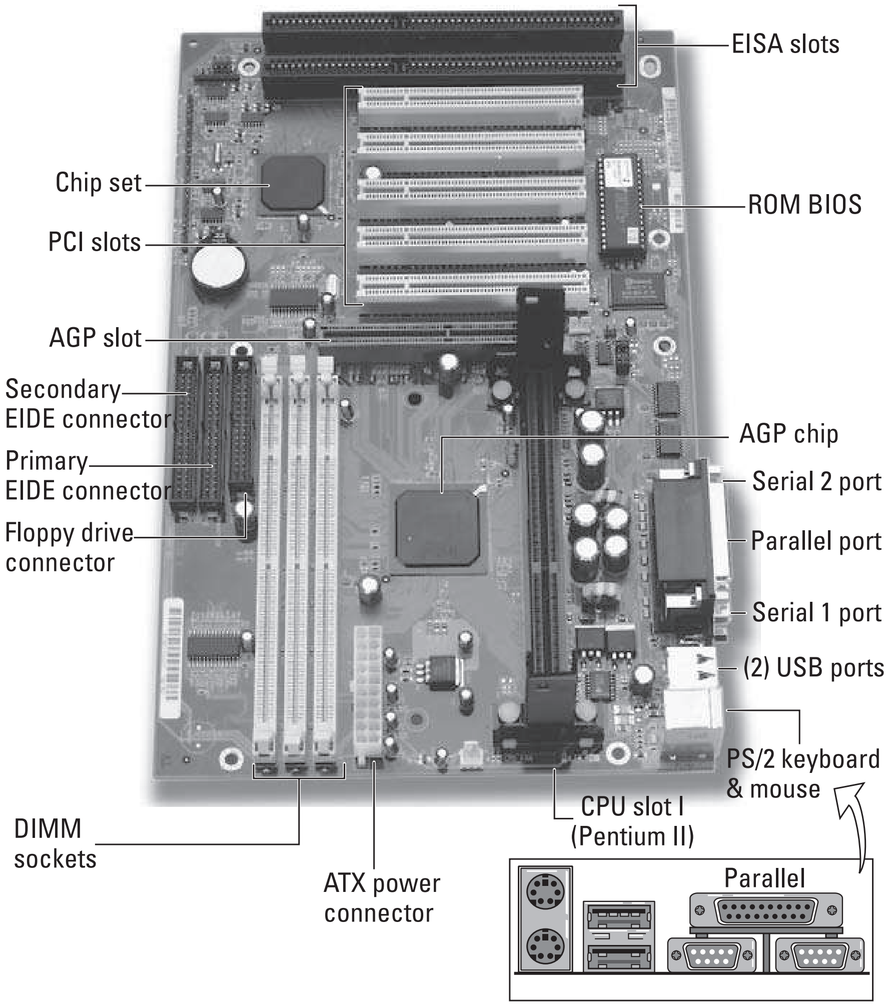

Figure 1-34: The ATX motherboard is very popular in today's systems.

MicroATX and FlexATX

Smaller versions of the ATX motherboard, known as MicroATX and FlexATX, have been developed. The *MicroATX* motherboard form factor is 9.6 inches by 9.6 inches and can fit in a MicroATX case or full ATX case. The *FlexATX* is smaller than the MicroATX (9 inches by 7.5 inches) and fits in an ATX and MicroATX case. FlexATX is not as popular because the size of the motherboard limits how much you can expand on the system. Figure 1-35 shows a MicroATX board.

