


“The memory be green.”

—WILLIAM SHAKESPEARE, *HAMLET*,
ACT I, SCENE 2



In this chapter, you will learn how to

- Identify the different types of DRAM packaging
- Explain the varieties of RAM
- Select and install RAM
- Perform basic RAM troubleshooting

Whenever people come up to me and start professing their computer savvy, I ask them a few questions to see how much they really know. In case you and I ever meet and you decide you want to “talk tech” with me, I’ll tell you my first two questions just so you’ll be ready. Both involve *random access memory (RAM)*, the working memory for the CPU.

1. “How much RAM is in your computer?”
2. “What is RAM and why is it so important that every PC has some?”

Can you answer either of these questions? Don’t fret if you can’t—you’ll know how to answer both of them before you finish this chapter. Let’s start by reviewing what you know about RAM thus far.



Cross Check

Dynamic and Random Access


You encountered DRAM back in Chapter 5, “Microprocessors,” so see if you can answer these questions: What makes DRAM dynamic? What does random access mean?

When not in use, programs and data are held in mass storage, which usually means a hard drive but could also mean a USB thumb drive, a CD-ROM, or some other device that can hold data when the computer is turned off. When you load a program by clicking an icon in Windows, the program is copied from the mass storage device to RAM and then run (Figure 6.1).

You saw in Chapter 5, “Microprocessors,” that the CPU uses dynamic random access memory (DRAM) as RAM for all PCs. Just like CPUs, DRAM has gone through a number of evolutionary changes over the years, resulting in improved DRAM technologies with names such as SDRAM, RDRAM, and DDR RAM. This chapter starts by explaining how DRAM works, and then moves into the types of DRAM used over the past few years to see how they improve on the original DRAM. The third section, “Working with RAM,” goes into the details of finding and installing RAM. The chapter finishes with troubleshooting RAM problems.



The CompTIA A+ certification domains use the term *memory* to describe the short-term storage used by the PC to load the operating system and running applications. The more common term in the industry is *RAM*, for random access memory, the kind of short-term memory you’ll find in every computer. More specifically, the primary system RAM is **dynamic random access memory (DRAM)**. For the most part, this book uses the terms RAM and DRAM.



• **Figure 6.1** Mass storage holds programs, but programs need to run in RAM.

Historical/Conceptual

■ Understanding DRAM

As discussed in Chapter 5, “Microprocessors,” DRAM functions like an electronic spreadsheet, with numbered rows containing cells and each cell holding a one or a zero. Now let’s look at what’s physically happening. Each spreadsheet cell is a special type of semiconductor that can hold a single bit—one or zero—by using microscopic capacitors and transistors. DRAM makers put these semiconductors into chips that can hold a certain number of bits. The bits inside the chips are organized in a rectangular fashion, using rows and columns.

Each chip has a limit on the number of lines of code it can contain. Think of each line of code as one of the rows on the electronic spreadsheet; one chip might be able to store a million rows of code while another chip can store over a billion lines. Each chip also has a limit on the width of the lines of code it can handle, so one chip might handle 8-bit-wide data while another might handle 16-bit-wide data. Techs describe chips by bits rather than bytes, so $\times 8$ and $\times 16$, respectively. Just as you could describe a spreadsheet by the number of rows



• **Figure 6.2** What do these numbers mean?

and columns—John’s accounting spreadsheet is huge, 48 rows \times 12 columns—memory makers describe RAM chips the same way. An individual DRAM chip that holds 1,048,576 rows and 8 columns, for example, would be a $1\text{ M} \times 8$ chip, with “M” as shorthand for “mega,” just like in megabytes (2^{20} bytes). It is difficult if not impossible to tell the size of a DRAM chip just by looking at it—only the DRAM makers know the meaning of the tiny numbers on the chips (see Figure 6.2), although sometimes you can make a good guess.

Organizing DRAM

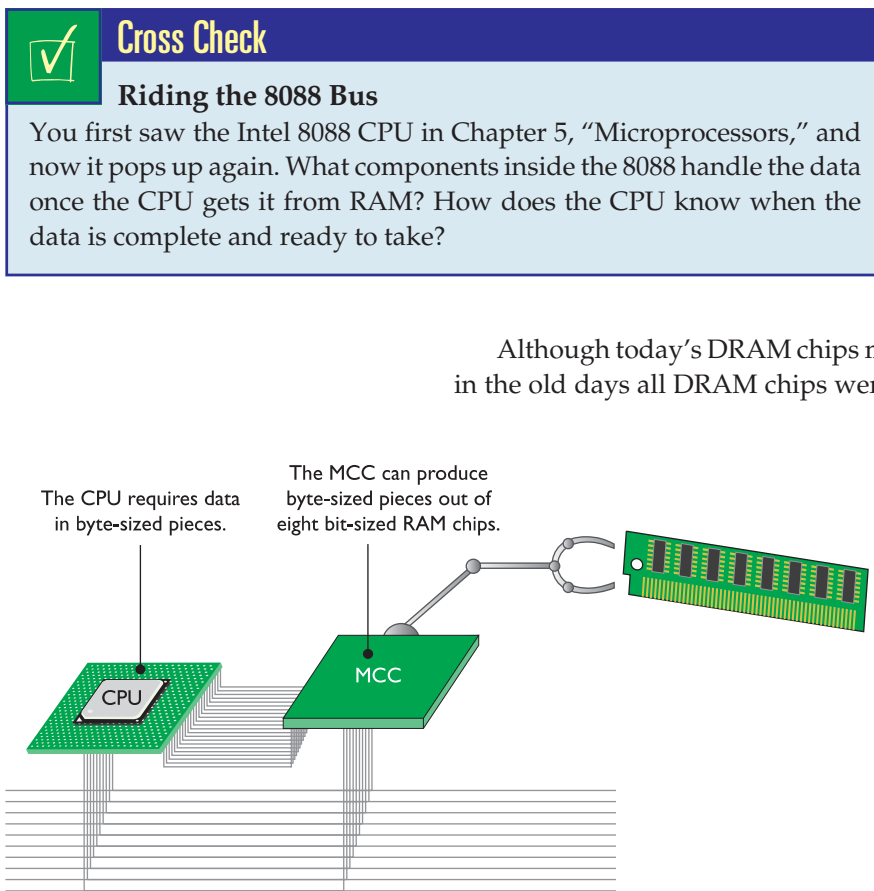
Because of its low cost, high speed, and capability to contain a lot of data in a relatively small package, DRAM has been the standard RAM used in all computers—not just PCs—since the mid-1970s. DRAM can be found in just about everything, from automobiles to automatic bread makers.

The PC has very specific requirements for DRAM. The original 8088 processor had an 8-bit frontside bus.

All the commands given to an 8088 processor were in discrete 8-bit chunks. You needed RAM that could store data in 8-bit (1-byte) chunks, so that each time the CPU asked for a line of code, the memory controller could put an 8-bit chunk on the data bus. This optimized the flow of data into (and out from) the CPU.

Although today’s DRAM chips may have widths greater than 1 bit, back in the old days all DRAM chips were 1 bit wide. That means you only had

such sizes as $64\text{ K} \times 1$ or $256\text{ K} \times 1$ —always 1 bit wide. So how was 1-bit-wide DRAM turned into 8-bit-wide memory? The answer was quite simple: Just take eight 1-bit-wide chips and electronically organize them with the memory controller chip to be eight wide. First, put eight 1-bit-wide chips in a row on the motherboard and then wire up this row of DRAM chips to the memory controller chip (which has to be designed to handle this) to make byte-wide memory (Figure 6.3). You just made eight 1-bit-wide DRAM chips look like a single 8-bit-wide DRAM chip to the CPU.



• **Figure 6.3** The MCC accessing data on RAM soldered onto the motherboard

Practical DRAM

Okay, before you learn more about DRAM, I need to make a critical point extremely clear. When you first saw the 8088’s machine language in Chapter 5,

“Microprocessors,” all the examples in the “codebook” were exactly 1-byte commands. Figure 6.4 shows the codebook again—see how all the commands are 1 byte?


Well, the reality is slightly different. Most of the 8088 machine language commands are 1 byte, but a few more complex commands need 2 bytes. For example, the following command tells the CPU to move 163 bytes “up the RAM spreadsheet” and run whatever command is there. Cool, eh?

```
1110100110100011
```

The problem here is that the command is 2 bytes wide, not 1 byte. So how did the 8088 handle this? Simple—it just took the command 1 byte at a time. It took twice as long to handle the command because the MCC had to go to RAM twice, but it worked.

