Chapter 3: What to Remember about Memory

Exam Objectives

- **Understanding memory terminology**
- **Identifying the types of RAM**
- **Understanding the types of DRAM**
- **Working with memory modules**
- **Identifying parity and non-parity memory**

Finding out how much memory a computer has is one popular way to measure the computer's power and capabilities. Think about it: If someone asked you what kind of computer you have, what would you say? Probably something like, "I have a Pentium 4 with 1024MB of RAM." But why do we measure the power of a computer based on the amount of memory it has?

In this chapter, you discover the purpose of memory and some of the different types of memory that are found in computers today. This chapter also discusses issues that affect the installation of memory in personal computers and laptops.

Understanding the Types of Memory

This section outlines different types of computer memory. The term *memory* refers to anything that stores information either permanently or temporarily. Computers have two different flavors of memory, ROM and RAM. From an exam perspective, make sure you fully understand the different types of memory and their uses because about 10% of the A+ exam focuses on memory.

Remembering the purpose of memory

Before we look at the different types of memory, let's first ensure that you understand the purpose of memory. We can compare memory to your desk at home or in the office. Whether sitting at your home or office desk (working on a proposal or preparing for your A+ Certification exam), chances are your desk is covered with documents, books, and papers. This desk is your work area, and its size dictates how many documents you can work on at any given time.

System memory works the same way. You have documents and applications stored on the hard drive. When you want to work on these documents, you open them and place them in the computer's work area. The work area (or desk space) for a computer is system memory. When you want to work with any application or document, the computer must retrieve that information from the hard drive and execute it from memory.

Assume, for instance, that your computer has 512MB of memory (not a lot in this day and age). You start up your system, which is running Windows XP, and decide to run Microsoft Word and Adobe Photoshop at the same time. Assume that you have opened two very large files in each application. Assume further that you are using 480MB of precious memory at this point — a few MBs for the operating system to load, and a few for each running application. As you can see, your memory usage adds up quickly.

In this scenario, you have already used 480MB of memory, which leaves 32MB of memory remaining. Assume that you are about to open up a Photoshop document and copy and paste information from one file to another. To put it simply, you are running out of desk space. You can solve the problem in one of two ways: You can either do less work (in other words, work on one application at a time — although this solution would not serve business users very well because they often need to run multiple applications simultaneously), or you can get a bigger desk, which in computer terminology means *installing more RAM.* When you install more RAM, you have a bigger desk to work on.

Now that you understand the general purpose of memory, let's dive into the different types of memory. There are many different types of memory you are required to know for the A+ exam that are outlined in the next sections.

Read-Only Memory (ROM)

Read-Only Memory (ROM) is a type of memory that you cannot write to. Information is written to ROM chips by the manufacturer, and this information cannot be changed. In the past, if ROM information needed to be updated, you had to remove the original ROM chip and replace it with an updated ROM chip from the manufacturer. Today you can update the ROM by running a special software program downloaded from the manufacturer's Web site, which means that you don't really have a ROM chip — you have an EEPROM (more on EEPROM in a bit).

TECHNICAL STUFF

Software written to a ROM chip is called *firmware*.

One of the major uses for ROM is storing the system *BIOS (Basic Input-Output System),* which contains Power-On Self-T*est (POST)* routines and other routines that initiate the loading of the operating system. The BIOS also contains the low-level code that allows the system to communicate with hardware devices.

For the exam, you need to know that the POST is part of the BIOS code stored in ROM. The POST contains routines that initiate the loading of the operating system and routines that make possible the communication between hardware devices.

EPROM

Erasable Programmable Read-Only Memory (EPROM) is a type of memory that normally cannot be written to because it is a variation of ROM. An EPROM chip is a special ROM chip that the manufacturer can reprogram by using a special programming device that uses ultraviolet light.

EEPROM

A new implementation of ROM is called *Electrically Erasable Programmable ROM (EEPROM),* or *flash ROM.* The manufacturer writes the software instructions into the ROM chip, but you can update these instructions by running a special software setup program provided by the manufacturer. The software setup program is usually provided to you through the manufacturer's Web site.

For the exam, remember that EEPROM, better known as flash ROM, is a ROM chip that can be rewritten with special EEPROM update software provided by the manufacturer of the chip.

EEPROM has become the typical way to update your BIOS. BIOS code is designed to work with certain hardware. As hardware improves, you need to update your BIOS code so that your system is aware of these hardware improvements. Therefore, the manufacturer places BIOS updates on its Web site for computer users running its particular BIOS to download. You just have to download the BIOS update program and then run the BIOS update on your system. The update rewrites the BIOS instructions, making the computer "more aware" of today's hardware.

Random Access Memory (RAM)

Of the two flavors of memory (ROM and RAM), RAM is probably the more fundamental. ROM is permanent memory, or permanent storage of information. As the computer's primary working memory, *RAM,* or *Random Access Memory,* stores information temporarily. RAM is volatile, meaning that it needs constant electrical current to maintain the information that resides in its chips. If the electrical current is lost, the contents of RAM are erased. When the computer is powered off, all the contents of RAM are flushed out.

The following sections discuss the different types of RAM. On the exams, you can expect a few questions about the different types of memory, so be sure that you are familiar with these different types of RAM.

DRAM

Dynamic RAM (DRAM) is probably the most popular type of memory today and the one that you are most often going to upgrade. When someone says to you, "I have 1024MB of Dynamic RAM," he or she is talking about DRAM.

Dynamic RAM gets its name from the fact that the information stored in DRAM needs to be constantly refreshed. Refreshing involves reading the bits of data stored in DRAM and then rewriting the same information back. DRAM is single ported — meaning that you can read and write to the memory but not at the same time.

Older implementations of RAM measured the memory's performance based on the time it took the CPU to access that data. The measurement used to determine the speed of memory is *nanoseconds (ns)* — one nanosecond equals a billionth of a second. If you have memory that is 50 ns, and your best friend has memory that is 70 ns, your memory is presumably faster. Your CPU receives the information from memory after waiting only 50 billionths of a second, whereas your best friend's CPU waits 70 billionths of a second.