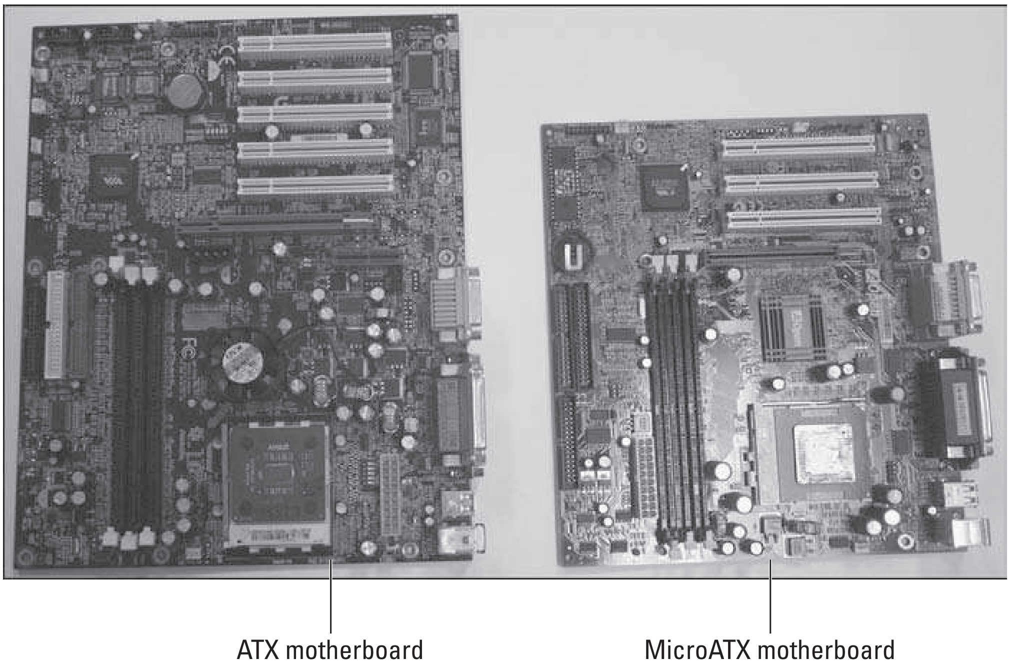

Figure 1-35: Comparing the size of the ATX motherboard (left) with the MicroATX motherboard (right).

The important point to make here is that when you purchase a motherboard, you must ensure that the motherboard you purchase fits the case you have. For example, if you have an ATX case, you now know that an ATX or MicroATX motherboard can fit in that case.

ON THE CD

Lab 1-3 will help you summarize distinguishing features of popular motherboard form factors. Lab 1-3 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Understanding Bus Architectures

The motherboard has a number of expansion slots that can expand the computer's capabilities. When the system is first purchased, a computer has only so many capabilities — the nice thing is that you can expand on those capabilities by purchasing cards to add to the expansion slots, or *bus architecture*.

Expansion slots expand on what the computer can do. The problem is that there are different types of expansion slots in the system, so when you go to purchase that sound card or network card, you have to make sure that you purchase the right type. In the following sections, I show you the different types of expansion slots and compare their characteristics.

Another term for the expansion slots is *bus architectures.* A number of different bus architectures have been developed over time. It is important to identify the differences between each of these architectures and also to know which ones are more popular today.

ISA

The *Industry Standard Architecture (ISA)* was the first major expansion bus architecture. It was originally developed as an 8-bit architecture and then evolved into a 16-bit architecture. The ISA bus architecture has a speed of 8 MHz, which is extremely slow by today's standards. Figure 1-36 shows two 16-bit ISA slots — note that the ISA slots are the black slots in the system.

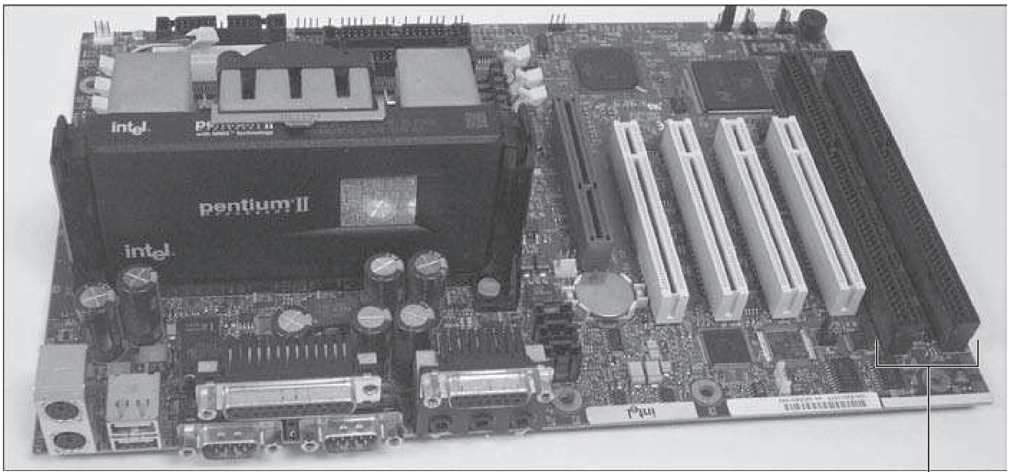

Figure 1-36: Identifying ISA slots on the system.

One of the reasons why you still see 16-bit ISA slots in some earlier Pentium or Pentium II systems is because companies typically had a number of ISA network cards in the office from previous systems. When the company upgraded to the Pentium or Pentium II it was nice that they did not have to purchase new network cards because earlier Pentiums had ISA slots. Most systems today no longer have ISA slots. Figure 1-37 shows a 16-bit ISA network card.