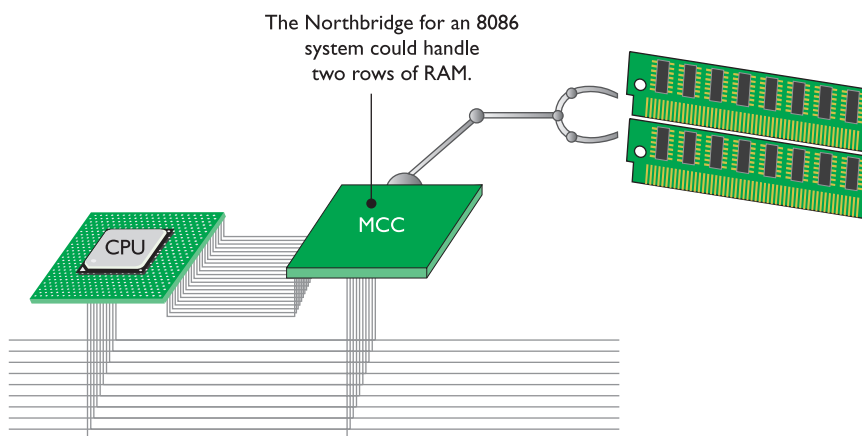
Okay, so if some of the commands are more than 1 byte wide, why didn’t Intel make the 8088 with a 16-bit frontside bus? Wouldn’t that have been better? Well, Intel did. Intel invented a CPU called the 8086. The 8086 actually predates the 8088 and was absolutely identical to the 8088 except for one small detail: it had a 16-bit frontside bus. IBM could have used the 8086 instead of the 8088 and used 2-byte-wide RAM instead of 1-byte-wide RAM. Of course, they would have needed to invent a memory controller chip that handled that kind of RAM (Figure 6.5).

Why didn’t Intel sell IBM the 8086 instead of the 8088? There were two reasons. First, nobody had invented an affordable MCC or RAM that handled 2 bytes at a time. Sure, chips were invented, but they were *expensive* and IBM didn’t think anyone would want to pay \$12,000 for a personal computer. So IBM bought the Intel 8088, not the Intel 8086, and all our RAM came in bytes. But as you might imagine, it didn’t stay that way too long.




LIGHTS	MEANING
10000000	The next line is a number, put it in the AX register
10010000	The next line is a number, put it in the BX register
10110000	Add AX to BX and put the result in AX
11000000	Place the value of AX on the External Data Bus
00000000	The number 0
00000001	The number 1
00000010	The number 2
00000011	The number 3
00000100	The number 4
00000101	The number 5
00000110	The number 6
00000111	The number 7
00001000	The number 8
00001001	The number 9

• Figure 6.4 Codebook again



• Figure 6.5 Pumped-up 8086 MCC at work



• **Figure 6.6** That's a lot of real estate used by RAM chips!



• **Figure 6.7** A 72-pin SIMM



Some MCCs are 128 bits wide.

DRAM Sticks

As CPU data bus sizes increased, so too did the need for RAM wide enough to fill the bus. The Intel 80386 CPU, for example, had a 32-bit data bus and thus the need for 32-bit-wide DRAM. Imagine having to line up 32 one-bit-wide DRAM chips on a motherboard. Talk about a waste of space! Figure 6.6 shows motherboard RAM run amuck.

DRAM manufacturers responded by creating wider DRAM chips, such as $\times 4$, $\times 8$, and $\times 16$, and putting multiples of them on a small circuit board called a **stick** or **module**. Figure 6.7 shows an early stick, called a **single inline memory module (SIMM)**, with eight DRAM chips. To add RAM to a modern machine means you need to get the right stick or sticks for the particular motherboard. Your motherboard manual tells you precisely what sort of module you need and how much RAM you can install.

Modern CPUs are a lot smarter than the old Intel 8088. Their machine languages have some commands that are up to 64 bits (8 bytes) wide. They also have at least a 64-bit frontside bus that can handle more than just 8 bits. They don't want RAM to give them a puny 8 bits at a time! To optimize the flow of data into and out of the CPU, the modern MCC provides at least 64 bits of data every time the CPU requests information from RAM.



Try This!


Dealing with Old RAM

Often in the PC world, old technology and ways of doing things are reimplemented with some newer technology. Learning how things worked back in the ancient days can stand a tech in good stead. Perhaps more importantly, many thousands of companies—including hospitals, auto repair places, and more—use very old, proprietary applications that keep track of medical records, inventory, and so on. If you're called to work on one of these ancient systems, you need to know how to work with old parts, so try this.

Obtain an old computer, such as a 386 or 486. Ask your uncle, cousin, or Great Aunt Edna if they have a PC collecting dust in a closet that you can use. Failing that, go to a secondhand store or market and buy one for a few dollars.

Open up the system and check out the RAM. Remove the RAM from the motherboard and then replace it to familiarize yourself with the internals. You never know when some critical system will go down and need repair immediately—and you're the one to do it!

Modern DRAM sticks come in 32-bit- and 64-bit-wide form factors with a varying number of chips. Many techs describe these memory modules by their width, so $\times 32$ and $\times 64$. Note that this number does *not* describe the width of the individual DRAM chips on the module. When you read or hear about *by whatever* memory, simply note that you need to know whether that person is talking about the DRAM width or the module width. When the CPU needs certain bytes of data, it requests those bytes via the address bus. The CPU does not know the physical location of the RAM that stores that data, nor the physical makeup of the RAM—such as how many DRAM chips work together to provide the 64-bit-wide memory rows. The MCC keeps track of this and just gives the CPU which-ever bytes it requests (Figure 6.8).



• **Figure 6.8** The MCC knows the real location of the DRAM.

Consumer RAM

If modern DRAM modules come in sizes much wider than a byte, why do people still use the word “byte” to describe how much DRAM they have? Convention. Habit. Rather than using a label that describes the electronic structure of RAM, common usage describes the *total capacity of RAM on a stick in bytes*. John has a single 512-MB stick of RAM on his motherboard, for example, and Sally has two 256-MB sticks. Both systems have a total of 512 MB of system RAM. That’s what your clients care about, after all, because having enough RAM makes their systems snappy and stable; not enough and the systems run poorly. As a tech, you need to know more, of course, to pick the right RAM for many different types of computers.

Essentials

■ Types of RAM

Development of newer, wider, and faster CPUs and MCCs motivate DRAM manufacturers to invent new DRAM technologies that deliver enough data at a single pop to optimize the flow of data into and out of the CPU.

SDRAM


Most modern systems use some form of **synchronous DRAM (SDRAM)**. SDRAM is still DRAM, but it is *synchronous*—tied to the system clock, just like the CPU and MCC, so the MCC knows when data is ready to be grabbed from SDRAM. This results in little wasted time.



Old RAM—really old RAM—was called *fast page mode* (FPM) RAM. This ancient RAM used a totally different technology that was not tied to the system clock. If you ever hear of FPM RAM, it’s going to be in a system that’s over a decade old. Be careful! CompTIA likes to use older terms like this to throw you off!



• **Figure 6.9** 144-pin micro-DIMM (photo courtesy of Micron Technology, Inc.)



• **Figure 6.10** A 168-pin DIMM above a 144-pin SO-DIMM

SDRAM made its debut in 1996 on a stick called a **dual inline memory module (DIMM)**. The early SDRAM DIMMs came in a wide variety of pin sizes. The most common pin sizes found on desktops were the 168-pin variety. Laptop DIMMs came in 68-pin, 144-pin (Figure 6.9), or 172-pin *micro-DIMM* packages; and the 72-pin, 144-pin, or 200-pin **small-outline DIMM (SO-DIMM)** form factors (Figure 6.10). With the exception of the 32-bit 72-pin SO-DIMM, all these DIMM varieties delivered 64-bit-wide data to match the 64-bit data bus of every CPU since the Pentium.

To take advantage of SDRAM, you needed a PC designed to use SDRAM. If you had a system with slots for 168-pin DIMMs, for example, your system used SDRAM. A DIMM in any one of the DIMM slots could fill the 64-bit bus, so each slot was called a **bank**. You could install one, two, or more sticks and the system would work. Note that on laptops that used the 72-pin SO-DIMM, you needed to install two sticks of RAM to make a full bank, because each stick only provided half the bus width.

SDRAM tied to the system clock, so its clock speed matched the frontside bus. Five clock speeds were commonly used on the early SDRAM systems: 66, 75, 83, 100, and 133 MHz. The RAM speed had to match or exceed the system speed or the computer would be unstable or wouldn't work at all. These speeds were prefixed with a "PC" in the front, based on a standard forwarded by Intel, so SDRAM speeds were PC66 through PC133. For a Pentium III computer with a 100-MHz frontside bus, you needed to buy SDRAM DIMMs rated to handle it, such as PC100 or PC133.