REMEMBER

The lower the number of nanoseconds, the better the performance.

The speed of older DRAM ranges from 60 ns to 80 ns. Today's implementations of DRAM measure the speed of memory in megahertz (MHz) — typically matching the motherboard speed. For example, my Pentium II system uses 100 MHz memory because it runs on a 100 MHz motherboard.

For more information on the types of DRAM, see the section, "Identifying the Types of DRAM," later in this chapter.

SRAM

Static RAM (SRAM) — so-called because the information held in its memory cells doesn't need to be refreshed — requires less overhead than DRAM to maintain the information stored in memory.

With speeds running from 10 ns to 20 ns, SRAM is much faster than DRAM. Because SRAM is faster memory than DRAM it is also more expensive, which is why people add DRAM to their systems more often than they add SRAM.

SRAM is typically used for cache memory. *Cache memory* stores frequently used data and program code after it is read from slower DRAM. Think of cache memory as a bucket that sits beside the CPU and stores frequently used information. After the system has searched through DRAM once for specific information, it can store that information in the bucket for easy access later. The next time the data is requested, it is read from cache instead of from system memory.

Because cache memory is much faster than DRAM, the CPU first tries to retrieve the information from cache, specifically L1 cache first and then L2 cache. If the information is not located in cache, the system then tries to retrieve the information from memory. If the information is not located in system memory, it then is retrieved from disk. Attempting to retrieve the requested information from cache first reduces wait time if the information actually resides there because of how fast cache is compared to DRAM.

Cache memory (SRAM) stores frequently used data and program code. Because cache memory is faster than DRAM, retrieving information from cache means that the processor does not have to wait for the slower DRAM, thus enhancing system performance.

CMOS RAM

The *Complementary Metal-Oxide Semiconductor (CMOS)* is the area where the computer stores its configuration information, such as whether or not the computer has a floppy drive, the amount of memory installed, the date and time for the system, and the number and size of the hard drives that are installed. Think of the CMOS information as an inventory list for the majority of components that are installed on the computer. For more information on CMOS, see Book 2, Chapter 4.

CMOS is the computer's inventory list. It tells the computer which devices reside in the system. For example, the CMOS information lists your hard drive, floppy drive, and other information — such as the date and time for the system.

Where is the CMOS information stored? Is the CMOS information stored in the BIOS chip, or perhaps another ROM chip? The answer is neither. In fact, if the information were stored in a ROM chip, you wouldn't be able to go into the CMOS setup program and change the configuration. The CMOS configuration information is stored in a type of RAM called *CMOS RAM.*

CMOS RAM is a special, volatile RAM chip that stores the CMOS information. *Volatile* means that if power is lost, the information is wiped out. This could present a problem with regard to CMOS configuration because if the CMOS RAM is wiped out, the computer forgets its inventory information and has to relearn it. Thus, the computer has a small battery on the motherboard that maintains enough of a charge to avoid CMOS RAM losing its data.

For the exams, remember that CMOS information is stored in CMOS RAM, which is volatile memory that maintains its information by using a battery stored on the system board.

Shadow RAM

Part of the boot process involves copying some of the BIOS instructions from ROM up to RAM and then executing those instructions from RAM rather than from the ROM chip. Why? Because ROM is much slower than RAM, performance speed increases when executing the instructions from RAM instead of from ROM. The process in which a copy of the BIOS instructions is shadowed, or copied, to an area of memory called *shadow RAM* is called *shadowing*.

VRAM

Video RAM (VRAM) is dual-ported memory, meaning it can be read from and written to at the same time. DRAM is single-ported, which means that the memory can be written to and read from, but not simultaneously — only one direction at a time. VRAM, however, lets you do both simultaneously.

VRAM is most commonly used on video accelerator cards to store the values of the pixels on the screen for refresh purposes. VRAM is the favored memory for video because it outperforms the other memory types by being dual ported.

WRAM

Window RAM (WRAM), also known as *Window Accelerator Card RAM,* is a modification of VRAM and is also used for video display purposes. Like VRAM, WRAM is dual-ported memory but runs about 25 percent faster. In general, WRAM offers better performance than VRAM.

Identifying the Types of DRAM

Dynamic RAM (DRAM) is the most popular type of memory used in systems today. It is also the most popular type of memory that computer users add to their computers for the purpose of upgrading memory. Therefore, you must understand the different types of DRAM and what types of DRAM outperform others.

Standard DRAM

Memory is organized into rows and columns, like a spreadsheet. The information is stored in the different cells, or blocks, that are created by the intersection of these rows and columns. With *standard DRAM,* the CPU requests data by sending the address of the row and the address of the column for every block of data that needs to be read to the memory controller. The memory controller then fetches the information from that memory location. Figure 3-1 shows two memory cells that hold data that the CPU wants to have.

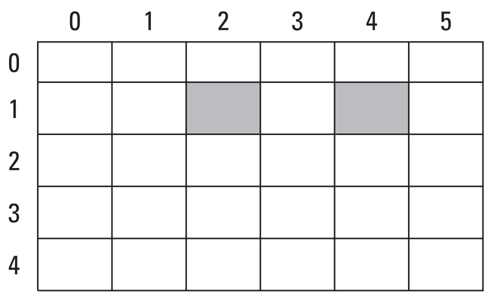

Figure 3-1: Looking at how data is accessed in memory.

To access the information shown in Figure 3-1, the CPU follows these basic steps to request information from standard DRAM:

1. In the first clock cycle, it sends the row address (1).
2. In the second clock cycle, it sends the column address (2).
3. On the third clock cycle, the memory controller reads the information (Address 1-2).
4. In the fourth clock cycle, the row address for the second memory cell is given (1).
5. In the fifth clock cycle, the column address for the second memory cell is given (4).
6. In the sixth clock cycle, the second memory cell is read (Address 1-4).