MCA

One of the major downfalls of the ISA bus architecture is its performance. It runs at only 8 MHz, and it is only a 16-bit architecture — that was fine years ago, but everything evolves, and new and improved standards arise.

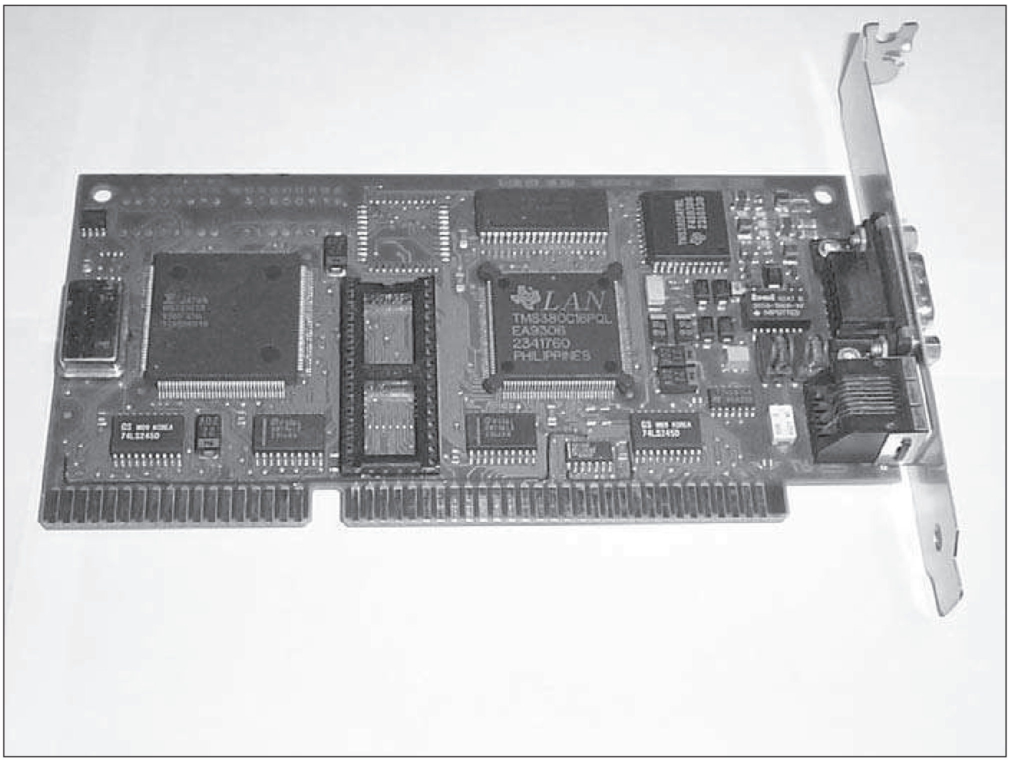

Figure 1-37: Looking at a 16-bit ISA network card.

The *Micro Channel Architecture (MCA),* which was developed by IBM, is a 32-bit architecture. The MCA architecture runs at 10 MHz and is not compatible with ISA. You usually find MCA slots in high-end IBM machines, such as those that might be used as a server.

Remember, ISA is an 8-bit or 16-bit technology that runs at 8 MHz. MCA transfers information in 32-bit chunks and runs at 10 MHz.

With MCA, IBM came up with a feature called *bus mastering.* Bus mastering works like this: Devices in the bus don’t have to send information through the CPU if they want to talk to one another — they just send the information directly. This takes some of the workload off the processor and allows it to perform other tasks. Bus mastering became an important feature in future bus architectures. Figure 1-38 shows an MCA card.

EISA

In 1988, the industry standard for expansion cards was still ISA, but bus architectures had already been created that performed better. So a number of companies got together with the goal of extending ISA while maintaining backward compatibility so that companies could use their existing ISA cards.