RDRAM

When Intel was developing the Pentium 4, they knew that regular SDRAM just wasn't going to be fast enough to handle the quad-pumped 400-MHz frontside bus. Intel announced plans to replace SDRAM with a very fast, new type of RAM developed by Rambus, Inc., called **Rambus DRAM**, or simply **RDRAM** (Figure 6.11). Hailed by Intel as the next great leap in DRAM technology, RDRAM could handle speeds up to 800 MHz, which gave Intel plenty of room to improve the Pentium 4.

RDRAM was greatly anticipated by the industry for years, but industry support for RDRAM proved less than enthusiastic due to significant delays in development and a price many times that of SDRAM. Despite this grudging support, almost



• **Figure 6.11** RDRAM



The 400-MHz frontside bus speed wasn't achieved by making the system clock faster—it was done by making CPUs and MCCs capable of sending 64 bits of data two or four times for every clock cycle, effectively doubling or quadrupling the system bus speed.



Cross Check

Double-Pumped and Quad-Pumped

You've seen double-pumped and quad-pumped frontside buses in Chapter 5, "Microprocessors," so see if you can answer these questions. What CPU—by codename—started the double-pumped bus bandwagon? What socket did it use?

all major PC makers sold systems that used RDRAM—for a while. From a tech’s standpoint, RDRAM shares almost all of the characteristics of SDRAM. A stick of RDRAM is called a **RIMM**. In this case, however, the letters don’t actually stand for anything; they just rhyme: SIMMs, DIMMs, and now RIMMs, get it?

RDRAM RIMMs came in two sizes: a 184-pin for desktops and a 160-pin SO-RIMM for laptops. RIMMs were keyed differently from DIMMs to ensure that even though they are the same basic size, you couldn’t accidentally install a RIMM in a DIMM slot or vice versa. RDRAM also had a speed rating: 600 MHz, 700 MHz, 800 MHz, or 1066 MHz. RDRAM employed an interesting **dual-channel architecture**. Each RIMM was 64 bits wide, but the Rambus MCC alternated between two sticks to increase the speed of data retrieval. You were required to install RIMMs in pairs to use this dual-channel architecture.

RDRAM motherboards also required that all RIMM slots be populated. Unused pairs of slots needed a passive device called a **continuity RIMM (CRIMM)** installed in each slot to enable the RDRAM system to terminate properly. Figure 6.12 shows a CRIMM.

RDRAM offered dramatic possibilities for high-speed PCs but ran into three roadblocks that Betamax’d it. First, the technology was owned wholly by Rambus; if you wanted to make it, you had to pay the licensing fees they charged. That led directly to the second problem, expense. RDRAM cost substantially more than SDRAM. Third, Rambus and Intel made a completely closed deal for the technology. RDRAM worked only on Pentium 4 systems using Intel-made MCCs. AMD was out of luck. Clearly, the rest of the industry had to look for another high-speed RAM solution.

DDR SDRAM

AMD and many major system and memory makers threw their support behind **double data rate SDRAM (DDR SDRAM)**. DDR SDRAM basically copied Rambus, doubling the throughput of SDRAM by making two processes for every clock cycle. This synchronized (pardon the pun) nicely with the Athlon and later AMD processors’ double-pumped frontside bus. DDR SDRAM could not run as fast as RDRAM—although relatively low frontside bus speeds made that a moot point—but cost only slightly more than regular SDRAM.


DDR SDRAM for desktops comes in 184-pin DIMMs. These DIMMs match 168-pin DIMMs in physical size but not in pin compatibility (Figure 6.13). The slots for the two types of RAM appear similar as well but have different guide notches, so you can’t insert either type of RAM into the other’s slot.



• **Figure 6.13** DDR SDRAM



On the CompTIA A+ exams, you’ll see RDRAM referred to as **RAMBUS RAM**. Don’t get thrown by the odd usage.



• **Figure 6.12** CRIMM



Betamax’d is slang for “made it obsolete because no one bought it, even though it was a superior technology to the winner in the marketplace.” Refers to the VHS versus Betamax wars in the old days of video cassette recorders.



Tech Tip

RAM Slang

Most techs drop some or all of the SDRAM part of DDR SDRAM when engaged in normal geekspeak. You’ll hear the memory referred to as DDR, DDR RAM, and the weird hybrid, DDRAM.



RAM makers use the term **single data rate SDRAM (SDR SDRAM)** for the original SDRAM to differentiate it from DDR SDRAM.



• **Figure 6.14** 172-pin DDR SDRAM micro-DIMM (photo courtesy of Kingston/Joint Harvest)

DDR SDRAM for laptops comes in either 200-pin SO-DIMMs or 172-pin micro-DIMMs (Figure 6.14).

DDR sticks use a rather interesting naming convention—actually started by the Rambus folks—based on the number of bytes per second of data throughput the RAM can handle. To determine the bytes per second, take the MHz speed and multiply by 8 bytes (the width of all DDR SDRAM sticks). So 400 MHz multiplied by 8 is 3200 megabytes per second. Put the abbreviation “PC” in the front to make the new term: PC3200. Many techs also use the naming convention used for the individual DDR chips; for example, *DDR400* refers to a 400-MHz DDR SDRAM chip running on a 200-MHz clock.

Even though the term *DDRxxx* is really just for individual DDR chips and the term *PCxxxx* is for DDR sticks, this tradition of two names for every speed of RAM is a bit of a challenge because you’ll often hear both terms used interchangeably. Table 6.1 shows all the speeds for DDR—not all of these are commonly used.

Following the lead of AMD, VIA, and other manufacturers, the PC industry adopted DDR SDRAM as the standard system RAM. In the summer of 2003, Intel relented and stopped producing motherboards and memory controllers that required RDRAM.

One thing is sure about PC technologies: Any good idea that can be copied will be copied. One of Rambus’ best concepts was the dual-channel architecture—using two sticks of RDRAM together to increase throughput. Manufacturers have released motherboards with MCCs that support dual-channel architecture using



Try This!

DRAM Availability

The DRAM market changes fairly rapidly at the higher end, but products tend to linger at the low end, seemingly well past their usefulness. What do your class systems need? What’s available in your area today? Try this.


Check out the RAM requirements for the PCs available in your class (or home or office). You can open them up for physical examination or read the motherboard books. Then go to your friendly neighborhood computer store and see what’s available. Does the store offer memory that your system cannot use, such as EDO DRAM on 72-pin SIMMs or registered DDR SDRAM? What does this tell you about the PCs available to you?

Table 6.1

DDR Speeds

Clock Speed	DDR Speed Rating	PC Speed Rating
100 MHz	DDR200	PC1600
133 MHz	DDR266	PC2100
166 MHz	DDR333	PC2700
200 MHz	DDR400	PC3200
217 MHz	DDR433	PC3500
233 MHz	DDR466	PC3700
250 MHz	DDR500	PC4000
275 MHz	DDR550	PC4400
300 MHz	DDR600	PC4800

Dual-channel DDR works like RDRAM in that you must have two identical sticks of DDR and they must snap into two paired slots. Unlike RDRAM, dual-channel DDR doesn't have anything like CRIMMs—you don't need to put anything into unused slot pairs. Dual-channel DDR technology is very flexible but also has a few quirks that vary with each system. Some motherboards have three DDR SDRAM slots, but the dual-channel DDR works only if you install DDR SDRAM in two of the slots. Other boards have four slots and you must install matching pairs in the same colored slots to run in dual-channel mode (Figure 6.15). If you populate a third slot, the system uses the full capacity of RAM installed but turns off the dual-channel feature.



The fastest versions of DDR RAM run at a blistering PC4800. That's 4.8 gigabytes per second (GBps) of data throughput! You'd think that kind of speed would satisfy most users, and to be honest, DRAM running at approximately 5 GBps really is plenty fast—for yesterday. However, the ongoing speed increases ensure that even these speeds won't be good enough in the future. Knowing this, the RAM industry came out with DDR2, the successor to DDR. *DDR2* is DDR RAM with some improvements in its electrical characteristics, enabling it to run even faster than DDR while using less power. The big speed increase from DDR2 comes by clock-doubling the input/output circuits on the chips. This does not speed up the core RAM—the part that holds the data—but speeding up the input/output and adding special buffers (sort of like a cache) make DDR2 run much faster than regular DDR. DDR2 uses a 240-pin DIMM that's not compatible with DDR (Figure 6.16). Likewise, the DDR2 200-pin SO-DIMM is incompatible with the DDR SO-DIMM. You'll find motherboards running both single-channel and dual-channel DDR2.