Fast page mode

Fast Page Mode (FPM) improves the performance of standard DRAM by not requiring a row address for each request to memory, assuming that the next block of data is on the same row (which in most cases will be true). The following list outlines the basic steps to access the same two blocks of data shown in Figure 3-1 via fast page mode memory:

1. In the first clock cycle, the CPU sends the row address (1).
2. In the second clock cycle, it sends the column address (2).
3. On the third clock cycle, the memory controller reads the information (Address 1-2).
4. In the fourth clock cycle, the column address is given (4).
5. In the fifth clock cycle, the second cell address is read (Address 1-4).

You can see in this example that it takes less time to read both blocks of data from memory with fast page mode DRAM. Therefore, FPM memory is a faster DRAM memory type than standard.

Extended data output

Extended Data Output (EDO) memory is about 10 to 15 percent faster than FPM memory and is usually found on 66 MHz motherboards. With EDO memory, the memory controller can read data from a memory block while listening for the next instruction. This capability increases performance because the memory controller doesn't have to wait for the next instruction after reading a block of memory; while it is reading one block of memory, it is receiving the next instruction. In contrast, with FPM DRAM, reading one memory block and listening for the next instruction are done in multiple steps.

Burst Extended Data Output

Burst Extended Data Output (BEDO) is a bursting-type technology. The word *burst* refers to the fact that when one memory address is requested and that address is retrieved, the system bursts into the next couple of blocks and reads those as well. The theory behind BEDO is that the system has already gone through the trouble of locating that block, and chances are that the next request will be for the next block, so why not take that information while the memory controller is already there? If that extra block is the next requested block from the CPU, the memory controller already has the data and can pass it to the CPU immediately.

BEDO is 50 percent faster than EDO. Because of lack of support from computer manufacturers, however, BEDO has not been used in many systems. It has been surpassed by SDRAM instead.

Synchronous DRAM

Synchronous DRAM (SDRAM) is memory synchronized to the system board speed. This synchronized speed means that the data stored in memory is refreshed at the system speed, and data is accessed in memory at the system speed as well.

SDRAM is one of the most popular types of DRAM found in later Pentium systems, such as the Pentium II. When you upgrade memory on your system, if you determine that you need SDRAM, you will then need to determine what speed SDRAM. Because you are running at the system speed, you must match the DRAM speed with the motherboard speed. Thus, if you have a 100 MHz motherboard, you need 100 MHz SDRAM. If you have a 133 MHz motherboard, you need 133 MHz SDRAM.

TIP

As mentioned, if you have a 100 MHz motherboard, you will purchase 100 MHz memory, typically labeled PC100. Be aware, however, that there is some flexibility when purchasing SDRAM. For example, I have a 100 MHz motherboard on an old Pentium II system. When I upgraded the DRAM on this system, I couldn't buy PC100 memory because PC133 (which is SDRAM that runs at 133 MHz) was the popular memory at that time. Not a problem! You can use faster memory than your motherboard speed as long as you are willing to accept that you have paid for memory that will not run to its full potential speed. In my example, the 133 MHz memory is only running at 100 MHz due to the speed of the motherboard.

Rambus DRAM

At the time that SDRAM was popular, there was a high-speed flavor of DRAM on the market called *Rambus DRAM (RDRAM),* which runs at speeds around 800 MHz! The RDRAM chips have a 16-bit internal bus width and are packaged together in a 184-pin, gold-plated memory module called a *Rambus Inline Memory Module (RIMM).* In order to take advantage of this type of memory, you need a motherboard and chipset that support RDRAM.

DDR

Double Data Rate (DDR) memory gets its name from the fact that it can transfer data twice during each clock cycle, whereas SDRAM can transfer data only once per clock cycle. DDR memory ships in 184-pin DIMM modules (see the section "DIMMs," later in this chapter) for desktop computers and 200-pin SO-DIMMs for laptop systems.

The speed of DDR memory is measured in MHz, like SDRAM is, and is labeled to indicate the speed. The labeling of DDR memory may look obscure at first because it also indicates the bandwidth by taking the speed and multiplying it by 8 bytes of data (64 bits). So if DDR memory is labeled PC1600, that label breaks down like this: If you divide the 1600 by 8 bytes, you get the speed of the memory; in this case, you're looking at 200 MHz memory. PC2700 runs at 333 MHz, while PC3200 runs at 400MHz. When you upgrade memory on systems that require DDR memory, you need to know the speed of the DDR memory.

DDR2

Improvements to DDR memory have already started with DDR2 memory. *DDR2* memory runs at speeds 400 MHz and higher, which is where DDR memory left off. DDR2 memory uses 240-pin memory modules and runs at 1.8 volts (as opposed to 2.5 volts for DDR memory). This results in less power consumption for more memory — which is great for laptop users.

Popular modules of DDR2 memory at the time of this writing are PC3200 (400 MHz), PC4200 (533 MHz), PC5300 (666 MHz), and PC6400 (800 MHz).

How Would You Like Your Chips Packaged?

Whether you're purchasing or installing RAM, understanding the different types of memory packages available is important. The following sections identify different memory packages used in desktop computers and laptop systems.

SIMMs

Single Inline Memory Modules (SIMMs) used to be one of the most popular types of memory modules, but they have been replaced by DIMMs (see the next section) in recent years. A *SIMM* is a card that holds a number of memory chips and has an edge connector containing a number of pins that make contact with the motherboard. This design makes it quite a bit easier to install memory than it was many years ago. In the past, you had to take a dual inline package (DIP) chip out of the system board and reinsert a new chip. Today, you purchase a card of chips (a SIMM) and install the SIMM into one of the SIMM sockets.