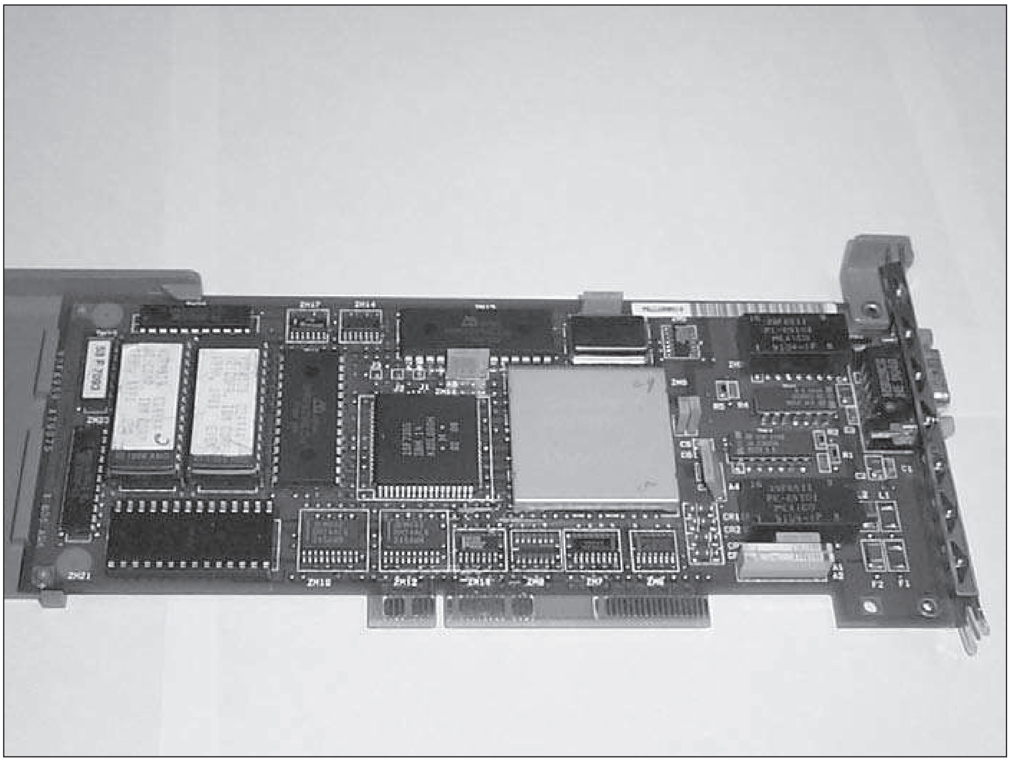

Figure 1-38: Looking at an MCA network card.

As a result, the *Extended Industry Standard Architecture (EISA)* was developed as a 16- and 32-bit architecture. The big advantage to EISA is that it maintains support for the ISA cards that some companies already have in large quantities, and it also supports 32-bit EISA cards. EISA also included the major advancement in expansion bus technology that MCA created, known as bus mastering. Because both ISA and EISA cards fit into the same slot, they keep the same speed of 8 MHz.

The bus architecture holds both 16- and 32-bit cards because the EISA slots have two levels. The EISA cards have very deep edge connectors that fill the two levels (32-bit) of the slot, but ISA cards only fill the top level (16-bit). Figure 1-39 shows an EISA slot and the two different levels in the slot: one level for the ISA card to fill and the other level for the EISA card to fill.

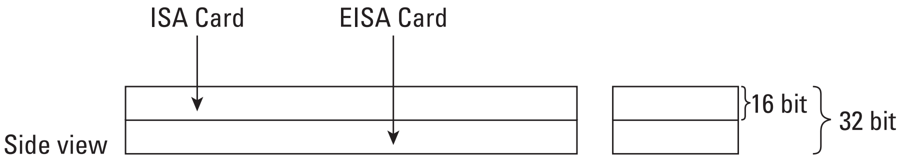

Figure 1-39: Looking at how the EISA slot is organized.

VESA

In 1992, the *Video Electronics Standard Association (VESA)* developed a bus architecture that outperformed ISA. VESA is a 32-bit architecture that supports bus mastering and runs at the same speed as the processor, which, when VESA was created, was around 25 to 33 MHz. Because the bus runs at the speed of the processor, developers called this *VESA local bus,* or *VLB.* VESA slots are typically used for video cards.

Remember, EISA is an extension on ISA and is a 16-bit or 32-bit technology. For backward compatibility, EISA runs at 8 MHz. VESA is a 32-bit architecture that runs at the processor's speed. It is generally used for video adapters.

VESA slots are extremely easy to identify because they are tan and act as an extension to the ISA slot. You will notice the black ISA slots and then right beside them may be a tan slot. The VESA card fills the entire ISA slot and the additional extension to make the full 32-bit path for VESA. This allows an ISA card to be inserted into the slot for backward compatibility or, with the extension slot, the VESA slot can hold a VESA card. Figure 1-40 shows a VESA slot.