Table 6.2 shows some of the common DDR2 speeds.

Table 6.2		DDR2 Speeds	
Core RAM Clock Speed	DDR I/O Speed	DDR2 Speed Rating	PC Speed Rating
100 MHz	200 MHz	DDR2-400	PC2-3200
133 MHz	266 MHz	DDR2-533	PC2-4200
166 MHz	333 MHz	DDR2-667	PC2-5300
200 MHz	400 MHz	DDR2-800	PC2-6400
250 MHz	500 MHz	DDR2-1000	PC2-8000

DDR3

DDR2 has been the standard for several years, but now there's a new kid on the block. *DDR3* boasts higher speeds, more efficient architecture, and around 30 percent lower power consumption than DDR2 RAM, making it a

compelling choice for system builders. Just like its predecessor, DDR3 uses a 240-pin DIMM, albeit one that is slotted differently to make it difficult for users to install the wrong RAM in their system without using a hammer (Figure 6.17). DDR3 SO-DIMMs for portable computers have 204 pins. Neither fits into a DDR2 socket.

DDR3 doubles the buffer of DDR2 from 4 bits to 8 bits, giving it a huge boost in bandwidth over older RAM. Not only that, but some DDR3 modules also include a feature called XMP, or extended mem-

ory profile, that enables power users to overclock their RAM easily, boosting their already fast memory to speeds that would make Chuck Yeager nervous. DDR3 modules also use higher-density memory chips, which means we may eventually see 16 GB DDR3 modules.

Some chipsets that support DDR3 also support a feature called *triple-channel memory*, which works a lot like dual-channel before it, but with three sticks of RAM instead of two. You'll need three of the same type of memory modules and a motherboard that supports it, but triple-channel memory can greatly increase performance for those who can afford it.

In keeping with established tradition, Table 6.3 is a chart of common DDR3 speeds. Note how DDR3 I/O speeds are quadruple the clock speeds, whereas DDR2 I/O speeds are only double the clock. This speed increase is due to the increased buffer size, which enables DDR3 to grab twice as much data every clock cycle as DDR2 can.



• **Figure 6.17** DDR2 DIMM on top of a DDR3 DIMM

Do not confuse DDR3 with GDDR3; the latter is a type of memory used solely in video cards. See Chapter 19, "Video," for the scoop on video-specific types of memory.


Table 6.3		DDR3 Speeds	
Core RAM Clock Speed	DDR I/O Speed	DDR3 Speed Rating	PC Speed Rating
100 MHz	400 MHz	DDR3-800	PC3-6400
133 MHz	533 MHz	DDR3-1066	PC3-8500
166 MHz	667 MHz	DDR3-1333	PC3-10667
200 MHz	800 MHz	DDR3-1600	PC3-12800

RAM Variations

Within each class of RAM, you'll find variations in packaging, speed, quality, and the capability to handle data with more or fewer errors. Higher-end systems often need higher-end RAM, so knowing these variations is of crucial importance to techs.

Double-Sided DIMMs

Every type of RAM stick, starting with the old FPM SIMMs and continuing through to 240-pin DDR3 SDRAM, comes in one of two types: **single-sided** and **double-sided**. As their name implies, single-sided sticks have chips on only one side of the stick. Double-sided sticks have chips on both sides (Figure 6.18). The vast majority of RAM sticks are single-sided, but plenty of double-sided sticks are out there. Double-sided sticks are basically two sticks of RAM soldered onto one board. There's nothing wrong with double-sided RAM other than the fact that some motherboards either can't use them or can only use them in certain ways—for example, only if you use a single stick and it goes into a certain slot.




• **Figure 6.18** Double-sided DDR SDRAM

Latency

If you've shopped for RAM lately, you may have noticed terms such as "CL2" or "low latency" as you tried to determine which RAM to purchase. You might find two otherwise identical RAM sticks with a 20 percent price difference and a salesperson pressuring you to buy the more expensive one because it's "faster" even though both sticks say DDR 3200 (Figure 6.19).

RAM responds to electrical signals at varying rates. When the memory controller starts to grab a line of memory, for example, a slight delay occurs; think of it as the RAM getting off the couch. After the RAM sends out the



• **Figure 6.19** Why is one more expensive than the other?



Tech Tip

Latency Is Relative

Latency numbers reflect how many clicks of the system clock it takes before the RAM responds. If you speed up the system clock, say from 166 MHz to 200 MHz, the same stick of RAM might take an extra click before it can respond. When you take RAM out of an older system and put it into a newer one, you might get a seemingly dead PC, even though the RAM fits in the DIMM slot. Many motherboards enable you to adjust the RAM timings manually. If so, try raising the latency to give the slower RAM time to respond. See Chapter 7, "BIOS and CMOS," to learn how to make these adjustments (and how to recover if you make a mistake).



CAS stands for *column array strobe*, one of the wires (along with the *row array strobe*) in the RAM that helps the memory controller find a particular bit of memory. Each of these wires requires electricity to charge up before it can do its job. This is one of the aspects of latency.

requested line of memory, there's another slight delay before the memory controller can ask for another line—the RAM sat back down. The delay in RAM's response time is called its **latency**. RAM with a lower latency—such as CL2—is faster than RAM with a higher latency—such as CL3—because it responds more quickly. The CL refers to clock cycle delays. The 2 means that the memory delays two clock cycles before delivering the requested data; the 3 means a three-cycle delay.

From a tech's standpoint, you need to get the proper RAM for the system you're working on. If you put a high-latency stick in a motherboard set up for a low-latency stick, you'll get an unstable or completely dead PC. Check the motherboard manual and get the quickest RAM the motherboard can handle, and you should be fine.

Parity and ECC

Given the high speeds and phenomenal amount of data moved by the typical DRAM chip, a RAM chip might occasionally give bad data to the memory controller. This doesn't necessarily mean that the RAM has gone bad. It could be a hiccup caused by some unknown event that makes a good DRAM chip say a bit is a zero when it's really a one. In most cases you won't even notice when such a rare event happens. In some environments, however, even these rare events are intolerable. A bank server handling thousands of online transactions per second, for example, can't risk even the smallest error. These important computers need a more robust, fault-resistant RAM.


The first type of error-detecting RAM was known as parity RAM (Figure 6.20). **Parity RAM** stored an extra bit of data (called the parity bit) that the MCC used to verify whether the data was correct. Parity wasn't perfect—it wouldn't always detect an error, and if the MCC did find an error, it couldn't correct the error. For years, parity was the only available way to tell if the RAM made a mistake.

Today's PCs that need to watch for RAM errors use a special type of RAM called **error correction code RAM (ECC RAM)**. ECC is a major advance in error checking on DRAM. First,

ECC detects any time a single bit is incorrect. Second, ECC fixes these errors on the fly. The checking and fixing come at a price, however, as ECC RAM is always slower than non-ECC RAM.

ECC DRAM comes in every DIMM package type and can lead to some odd-sounding numbers. You can find DDR2 or DDR3 RAM sticks, for example, that come in 240-pin, 72-bit versions. Similarly, you'll see 200-pin, 72-bit SO-DIMM format. The extra 8 bits beyond the 64-bit data stream are for the ECC.

You might be tempted to say "Gee, maybe I want to try this ECC RAM." Well, don't! To take advantage of ECC RAM, you need a motherboard with an MCC designed to use ECC. Only expensive motherboards for high-end systems use ECC. The special-use-only nature of ECC makes it fairly rare. Plenty of techs with years of experience have never even seen ECC RAM.



• **Figure 6.20** Ancient parity RAM stick



Some memory manufacturers call the technology *Error Checking and Correction (ECC)*. Don't be thrown off if you see the phrase—it's the same thing, just a different marketing slant for error correction code.



Tech Tip

Counting Chips

To determine if a DRAM stick is parity or ECC, count the individual chips on the stick. If you have an odd number of chips on the stick, you have ECC or parity DRAM.

Buffered/Registered DRAM

Your average PC motherboard accepts no more than four sticks of DRAM, because more than four physical slots for sticks gives motherboard designers some serious electrical headaches. Yet some systems that use a lot of RAM need the capability to use more DRAM sticks on the motherboard, often six or eight. To get around the electrical hassles, special DRAM sticks add a buffering chip to the stick that acts as an intermediary between the DRAM and the MCC. These special DRAM sticks are called **buffered** or **registered DRAM** (Figure 6.21).