SIMMs come in two flavors, 30-pin and 72-pin, which describe the number of connectors that make contact with the motherboard. Before buying a SIMM to install in a computer, review the documentation for the computer or look at the system board to determine what size SIMM module you need. Figure 3-2 shows a 30-pin SIMM, a 72-pin SIMM, and a 168-pin *dual inline memory module (DIMM).*

The 30-pin SIMMs have an 8-bit data path, meaning they supply information in 8-bit blocks. When installing memory into a system, you *must* install enough SIMMs to fill a memory bank. A *memory bank* is the number of SIMMs it takes to fill the data path of the processor. For example, if you have a system with a 486 processor, the processor is a 32-bit processor. Therefore, the processor wants to deal with information in 32-bit chunks. When using 30-pin SIMMs, you need to install four of them at a time to fill a memory bank because each 30-pin SIMM only supplies 8 bits of data (8 bits × 4 SIMMs = 32-bit chunks).

The 72-pin SIMMs supply information in 32-bit chunks. Therefore, if you are installing 72-pin SIMMs on a system using a 32-bit 486 chip, you need just one SIMM to fill a memory bank and the data path. If you're installing 72-pin SIMMs in a Pentium system, you must install SIMMs in pairs because the Pentium data path is 64-bit; to fill a bank on these systems, you need two 32-bit modules (72-pin SIMMs).

FOR THE EXAM

For the exam, remember the data path of the SIMM modules. You should also know how many SIMMs it takes to fill a memory bank on different systems. Remember that a *memory bank* is the number of memory slots needed to fill the data path of the processor.

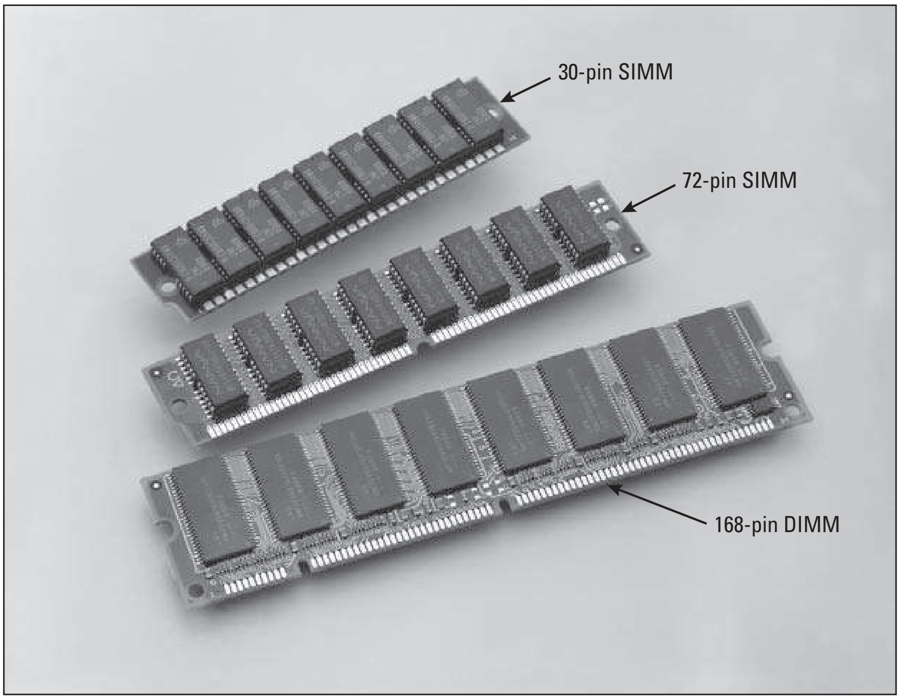

Figure 3-2: Looking at SIMM and DIMM memory modules.

In this day and age, you most likely will not see SIMMs in a system unless you are supporting older computers. If you see a system that uses SIMMs, it most likely conforms to the 72-pin format.

TIP

You can easily distinguish what size SIMM a system uses, even if you don't have the documentation for that system. The 72-pin SIMMs have a notch close to the center of the module. If there are SIMMs already installed in the system, you can take them out and examine them. They usually have a label with a 1 or a 72, representing the pin numbers, at either end of the module — so if you see a number 72, you know you have a 72-pin SIMM.

DIMMs

Dual Inline Memory Modules (DIMMs) are like SIMMs, only they supply information in 64-bit chunks. DIMMs use 168 pins on the module and are a little larger than the 72-pin SIMMs. (Refer to Figure 3-2.)

FOR THE EXAM

For the exam, remember that SIMMs come in 30-pin and 72-pin flavors, whereas DIMMs have 168 pins. DIMMs are also the most popular type of memory module that you will find in systems today.

Consider the memory bank issue again. Because the DIMM supplies data in 64-bit chunks, and the data path of a Pentium processor is 64-bit, you can install DIMMs singly in a Pentium system. On the other hand, you must install SIMMs in pairs in a Pentium system. Figure 3-3 shows what 72-pin SIMM and 168-pin DIMM sockets look like.

SODIMM

Small Outline Dual Inline Memory Modules (SODIMMs) are memory modules that are smaller than normal DIMMs and are used in laptops. A SODIMM comes in three different-sized modules: a 32-bit 72-pin module; a 64-bit 144-pin module (SDRAM); and a 64-bit 200-pin module (DDR). Figure 3-4 compares a SODIMM and a DIMM.

MicroDIMM

A *Micro Dual Inline Memory Module (MicroDIMM)* is another memory module that is used in laptop computers. The MicroDIMM is smaller than the SODIMM and comes in a 144-pin module for SDRAM and a 172-pin module for DDR memory.