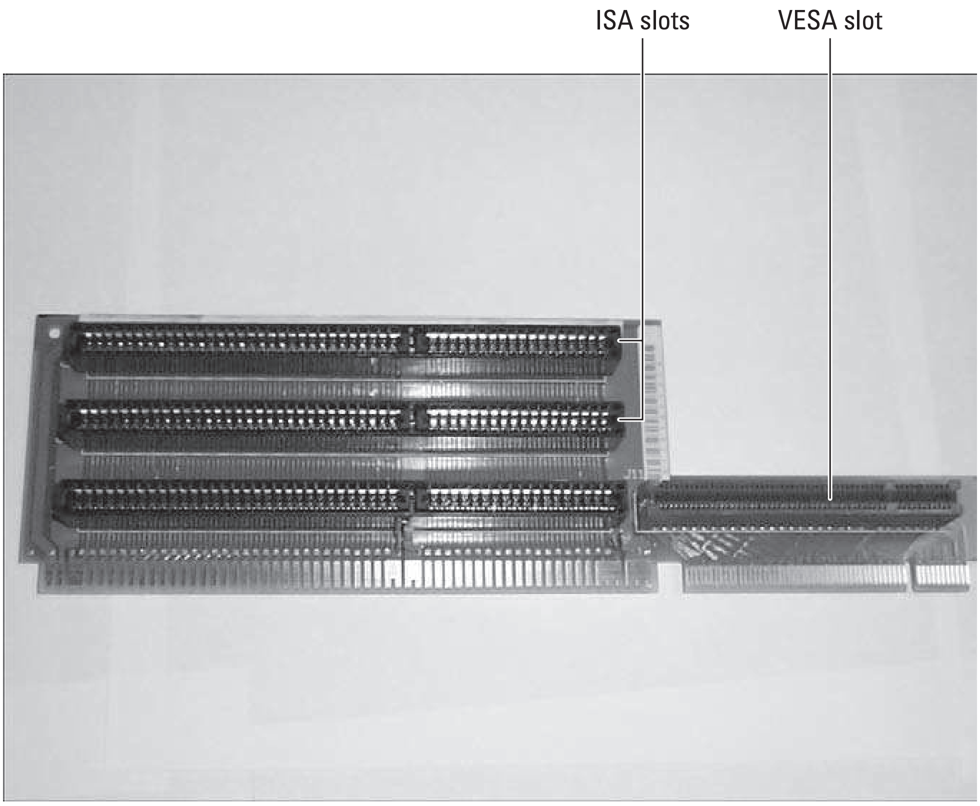

Figure 1-40: Looking at a VESA slot, which is an extension of the ISA slot.

PCI

Peripheral Component Interconnect (PCI) is one of the newer bus architectures to hit the market. PCI has two flavors: 32-bit cards and 64-bit cards. When Pentium systems hit the market, their motherboards featured both ISA/EISA slots and PCI slots. If you want to buy a new card today, you would most likely buy a PCI device for one of the PCI slots in your system.

The 32-bit version of PCI has a speed of 33 MHz, while the 64-bit version of PCI runs at 66 MHz. PCI also supports bus mastering. One of the other major benefits of PCI is that it is a Plug and Play architecture. If you are running a Plug and Play operating system like Windows 2000 or XP and your computer has a Plug and Play BIOS, then the system resources like IRQs and I/O addresses are dynamically assigned for PCI components.

PCI slots are easily identified on the motherboard as the small white slots, usually alongside the AGP slot. Figure 1-41 identifies PCI slots on a motherboard.