Like ECC, you must have a motherboard with an MCC designed to use this type of DRAM. Rest assured that such a motherboard has a large number of RAM slots. Buffered/registered RAM is rare (maybe not quite as rare as ECC RAM), and you'll never see it in the typical desktop system.




• Figure 6.21 Buffered DRAM

■ Working with RAM

Whenever someone comes up to me and asks what single hardware upgrade they can do to improve their system performance, I always tell them the same thing—add more RAM. Adding more RAM can improve overall system performance, processing speed, and stability—if you get it right. Botching the job can cause dramatic system instability, such as frequent, random crashes and reboots. Every tech needs to know how to install and upgrade system RAM of all types.

To get the desired results from a RAM upgrade, you must first determine if insufficient RAM is the cause of system problems. Second, you need to pick the proper RAM for the system. Finally, you must use good installation practices. Always store RAM sticks in anti-static packaging whenever they're not in use, and use strict ESD handling procedures. Like many other pieces of the PC, RAM is *very* sensitive to ESD and other technician abuse (Figure 6.22)!



• Figure 6.22 Don't do this! Grabbing the contacts is a


Do You Need RAM?

Two symptoms point to the need for more RAM in a PC: general system sluggishness and excessive hard drive accessing. If programs take forever to load and running programs seem to stall and move more slowly than you would like, the problem could stem from insufficient RAM. A friend with a new Windows Vista system complained that her PC seemed snappy when she first got it but takes a long time to do the things she wants to do with it, such as photograph retouching in Adobe Photoshop and document layout for a print zine she produces. Her system had only 1 GB of RAM, sufficient to run Windows Vista, but woefully insufficient for her tasks—she kept

maxing out the RAM and thus the system slowed to a crawl. I replaced her stick with a pair of 2-GB sticks and suddenly she had the powerhouse workstation she desired.

Excessive hard drive activity when you move between programs points to a need for more RAM. Every Windows PC has the capability to make a portion of your hard drive look like RAM in case you run out of real RAM. This is called the **page file** or **swap file**, as you'll recall from Chapter 4, "Understanding Windows." If you fill your RAM up with programs, your PC automatically starts loading some programs into the page file. You can't see this process taking place just by looking at the screen—these swaps are done in the background. But you will notice the hard drive access LED going crazy as Windows rushes to move programs between RAM and the page file in a process called **disk thrashing**. Windows uses the page file all the time, but excessive disk thrashing suggests that you need more RAM.

You can diagnose excessive disk thrashing through simply observing the hard drive access LED flashing or through various third-party tools. I like FreeMeter (www.tiler.com/freemeter/). It's been around for quite a while, runs on all versions of Windows, and is easy to use (Figure 6.23). Notice on the FreeMeter screenshot that some amount of the page file is being used. That's perfectly normal.



• Figure 6.23 FreeMeter

System RAM Recommendations

Microsoft sets very low the minimum RAM requirements listed for the various Windows operating systems to get the maximum number of users to upgrade or convert, and that's fine. A Windows XP Professional machine runs well enough on 128 MB of RAM. Just don't ask it to do any serious computing, such as running Doom III! Windows



Try This!

Checking the Page File

How much of your hard drive does Windows use for a page file? Does the level change dramatically when you open typical applications, such as Microsoft Word, Solitaire, and Paint Shop Pro? The answers to these questions can give a tech a quick estimation about RAM usage and possibly RAM needs for a particular system, so try this.

In Windows 2000/XP you can easily glance at your page file usage through the Task Manager. To access the Task Manager, press CTRL-ALT-DEL simultaneously once. Click the Performance tab. The second box on the left, titled PF Usage, displays the amount of hard drive the page file is currently using. In Windows Vista, the process of viewing the page file is exactly the same, but Vista uses a less descriptive way of displaying its usage. In the bottom-right corner of the Performance tab, you'll see the title "Page File," which gives you a numeric representation.

1. How big is the page file when you have no applications open?
2. How much does it change when you open applications?

Vista has raised the bar considerably, especially with the 64-bit version of the operating system. Here are my recommendations for system RAM.

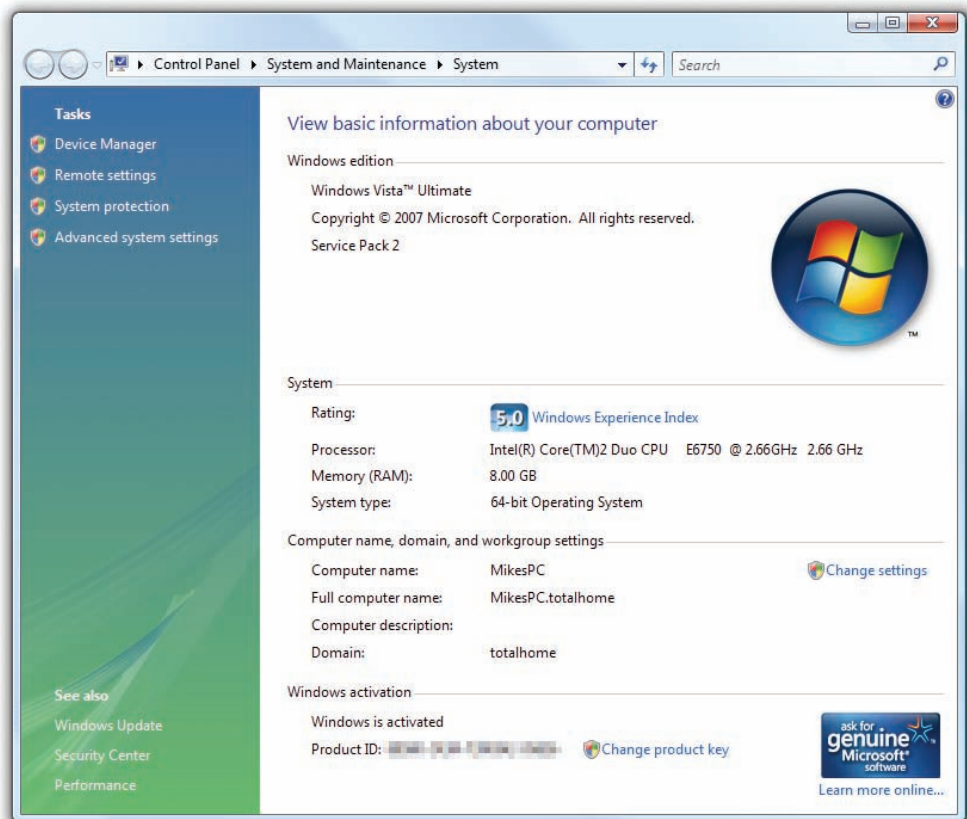
Operating System	Reasonable Minimum	Solid Performance	Power User
Windows 2000	128 MB	256 MB	512 MB
Windows XP	256 MB	1 GB	2 GB
Windows Vista	2 GB	4 GB	8 GB

Determining Current RAM Capacity


Before you go get RAM, you obviously need to know how much RAM you currently have in your PC. Every version of Windows works the same way. Just select the Properties for My Computer or Computer to see how much RAM is in your system (Figure 6.24). If you have a newer keyboard, you can access the screen with the WINDOWS-PAUSE/BREAK keystroke combination. Windows 2000, XP, and Vista come with the handy Performance tab under the Task Manager (as shown in Figure 6.25).

Getting the Right RAM

To do the perfect RAM upgrade, determine the optimum capacity of RAM to install and then get the right RAM for the motherboard. Your first two stops toward these goals are the inside of the case and your motherboard manual. Open the case to see how many sticks of RAM you have installed currently and how many free slots you have open. Check the motherboard book to determine the total capacity of RAM the system can handle and what specific technology works with your system. You can't put DDR2 into a system that can only handle DDR SDRAM, after all, and it won't do you much good to install a pair of 2-GB DIMMs when your system tops out at 1.5 GB. Figure 6.26 shows the RAM limits for my ASUS Crosshair motherboard.



• **Figure 6.24** Mike has a lot of RAM!



• **Figure 6.25** Performance tab in Windows XP Task Manager

Tech Tip

CPU-Z

The freeware CPU-Z program tells you the total number of slots on your motherboard, the number of slots used, and the exact type of RAM in each slot—very handy. CPU-Z not only determines the latency of your RAM, but also lists the latency at a variety of motherboard speeds. The CD accompanying this book has a copy of CPU-Z, so check it out.