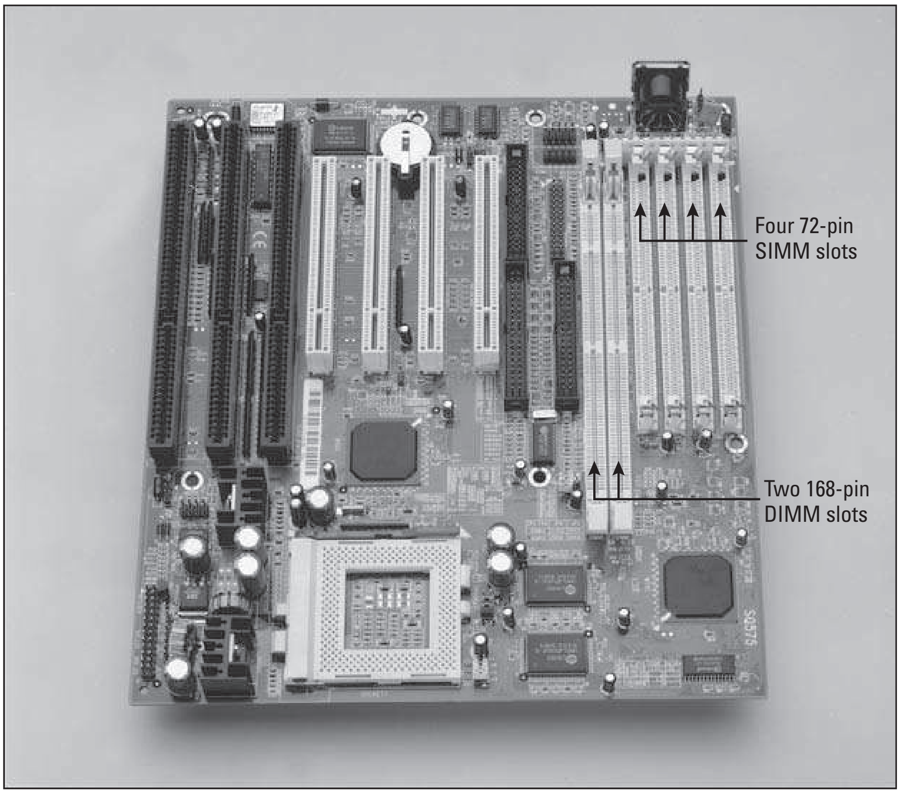

Figure 3-3: Looking at memory sockets on a motherboard.

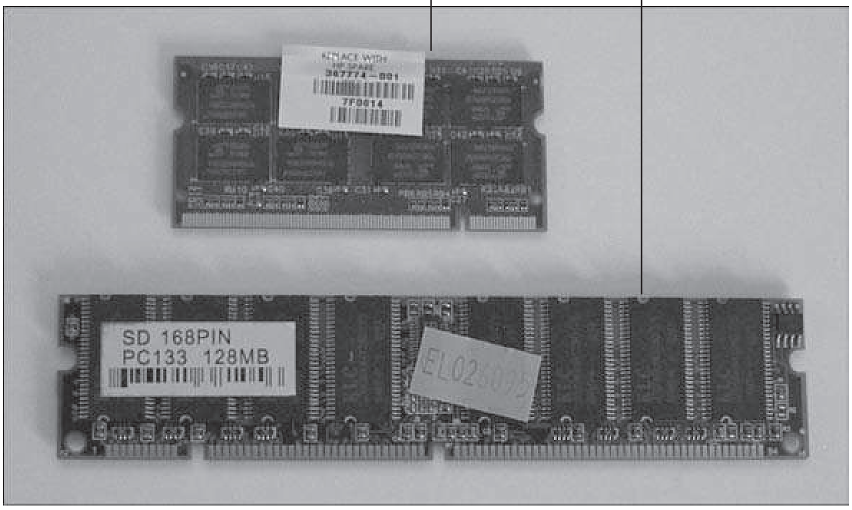

Figure 3-4: Comparing a SODIMM and a DIMM.

Understanding Error-Checking Memory

There are two primary types of error-checking memory that have been used in systems over the years. The following sections introduce you to these two types of error-checking memory — be sure to become familiar with them for the exam.

Parity versus non-parity

In this section you learn about parity versus non-parity memory. *Parity memory* is a type of error-checking memory, which is memory that verifies the information stored in memory is what is actually read from memory at a later time. *Non-parity memory* is simply memory that does not perform any kind of error checking to ensure that the data written to memory is what is actually read when it is retrieved. Let's look at how parity memory works!

There are two types of parity memory: *odd parity* and *even parity.* Both parity methods function the same way but differ in the sense of whether they look for an odd number of bits or an even number of bits. This discussion uses odd parity as the example.

With parity memory, for every byte (8 bits) of data written to memory, there is an additional 9th bit known as the parity bit. When storing information to memory, the number of the enabled data bits (bits set to 1) written to memory are added up.

With *odd parity,* if an even number of data bits are enabled, the parity bit is set to 1 (enabled) so that there is an odd number of enabled bits in total written to memory. If the result of all the enabled data bits is odd, the parity bit is set to 0 (disabled) so that the odd number of enabled bits is retained.

After the parity bit has been set, the byte of data and the parity bit are written to memory. Note that *even parity* works the same way, only it looks for an even number of enabled bits; if the number of enabled bits is odd, then the parity bit is enabled.

When the CPU requests data from memory, the data byte is retrieved along with the parity bit that was generated when the byte of information was stored in memory. The system looks at the data byte and calculates whether the parity bit stored in memory should be set to 1 or 0. It then compares the answer it has just generated with the value of the parity bit stored in memory. If the two match, the integrity of the information in memory is considered okay, the parity bit is stripped from the data byte, and the data is delivered to the CPU. If the two differ, you have a *parity error*, meaning that there is a problem with the integrity of the data stored in memory.

REMEMBER

Note that parity memory cannot correct the error; it just reports that an error exists.

FOR THE EXAM

For the exam, remember that parity memory has an extra bit (the parity bit) for every 8 bits of data. SIMMs with parity come in 9-bit (30-pin SIMM) or 36-bit (72-pin SIMM) flavors. Also, remember that a parity error indicates that there's something wrong with the integrity of data stored in memory.

ECC memory

Error-checking and correction (ECC) memory is memory that can detect data integrity problems the way that parity memory can, the difference being that ECC memory can recover from the error and attempt to fix the problem with the data being read.

Working with Cache Memory

Cache memory stores frequently used data and program code after it is read from slower DRAM. Cache memory is made up of SRAM, which is much faster than DRAM. The average speed of DRAM is 60 ns, whereas the average speed of SRAM is 20 ns. If at all possible, you want the CPU's request for information to be serviced by cache memory for a quicker response. To help service these responses, the system has two major levels of cache memory: L1 and L2 that are popular, and also an L3 cache that is making its way on systems today.