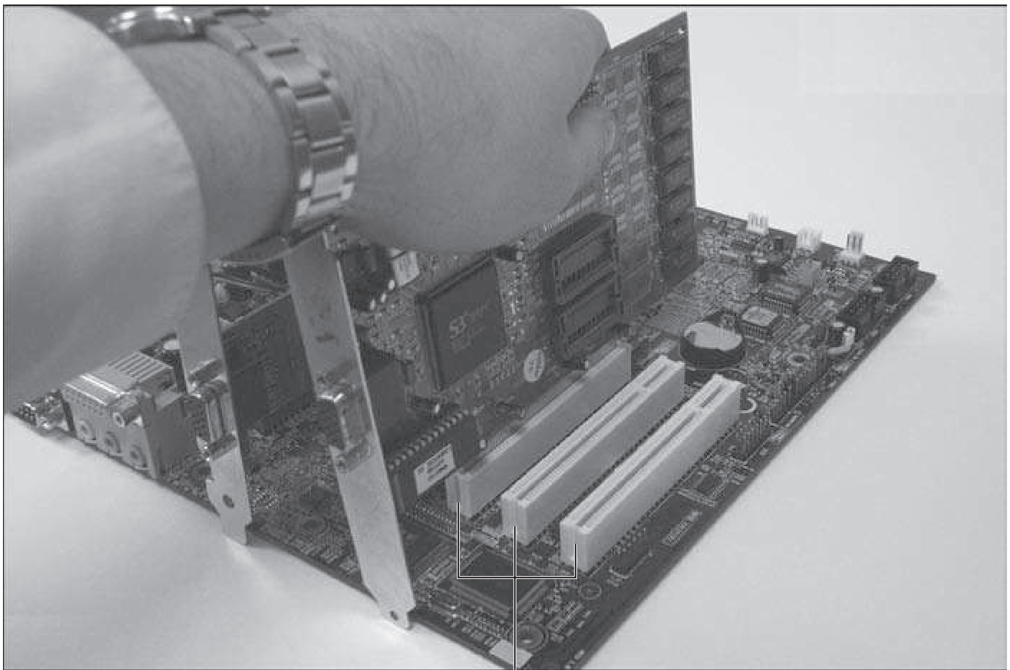

Figure 1-41: Installing a card into a PCI slot located on the motherboard.

PCMCIA

Personal Computer Memory Card Industry Association (PCMCIA) is a unique type of expansion bus architecture because of its small size. PCMCIA is popular in laptop computers. How are you going to get a big network card like the one that is used in a desktop computer into a little laptop to add network support? The answer is that you can't; you have to purchase a PCMCIA network card for the laptop to add network support. *PCMCIA cards,* also known as *PC Cards*, are a little bit larger than a credit card and can fit into your back pocket (though I don't suggest that you put one there). Figure 1-42 shows a PCMCIA network card.

Figure 1-42: Looking at a PCMCIA network card.

PCMCIA (say that five times fast!) is a 16-bit architecture that runs at 33 MHz and supports Plug and Play as well. Not only is PCMCIA a Plug and Play technology, but it is also a hot swappable technology. Hot swappable means that you can insert and remove PCMCIA cards without shutting down the system first. PCMCIA has three different types of slots named type 1, type 2, and type 3.

Table 1-1 shows the different PCMCIA slot types and the types of devices you can find in the different types of slots.

Table 1-1 Types of PCMCIA Slots

Slot Name	*Thickness*	*Types of Devices*
Type 1	3.3 mm	Memory cards
Type 2	5.0 mm	Modems/Network cards
Type 3	10.5 mm	Removable drives

Type 1 cards were originally used to add memory to laptop computers or personal computers. This is where the "personal computer memory card" part of the PCMCIA name comes from.

FOR THE EXAM

Remember, PCI is a 32- or 64-bit technology, runs at 33 MHz, and supports Plug and Play. PCMCIA is the expansion bus architecture used by laptop computers and is a 16-bit architecture that runs at 33 MHz.

AGP

Advanced Graphics Port (AGP) has been around since the Pentium II processor appeared in 1997. It's a 32-bit bus architecture that runs at 66 MHz — which is twice the speed of the PCI bus. Today's motherboards have one AGP slot to hold an AGP video card. The performance gain from the AGP port not only comes from the increase in speed, but also from the fact that the AGP bus has a direct path to the processor so that information travels quickly from the processor to the AGP card. Figure 1-43 shows an AGP slot beside some PCI slots.

AGP can run in different modes, and the different modes dictate the speed of the bus. 1x mode runs at 66 MHz (266 MBps), 2x runs at 133 MHz (533 MBps), 4x runs at 266 MHz (1.07 GBps), and 8x runs at 533 MHz (2.2 GBps)!

PCI-X

A fairly new bus architecture is the PCI-X bus architecture. Because PCI-X uses the same connector style as PCI, it is 100-percent compatible with PCI in the sense that it can hold PCI cards. So, a motherboard that has PCI-X slots can also house older PCI cards — that is a great feature!

Like PCI, PCI-X is a 32-bit and 64-bit bus architecture and is available in four different speeds. PCI-X runs at speeds of 66 MHz, 133 MHz, 266 MHz, and 533 MHz.