Crosshair specifications summary	
CPU	Support AMD® Socket AM2 Athlon 64 X2 / Athlon 64 FX / Athlon 64 / Sempron AMD Cool'n'Quiet™ Technology AMD64 architecture enables simultaneous 32-bit and 64-bit computing AMD Live!™ Ready
Chipset	NVIDIA nForce® 590 SLI™ MCP NVIDIA LinkBoost™ Technology
System bus	2000 / 1600 MT/s
Memory	Dual channel memory architecture 4 x DIMM, max. 8GB, DDR2-800/667/533, ECC and non-ECC, un-buffered memory
Expansion slots	2 x PCI Express x16 slot with NVIDIA® SLI™ technology support, at full x16, x16 speed 1 x PCI Express x4 3 x PCI 2.2
Scalable Link Interface (SLI™)	Support two identical NVIDIA SLI-Ready graphics cards (both at x16 mode) ASUS two-slot thermal design ASUS PEG Link
High Definition Audio	SupremeFX Audio Card featuring ADI 1988B 8-channel High Definition Audio CODEC Support Jack-Sensing, Enumeration, Multi-streaming and Jack-Retasking 8 channel audio ports Coaxial, Optical S/PDIF out on back I/O port * ASUS Array Mic * Noise Filter
Storage	NVIDIA nForce® 590 SLI™ MCP supports: * 1 x Ultra DMA 133 / 100 / 66 / 33 * 6 x Serial ATA 3.0Gb/s with NCQ * NVIDIA MediaShield™ RAID supports RAID 0, 1, 0+1, 5 and JBOD span cross Serial ATA drives Silicon Image® 3132 SATA controller supports: * 2 x External Serial ATA 3.0Gb/s port on back I/O (SATA On-the-Go) * Support RAID 0, 1, JBOD, RAID 0+1(10) and 5 through multiplier

• **Figure 6.26** The motherboard book shows how much RAM the motherboard will handle.

Mix and Match at Your Peril

All motherboards can handle different capacities of RAM. If you have three slots, you may put a 512-MB stick in one and a 1-GB stick in the other with a high chance of success. To ensure maximum stability in a system, however, shoot for as close as you can get to uniformity of RAM. Choose RAM sticks that match in technology, capacity, and speed. Even on motherboards that offer slots for radically different RAM types, I recommend uniformity.

Mixing Speeds

With so many different DRAM speeds available, you may often find yourself tempted to mix speeds of DRAM in the same system. Although you may get away with mixing speeds on a system, the safest, easiest rule to follow is


to use the speed of DRAM specified in the motherboard book, and make sure that every piece of DRAM runs at that speed. In a worst-case scenario, mixing DRAM speeds can cause the system to lock up every few seconds or every few minutes. You might also get some data corruption. Mixing speeds sometimes works fine, but don't do your income tax on a machine with mixed DRAM speeds until the system has proven to be stable for a few days. The important thing to note here is that you won't break anything, other than possibly data, by experimenting.

Okay, I have mentioned enough disclaimers. Modern motherboards provide some flexibility regarding RAM speeds and mixing. First, you can use RAM that is faster than the motherboard specifies. For example, if the system needs PC3200 DDR2 SDRAM, you may put in PC4200 DDR2 SDRAM and it should work fine. Faster DRAM is not going to make the system run any faster, however, so don't look for any system improvement.

Second, you can sometimes get away with putting one speed of DRAM in one bank and another speed in another bank, as long as all the speeds are as fast as or faster than the speed specified by the motherboard. Don't bother trying to put different-speed DRAM sticks in the same bank with a motherboard that uses dual-channel DDR. Yes, it works once in a while, but it's too chancy. I avoid it.

Installing DIMMs and RIMMs

Installing DRAM is so easy that it's one of the very few jobs I recommend to non-techie folks. First, attach an anti-static wrist strap or touch some bare metal on the power supply to ground yourself and avoid ESD. Then swing the side tabs on the RAM slots down from the upright position. Pick up a stick of RAM—don't touch those contacts—and line up the notch or notches with the raised portion(s) of the DIMM socket (Figure 6.27). A good hard push down is usually all you need to ensure a solid connection. Make sure that the DIMM snaps into position to show it is completely seated. Also, notice that the two side tabs move in to reflect a tight connection.




• Figure 6.27 Inserting a DIMM

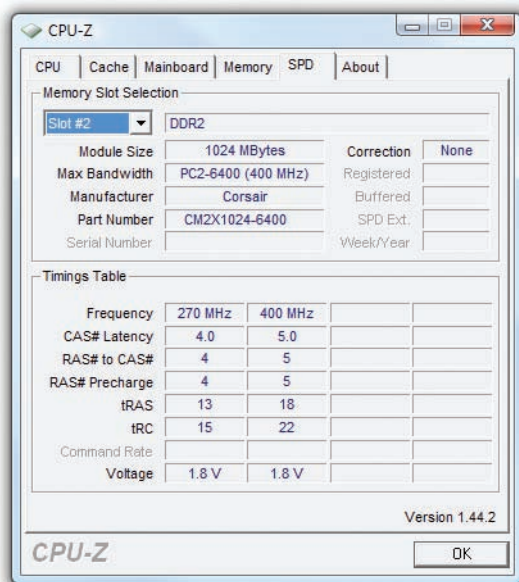
Serial Presence Detect (SPD)

Your motherboard should detect and automatically set up any DIMM or RIMM you install, assuming you have the right RAM for the system, using a technology called **serial presence detect (SPD)**. RAM makers add a handy chip to modern sticks called the SPD chip (Figure 6.28). The SPD chip stores all the information about your DRAM, including size, speed, ECC or non-ECC, registered or unregistered, and a number of other more technical bits of information.

When a PC boots, it queries the SPD chip so that the MCC knows how much RAM is on the stick, how fast it runs, and other information.



• Figure 6.28 SPD chip on a stick



• **Figure 6.29** CPU-Z showing RAM information

```
Award Modular BIOS v6.00PG, An Energy Star Ally
Copyright (C) 1984-2005, Award Software, Inc.

GA-K8NP F13

Processor : AMD Athlon(tm) 64 Processor 3200+
<CPUID:0000F4A Patch ID:003A>
Memory Testing : 1048576K OK
CPU clock frequency : 200 Mhz

Detecting IDE drives ...
```

• **Figure 6.30** Hey, where's the rest of my RAM?!

```
Award Modular BIOS v6.00PG, An Energy Star Ally
Copyright (C) 1984-2005, Award Software, Inc.

GA-K8NP F13

Processor : AMD Athlon(tm) 64 Processor 3200+
<CPUID:0000F4A Patch ID:003A>
Memory Testing : 3145728K OK
CPU clock frequency : 200 Mhz

Detecting IDE drives ...
```

• **Figure 6.31** RAM count after proper insertion of DIMMs

Any program can query the SPD chip. Take a look at Figure 6.29 with the results of the popular CPU-Z program showing RAM information from the SPD chip.

All new systems count on SPD to set the RAM timings properly for your system when it boots. If you add a RAM stick with a bad SPD chip, you'll get a POST error message and the system will not boot. You can't fix a broken SPD chip; you just buy a new stick of RAM.

The RAM Count

After installing the new RAM, turn on the PC and watch the boot process closely. If you installed the RAM correctly, the RAM count on the PC reflects the new value (compare Figures 6.30 and 6.31). If the RAM value stays the same, you probably have installed the RAM in a slot the motherboard doesn't want you to use (for example, you may need to use a particular slot first) or have not installed the RAM properly. If the computer does not boot and you've got a blank screen, you probably have not installed all the RAM sticks correctly. Usually, a good second look is all you need to determine the problem. Reseat or reinstall the RAM stick and try again.

RAM counts are confusing because RAM uses megabytes and gigabytes as opposed to millions and billions. Here are some examples of how different systems would show 256 MB of RAM:

268435456 (exactly $256 \times 1 \text{ MB}$)

256M (some PCs try to make it easy for you)

262,144 (number of KB)

You should know how much RAM you're trying to install and use some common sense. If you have 512 MB and you add another 512-MB stick, you need a number that looks like one gigabyte. After you add the second stick, if you see a RAM count of 524582912—that sure looks like 512 MB, not the one gigabyte!

Installing SO-DIMMs in Laptops

It wasn't that long ago that adding RAM to a laptop was either impossible or required you to send the system back to the manufacturer. For years, every laptop maker had custom-made, proprietary RAM packages that were difficult to handle and staggeringly expensive. The wide acceptance of SO-DIMMs over the last few years has virtually erased these problems. All laptops now provide relatively convenient access to their SO-DIMMs, enabling easy replacement or addition of RAM.