L1 cache

Level 1 cache, or *L1 cache,* is "internal cache" integrated into the CPU. This memory is typically a small amount of SRAM integrated into the processor's chip, giving the processor instant access to this memory with no wait time. *Wait time* is the amount of time it takes between the processor requesting information stored in memory and actually receiving the information.

Every processor before the Pentium processor has L1 cache integrated into the processor chip, but the amount of L1 cache can vary. For example, the 486 chips had 8K of L1 cache, whereas the early Pentium processors had 16K of L1 cache. Newer processors have doubled that amount to 32K of L1 cache.

L2 cache

Level 2 cache, also known as *L2 cache,* exists outside the CPU, usually on the motherboard or just outside the processor but in the processor casing. Therefore, some delay occurs when the processor accesses the information in L2 cache due to the distance between the processor and the L2 cache.

FOR THE EXAM

For the exam, remember that L1 cache is SRAM integrated into the processor's chip, whereas L2 cache is SRAM located outside the CPU, usually on the system board or in the casing of the processor.

One of the selling points of different processors is the amount of cache memory that comes with the processor. Many processors today typically have at least 32K of L1 cache and 512K, 1MB, or 2MB of L2 cache inside the casing of the processor. The more cache memory a system has, the bigger the bucket to store more frequently used information.

When the processor retrieves information, it first checks to see whether the information it needs is stored in L1 cache (because L1 cache has no wait time). If the processor does not find the information in L1 cache, it checks the L2 cache. If the information cannot be found in either L1 or L2 cache, the processor finally retrieves the information from RAM. Figure 3-5 shows the steps the processor takes to retrieve information.

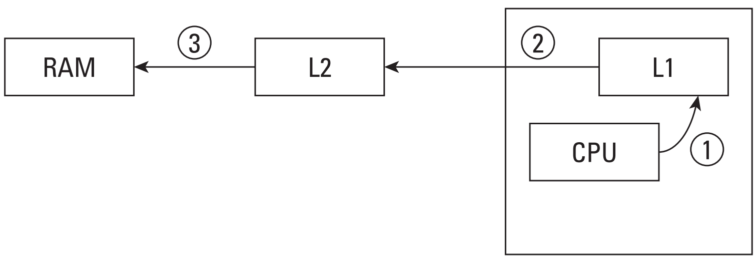

Figure 3-5: How the CPU retrieves information from memory.

L3 cache

Because processors today provide a small amount of L1 cache *and* a large amount of L2 cache, some people are now using the term *L3 cache* to identify cache that resides on the motherboard.

If you are confused by L3 cache I can see why — we have changed the terminology on you. In the past, L2 cache resided on the motherboard but now that processors include L2 cache on the CPU, the term for cache memory located on the motherboard is L3 cache!

Installing or Upgrading Memory

The discussion in this section focuses on issues related to memory upgrades. In general, upgrading memory is a simple task — assuming you purchase the proper type of memory for the upgrade. Factors that affect the proper type of memory are

- Type of memory (FPM, EDO, SDRAM, DDR)
- The speed of the memory
- Pin connector type
- Parity versus non-parity

The following sections discuss each of the factors that affect how you upgrade your system's memory.

Type of memory

The first thing you need to know to upgrade your computer's memory is which type of memory you need. You first need to figure out whether you need to install a SIMM, DIMM, or SODIMM. Pentium II and later desktop computers usually need a DIMM; laptops use a SODIMM.

After you have determined the memory module type, you need to determine the type of memory to install, such as SDRAM or DDR memory.

Speed

When you buy memory, you need to take into account the speed of the memory. Older memory, such as FPM or EDO memory, is typically measured in nanoseconds (ns) and ranges from 60 ns to 80 ns. With these types of memory, it is important not to mix speeds or the system will become unreliable. The speed of a SIMM is usually indicated on the chips themselves (displaying either a numeric value or a simple minus sign with a number). For example, a memory module running at 70 ns would show either "70" or "–7" on the chips.

The speed of newer memory types, such as DIMMs and SODIMMs, is measured in MHz. If you're buying SDRAM or DDR memory, make sure you verify that you get the correct memory speed. For example, I recently upgraded the memory in the laptop I'm using to write this book, and when I went to the store to buy a memory module, the first questions the in-store computer geeks asked were, "What type of memory?" and "What speed?" The system documentation can help you determine the speed of memory needed. If you don't have the documentation, be sure to look up the information on the manufacturer's Web site.

Connectors

Another important issue with regard to memory installation is with the metal used on the memory modules. You need to purchase memory modules that use pins plated with the same metal used in the memory socket on the motherboard. Memory modules use silver or gold plated pins. If the SIMM socket, for example, uses silver-plated connectors, the memory module you purchase must use silver-plated pins. If you mix metal types, you'll eventually have an unstable system.

WARNING!

Today's preferred memory modules are gold-plated DIMM modules. It is best not to assume you are using gold plated DIMMs, but to ensure you are installing the correct type of memory by checking the documentation for your system.

Parity versus non-parity

The final issue with regard to memory upgrades is whether the system uses parity or non-parity memory. This information can be determined by checking the documentation that came with the system or by checking the system summary in CMOS. If you can't find the information in CMOS or have misplaced the documentation, you can try to find the information on the Internet at the vendor's Web site. When you locate the information, use Table 3-1 to record the type of memory your system has so that you may refer to this when you perform a memory upgrade.

Table 3-1	Identifying Memory Used by Your System
Memory Type (SDRAM, DDR)	
Memory Speed (60 ns, 100MHz)	
Gold or Silver pins	
SIMM or DIMM	

Installing memory on desktop PCs

Now that you have purchased the correct type of memory — meaning the correct type at the correct speed with the correct number of pins — you are ready to install the memory!

Take off the computer's cover. You should see either SIMM sockets or DIMM sockets — maybe even both types on an older system.