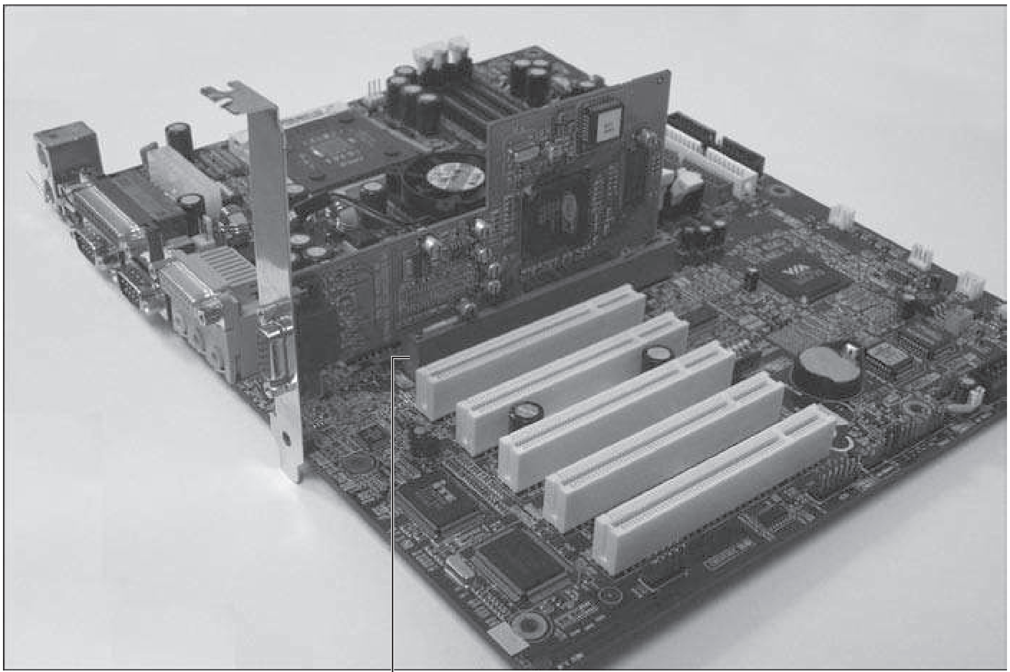

Figure 1-43: Looking at an AGP card in an AGP slot.

PCI Express

While PCI-X is compatible with PCI by being able to hold PCI cards and also sending data in parallel (multiple bits at one time), the PCI Express bus architecture takes a totally different approach. PCI Express is a serial bus that does not support existing PCI cards. The PCI Express slot, shown in Figure 1-44, is the smaller black slot and is much smaller than a normal PCI slot, so it can't possibly house a PCI card.

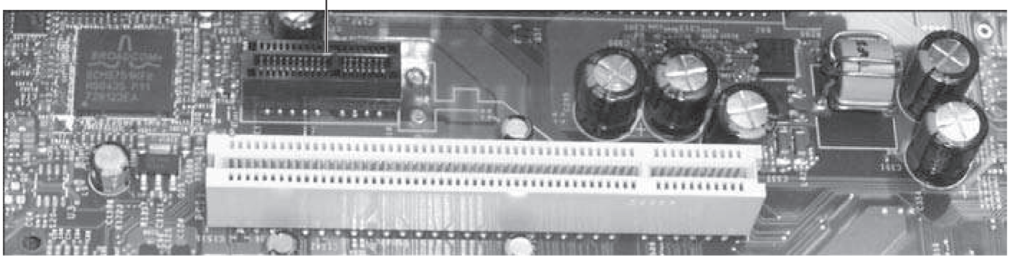

Figure 1-44: Identifying a PCI Express slot on the mother-board.

PCI Express uses data lanes to transfer the information within the bus architecture. A data lane delivers an amazing transfer rate of 250 MBps per lane. PCI Express has different implementations, with each implementation having a different number of lanes identified by a multiplier. For example PCI Express with only one lane is known as x1 while a PCI Express bus with eight lanes is known as x8. The implementation of lanes allows PCI Express to reach fast transfer rates by implementing additional lanes. For example, current graphics cards for PCI Express have 16 lanes which provide a transfer rate of 4 GBps (16 x 250 MBps) — which is twice the rate of AGP 8x, which runs at 2 GBps).

There are currently systems with PCI Express at x1, x2, x4, x8, x16, and x32. The PCI Express slot gets bigger with each multiplier — for example, Figure 1-44 is displaying a PCI Express x1 slot which is the black slot only about one inch in length.

AMR and CNR

Audio/Modem Riser (AMR) is a newer bus architecture that adds a modem and audio card to the system. AMR allows the two components to be incorporated into a single card to reduce cost. Figure 1-45 shows an AMR slot on a motherboard.