Access to RAM usually requires removing a panel or lifting up the keyboard—the procedure varies among laptop manufacturers. Figure 6.32 shows a typical laptop RAM access panel. You can slide the panel off to reveal the SO-DIMMs. SO-DIMMs usually insert exactly like the old SIMMs; slide the pins into position and snap the SO-DIMM down into the retaining clips (Figure 6.33).



• **Figure 6.32** A RAM access panel on a laptop



• **Figure 6.33** Snapping in an SO-DIMM

Before doing any work on a laptop, turn the system off, disconnect it from the AC wall socket, and remove all batteries. Use an anti-static wrist strap because laptops are far more susceptible to ESD than desktop PCs.

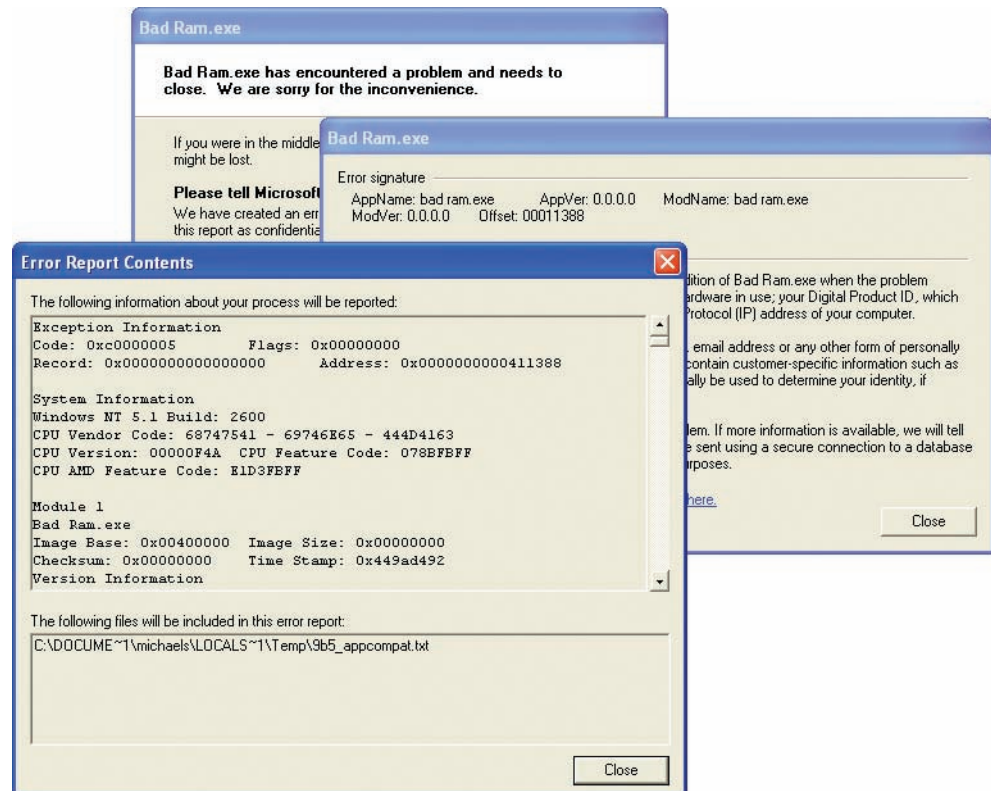
Practical Application

■ Troubleshooting RAM

“Memory” errors show up in a variety of ways on modern systems, including parity errors, ECC error messages, system lockups, page faults, and other error screens in Windows. These errors can indicate bad RAM but often point to something completely unrelated to RAM. This is especially true with intermittent problems. The challenge for techs is to recognize these errors and then determine which part of the system caused the memory error.

You can get two radically different types of parity errors: real and phantom. Real parity errors are simply errors that the MCC detects from the parity or ECC chips (if you have them). The operating system then reports the problem in an error message, such as “Parity error at xxxx:xxxxxxx,” where xxxx:xxxxxxx is a hexadecimal value (a string of numbers and letters, such as A5F2:004EEAB9). If you get an error like this, write down the value (Figure 6.34). A real parity/ECC error shows up at the same place in memory each time and almost always indicates that you have a bad RAM stick.

Phantom parity errors show up on systems that don’t have parity or ECC memory. If Windows generates parity errors with different addresses, you most likely do *not* have a problem with RAM. These phantom errors can occur for a variety of reasons, including software problems, heat or dust, solar flares, fluctuations in the Force...you get the idea.



• **Figure 6.34** Windows error message

System lockups and page faults (they often go hand in hand) in Windows can indicate a problem with RAM. A system lockup is when the computer stops functioning. A **page fault** is a milder error that can be caused by memory issues but not necessarily system RAM problems. Certainly page faults *look* like RAM issues because Windows generates frightening error messages filled with long strings of hexadecimal digits, such as “KRNL386 caused a page fault at 03F2:25A003BC.” Just because the error message contains a memory address, however, does not mean that you have a problem with your RAM. Write down the address. If it repeats in later error messages, you probably have a bad RAM stick. If Windows displays different memory locations, you need to look elsewhere for the culprit.

Every once in a while, something potentially catastrophic happens within the PC, some little electron hits the big red panic button, and the operating system has to shut down certain functions running before it can save data. This panic button inside the PC is called a **non-maskable interrupt (NMI)**, more simply defined as an interruption the CPU cannot ignore. An NMI manifests to the user as what techs lovingly call the **Blue Screen of Death (BSoD)**—a bright blue screen with a scary-sounding error message on it (Figure 6.35).

Bad RAM sometimes triggers an NMI, although often the culprit lies with buggy programming or clashing code. The BSoD varies according to the operating system, and it would require a much lengthier tome than this one to cover all the variations. Suffice it to say that RAM *could* be the problem when that delightful blue screen appears.

Finally, intermittent memory errors can come from a variety of sources, including a dying power supply, electrical interference, buggy applications, buggy hardware, and so on. These errors show up as lockups, general protection faults, page faults, and parity errors, but they never have the same address or happen with the same applications. Try the power supply first with non-application-specific intermittent errors of any sort.

Testing RAM

Once you discover that you may have a RAM problem, you have a couple of options. First, several companies manufacture hardware RAM-testing devices, but unless you have a lot of disposable income, they're probably priced way too high for the average tech (\$1,500 and higher). Second, you can use the method I use—*replace and pray*. Open the system case and replace each stick, one at a time, with a known good replacement stick. (You have one of those lying around, don't you?) This method, although potentially time-consuming, certainly works. With PC prices as low as they are now, you could simply replace the whole system for less than the price of a dedicated RAM tester.

Third, you could run a software-based tester on the RAM. Because you have to load a software tester into the memory it's about to scan, there's always a small chance that simply starting the software RAM tester might cause an error. Still, you can find some pretty good free ones out there. My favorite is the venerable Memtest86 written by Mr. Chris Brady (www.memtest86.com). Memtest86 exhaustively checks your RAM and reports bad RAM when it finds it (Figure 6.36).

```
A problem has been detected and windows has been shut down to prevent damage
to your computer.

The problem seems to be caused by the following file: SPCMDCON.SYS

PAGE_FAULT_IN_NONPAGED_AREA

If this is the first time you've seen this Stop error screen,
restart your computer. If this screen appears again, follow
these steps:

check to make sure any new hardware or software is properly installed.
If this is a new installation, ask your hardware or software manufacturer
for any windows updates you might need.

If problems continue, disable or remove any newly installed hardware
or software. Disable BIOS memory options such as caching or shadowing.
If you need to use Safe Mode to remove or disable components, restart
your computer, press F8 to select Advanced Startup options, and then
select Safe Mode.

Technical information:

*** STOP: 0x00000050 (0xFD3094C2,0x00000001,0xFBFE7617,0x00000000)

*** SPCMDCON.SYS - Address FBFE7617 base at FBFE5000, DateStamp 3d6dd67c
```

• Figure 6.35 Blue Screen of Death



A **general protection fault (GPF)** is an error that can cause an application to crash. Often they're caused by programs stepping on each other's toes. Chapter 17, "Maintaining and Troubleshooting Windows," goes into more detail on GPFs and other Windows errors.

```
Memtest86 v1.65 : Pass 11% ####
Athlon 64 (8.09) 2809 MHz : Test 12% ####
L1 Cache: 128K 16468MB/s : Test #4 [Moving inversions, random pattern]
L2 Cache: 512K 16468MB/s : Testing: 108K - 256M 256M
Memory : 256M 11224MB/s : Pattern: 94a989c0
Chipset: Intel i440BX

WallTime  Cached  RsvdMem  MemMap  Cache  ECC  Test  Pass  Errors ECC Errs
-----
0:00:05  256M    216K    e820-Std  on  off  Std    0    0

(ESC)Reboot (c)configuration (SP)scroll_lock (CR)scroll_unlock
```

• Figure 6.36 Memtest86 in action

Chapter 6 Review

■ Chapter Summary

After reading this chapter and completing the exercises, you should understand the following about RAM.