WARNING!

Look at the sockets and determine whether the memory modules will sit diagonally or vertically. This step is very important. I have seen many people struggle to install memory because they didn't understand how to correctly place the modules in the sockets.

If the socket is on a diagonal, lightly place the memory module vertically and then lay it back diagonally; it should just snap in. If the socket is vertical, place the memory module at a 45-degree angle and lightly lay it back to the vertical position; it should snap in. Figure 3-6 shows the installation of memory modules.

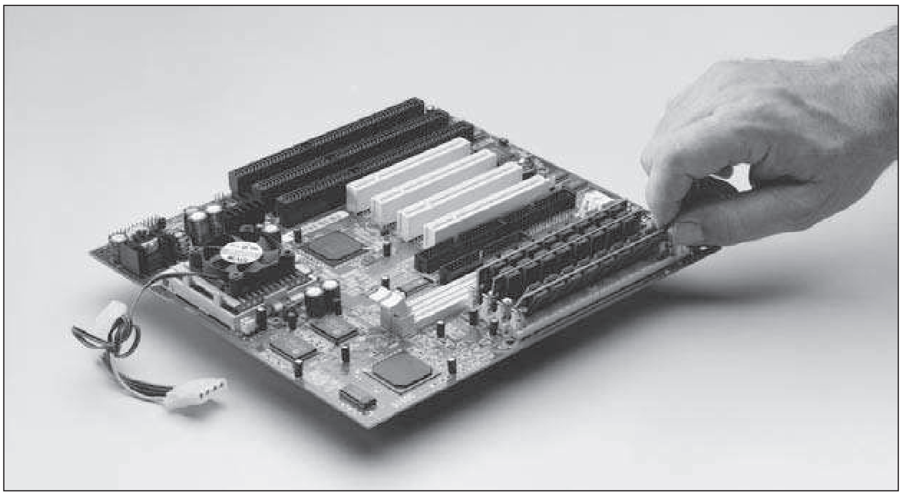

Figure 3-6: Installing memory modules on a system board.

When you install the memory module, make sure you line up Pin 1 on the memory module with Pin 1 in the socket. To locate Pin 1 on a SIMM, look for the cutout on the memory module (shown in Figure 3-7) and place it over the shoulder of the SIMM socket.

To locate Pin 1 on a DIMM, simply look at the memory module; the pins are labeled. The DIMM cannot be placed in backward because it's *keyed* (meaning that it has cutouts so that the memory can only be inserted one way). When installing the DIMMs, you will push down firmly on the module to place the DIMM deep in the socket.

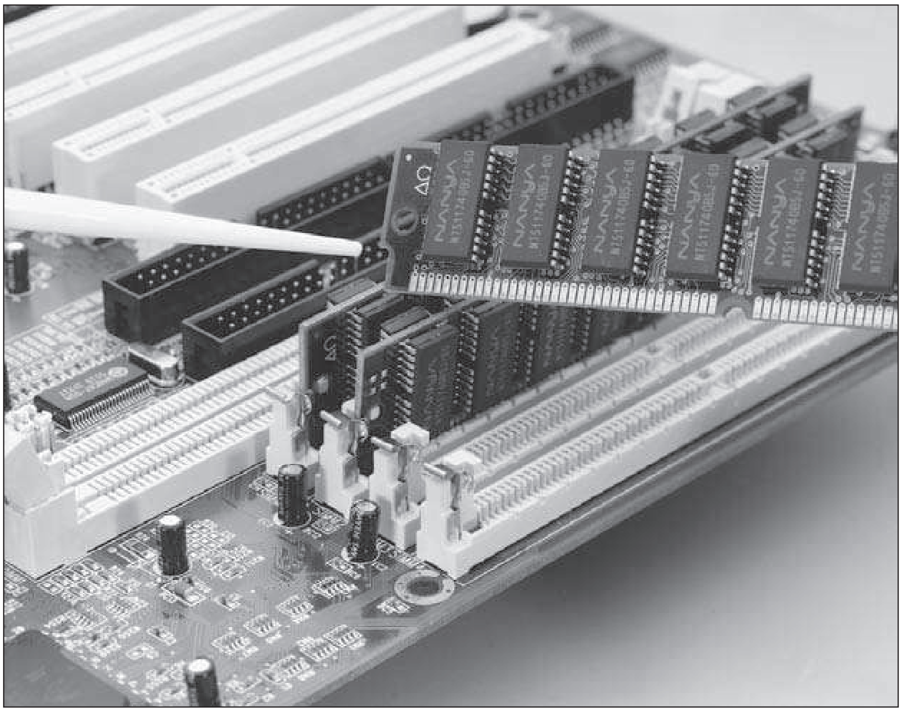

Figure 3-7: Locating Pin 1 on a SIMM memory module.

ON THE CD

Lab 3-1 and lab 3-2 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM. Lab 3-1 will allow you to practice finding out what type of memory you have, while Lab 3-2 gives you practice installing the memory on a desktop system.

Installing memory on laptop systems

Installing memory on today's laptop systems is just as easy as installing memory on a desktop PC. The first step is to make sure you buy the correct type of memory for the laptop — again, check the laptop documentation before heading to the store!

To install your newly purchased memory, flip the laptop over. You'll notice a door on the bottom of the laptop that can be removed (shown in Figure 3-8). This is where you add memory.

After the cover has been removed, you can insert the SODIMM into an empty slot by lightly placing the SODIMM into the memory slot on a slight angle and then clamping it back into place, as shown in Figure 3-9. If you are replacing memory, you will need to first remove the old memory module by pressing on the clips on the side and then lifting the memory out.

When you have the SODIMM locked in place, put the cover back on and you're ready to use your laptop with the new memory!