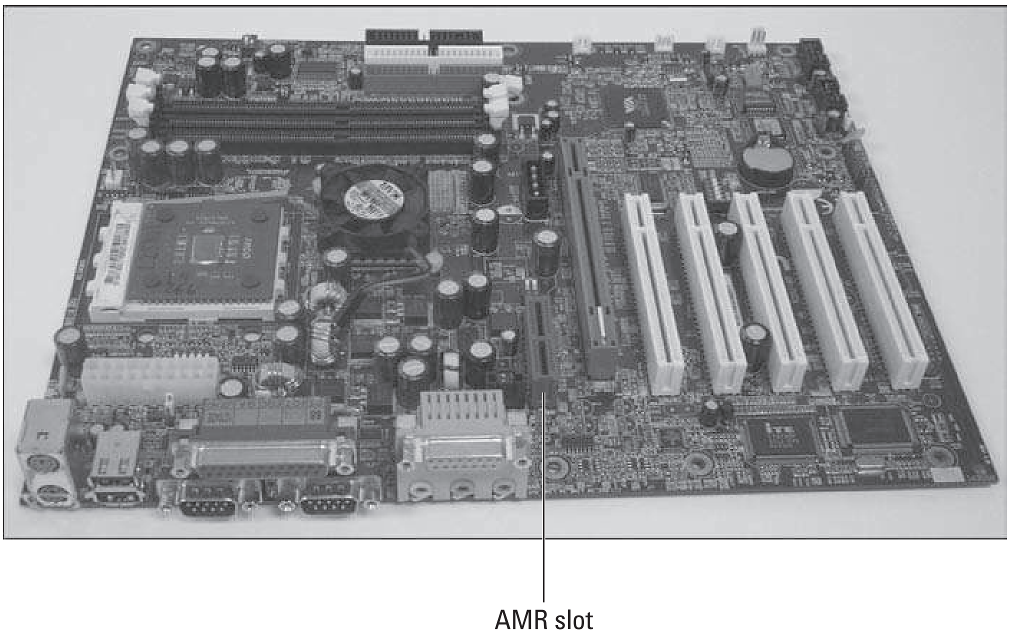

Figure 1-45: An AMR slot on a motherboard.

Communication and Network Riser (CNR) is another new architecture that is used to implement LAN, audio, and modem functionality all in one.

As far as the "real world" and the exam are concerned, you need to be extremely strong in the area of bus architectures. A big part of servicing computers is installing network cards, sound cards, and video cards — these components come as ISA, PCI, or AGP cards today. You need to know how to look at a system and say, "We are going to buy a PCI network card for this system."

Lab 1-4 will help you identify the different performance characteristics of each of the standard bus architectures. Lab 1-4 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Performance Considerations

When you want to improve a motherboard's performance, one of the first things you should do is check the speed of the motherboard. For example, some systems today have 400 MHz motherboards, and some have 533 MHz motherboards. To find out the speed of the motherboard check the documentation for the board. You can get a performance increase from a faster motherboard.

Another performance consideration occurs when you add expansion cards to the system. You should first evaluate what expansion slots are free and then purchase a card that will give you the best performance. For example, if you need to buy a network card for your computer, you start by looking at what expansion slots are free. If only two slots are free, an ISA slot and a PCI slot, then your choices are limited to an ISA network card or a PCI network card. Because PCI outperforms ISA, however, you would be better off purchasing a PCI network card.

You can also get a performance increase from motherboards that have more cache memory. Look at the motherboard to see if there is a place to install some level 2 (L2) cache. L2 cache can dramatically increase performance because it is generally closer to the processor than the system memory (RAM) and is a faster type of memory than the system memory. Bottom line: the more cache memory you install, the better the motherboard's performance.

Getting an A+

This chapter introduces you to a number of key components of the motherboard and different motherboard form factors. The following is a list of the key points to remember when dealing with motherboards:

- ✦ The *motherboard* (or *system board*) is the computer component that interconnects all other components.
- ✦ Serial (COM) ports come in two flavors: DB9-male and DB25-male. Parallel ports come in only a DB25-female port.
- ✦ There are two main types of cache memory: L1 and L2 cache. L1 cache memory is integrated into the processor, while L2 cache is contained outside the processor but in the processor casing or on the motherboard.
- ✦ IDE supports two devices in the IDE chain, whereas EIDE has two channels with two devices in each channel (a total of four devices).
- ✦ There are a number of major motherboard form factors: Full AT, Baby AT, NLX, and ATX, to name a few. Motherboard form factors differ in the size of the board and the layout of the components stored on the board.
- ✦ You may add components, such as a sound card or network card, to the computer by inserting an expansion card into one of the expansion slots in the system.
- ✦ ISA was the popular bus architecture for years, but because of its limitations (16-bit architecture and a speed of 8 MHz), it has been replaced by the PCI bus architecture. PCI is a 32-bit/64-bit architecture with a speed of 33 MHz.
- ✦ AGP is the common bus architecture used to insert a video card in today's systems.
- ✦ You may increase the performance of the system by using a faster motherboard or by purchasing better-performing expansion cards.