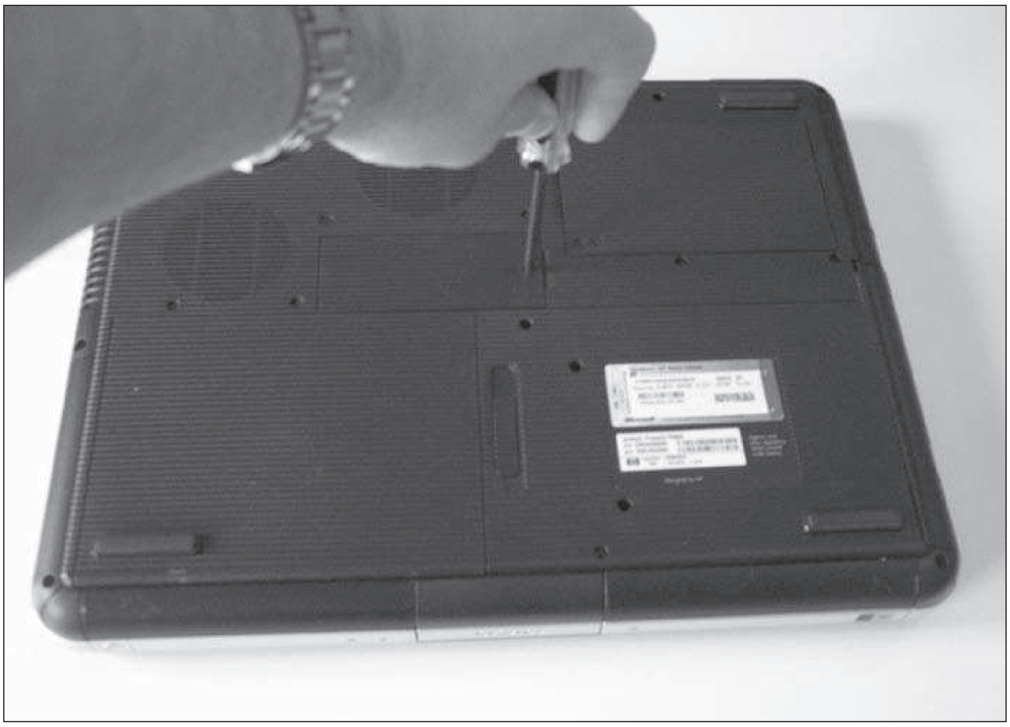

Figure 3-8: Removing the cover to add memory to a laptop.

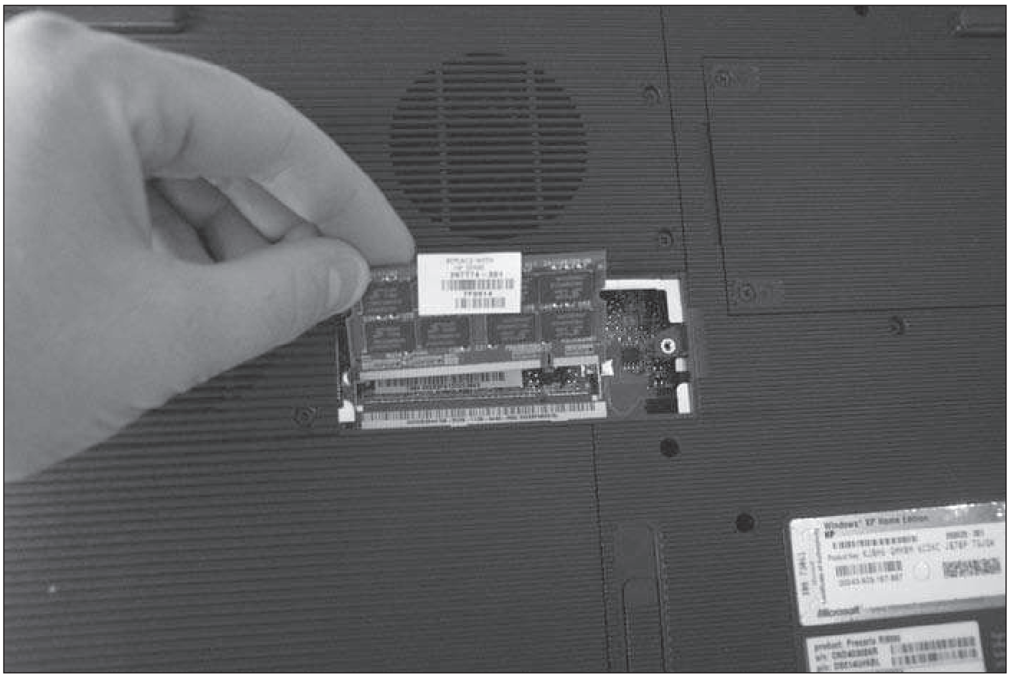

Figure 3-9: Inserting the SODIMM module into the memory slot.

ON THE CD

Lab 3-3 lets you practice finding out what type of memory is needed by a specific laptop. In this lab you will need to research the laptop on the Internet. Lab 3-3 can be found in the `Labs.pdf` file in the Author directory of the CD-ROM.

Getting an A+

This chapter provides an overview of the different types of memory and installation of memory. The following points are touched on:

- *Read-Only Memory (ROM)* is memory that can be read from but not written to, and the information stored there is permanent.
- *Random Access Memory (RAM)* is volatile memory that can be written to and read from; information stored there is flushed out when power is lost.
- *Dynamic RAM (DRAM)* is memory that needs constant refreshing of its memory cells; it is also the type of memory that is typically upgraded on systems.
- *Static RAM (SRAM)* is static memory, meaning that it does not need refreshing as often as DRAM does. SRAM is also faster than DRAM. SRAM is typically used for cache memory.
- *L1 cache* is cache memory integrated into the processor, whereas *L2 cache* is cache memory that exists outside the processor's chip.
- Memory is installed with *memory modules.* 72-pin SIMMs were popular in original Pentium systems, but 168-pin DIMMs are now the popular memory module for desktop systems. Laptops use SODIMMs as their memory module type.
- You need to know the *type of memory* your system uses before you upgrade your system. For example, you need to note whether your system uses SDRAM or DDR memory. Also be sure to note the speed of the memory.