Identify the different types of DRAM packaging

- The sticks of RAM snapped into your motherboard are DRAM. Think of DRAM as a spreadsheet in which each cell holds a one or a zero. Each cell represents a single bit. The number of columns and rows are finite.
- If a chip contains 1,048,576 rows and eight columns, it can be described as a 1 M x 8 chip. A chip containing 2,097,152 rows and 16 columns can be described as a 2 M x 16 chip.
- DRAM needs to be able to fill the data bus. A 1 M x 8 DRAM chip is only 8 bits wide and would provide the CPU with only 8 bits of data. You would need eight of the 1 M x 8 chips to fill a 64-bit data bus.
- DRAM chips are soldered to a circuit board called a stick or module, and when these chips are combined, a single stick can fill the data bus. Modern DRAM sticks come in 32-bit- and 64-bit-wide varieties.

Explain the varieties of RAM

- Synchronous dynamic random access memory, or SDRAM, is tied to the system clock. The CPU, memory controller, and RAM all work in tandem with the clock pulse. This results in little wasted time when the CPU requests data from RAM.
- SDRAM comes in sticks called dual inline memory modules (DIMMs). SDRAM DIMMs come in a variety of pin sizes, with 168 pins being the most common for desktop systems. DIMMs are 64 bits wide so that they match the 64-bit data bus of modern CPUs. DIMMs snap into slots on the motherboard. Each DIMM is 64 bits, so snapping in a single DIMM fills the data bus; therefore, each slot is called a bank.
- SDRAM is advertised by capacity and speed. The speed is prefixed with “PC”; for example, PC66 runs at 66 MHz. A stick of SDRAM advertised as 256 MB PC133 has a capacity of 256 MB and runs at

133 MHz. Five clock speeds are commonly used: 66, 75, 83, 100, and 133 MHz.

- RAM speed must match or exceed that of the system. For example, PC66 RAM can only be used on a motherboard running at (or below) 66 MHz. A motherboard with a frontside bus of 100 MHz needs at least PC100 RAM.
- Rambus DRAM (RDRAM) was created by Rambus, Inc., and came in sticks called RIMMs. Due to cost, licensing issues, and compatibility, RDRAM is now obsolete.
- RIMMs were 64 bits wide and came in speeds of 600 MHz, 700 MHz, 800 MHz, and 1066 MHz. Differences in the notches prevented you from snapping a RIMM into a motherboard DIMM slot or vice versa.
- The RIMM memory controller uses a dual-channel architecture: it alternates between two banks to increase the speed of data retrieval. You need to install RIMMs in pairs to accommodate this feature.
- All RIMM slots on the motherboard must be filled with either a RIMM or a CRIMM. A CRIMM, or continuity RIMM, is simply a stick with no memory chips on it.
- Double data rate SDRAM (DDR RAM) is faster than regular SDRAM because it doubles the throughput. Desktop PCs use 184-pin DDR DIMMs, whereas laptops use either 200-pin or 172-pin small-outline DIMMs (SO-DIMM). DDR RAM modules have notches that differ from both regular SDRAM and RDRAM, so DDR RAM only snaps into DDR slots on a motherboard.
- DDR RAM is advertised one of two ways. One way is with “DDR” followed by the speed; for example, DDR400 for a stick of RAM runs at 400 MHz. Because DDR RAM runs at twice the clock speed, DDR400 is meant for a 200-MHz frontside bus. DDR RAM is also advertised by “PC” followed by the throughput; for example, PC3200. The number is determined by multiplying the speed by eight (as all DDR sticks are 64 bits, or 8 bytes, wide). Thus, DDR400 has a throughput of 3200 bytes per second (400 MHz x 8 bytes). DDR400 and PC3200 are equivalent.

- Dual-channel DDR improves upon DDR by using the DIMMs in pairs, much like RDRAM. For systems to use dual-channel DDR, the motherboard and memory controller must support dual-channel DDR RAM, and the DDR DIMMs must be identical. Dual-channel DDR does not require dummy modules such as CRIMMs. Dual-channel slots are identified on a motherboard by their blue color. Some motherboards offer a third slot in black, but filling that third slot to increase your total system RAM also turns off the dual-channel feature.
- DDR2 is an improvement of DDR. DDR2 runs faster than DDR and uses a 240-pin form factor that is not compatible with DDR. You need a motherboard with DDR2 slots to use DDR2 RAM. DDR2 is rated by speed or throughput. Thus DDR2-400 is the same as PC2-3200 and is meant for a 200-MHz frontside bus.
- DDR3 boasts higher speeds, more efficient architecture, and around 30 percent lower power consumption than DDR2 RAM. DDR3 doubles the buffer of DDR2 from 4 bits to 8 bits, giving it a huge boost in bandwidth over older RAM. Some chipsets that support DDR3 also support triple-channel memory, which works a lot like dual-channel but with three sticks of RAM instead of two.
- RAM modules are either single-sided or double-sided. Single-sided sticks have chips on only one side, whereas double-sided sticks have chips on both sides. Most RAM is single-sided.
- Latency refers to the time lag between when the memory controller starts to fetch data from RAM and when RAM actually sends out the requested data. Another lag occurs when the memory controller asks for the next line from RAM. The shorter this lag time, the faster the system. CL2 indicates a shorter lag time, or lower latency, than CL3, so CL2 is faster. Check your motherboard documentation and get the lowest latency RAM your system will support.
- Error correction code (ECC) RAM improved upon parity RAM by being able to fix single-bit errors on the fly. ECC RAM is always slower than non-ECC RAM due to the overhead of the correcting code. Only high-end motherboards and memory controllers can use ECC RAM.

- Buffered or registered RAM adds a buffering chip to compensate for the electrical interference that can result from using more than four DIMMs on a single motherboard. The motherboard and memory controller must be designed specifically to support buffered or registered RAM. You are not likely to see it in a typical desktop system.

Select and install RAM

- Disk thrashing is constant hard drive activity symptomatic of insufficient RAM. It occurs when Windows repeatedly uses up all available RAM space and has to move data not immediately needed out of the RAM into a temporary file on the hard drive called a swap file or page file and then swap the data back into RAM when it is needed by the program. You can monitor the size of your swap file in the Task Manager.
- You must know what type of RAM (such as regular SDRAM, DDR, DDR2, or DDR3) your motherboard accepts before you purchase a RAM upgrade. You also need to know the maximum amount of RAM your motherboard supports and the maximum supported per slot.
- Though not required, it is good practice to make sure all sticks of RAM in any system are as close to identical as possible. Matching your RAM modules in technology, capacity, speed, and manufacturer lessens the chance of problems and incompatibility.
- The serial presence detect (SPD) chip on modern DIMMs automatically supplies all the information about the RAM to the system, such as the size, speed, ECC or non-ECC, registered or unregistered, and other details.
- To install SO-DIMMs in a laptop, you must remove a panel on the underside of the laptop or remove the keyboard to find the RAM slots. The SO-DIMM slides into the slot and snaps down into position. Unplug the laptop and remove the battery before attempting a RAM upgrade, and protect the RAM from ESD by wearing an anti-static wristband.

Perform basic RAM troubleshooting

- Symptoms of bad RAM include parity errors, system lockups, page faults, and other error screens in Windows. However, other failing components can cause similar problems. Bad RAM usually results in error screens displaying messages such as "Parity error at xxxx:xxxxxxx"

where `xxxx:xxxxxxxx` is a hexadecimal value such as `A5F2:004EEAB9`. A real parity error shows up in the same place in memory each time—if that hexadecimal code is always the same, you probably have bad RAM.

- Page faults result in error screens such as “`KRNL386` caused a page fault at `03F2:25A003BC`.” The process that caused the page fault (in this case, `KRNL386`) may change, but if the hexadecimal address is the same across numerous error screens, you probably have bad RAM.

- A non-maskable interrupt (NMI) results in a Blue Screen of Death (BSOD). Although BSODs are often blamed on bad RAM, they are more often caused by buggy application program code.

- If you suspect you have bad RAM and you don’t have a hardware RAM-testing device, swap one of the sticks in your system with a known good stick. If the system works, you’ve found the bad stick. If the system still has errors, replace the stick you removed and swap a different stick for the known good stick. Another option is to use a software RAM tester such as Memtest86.

■ Key Terms

bank (182)

Blue Screen of Death (BSOD) (196)

buffered/registered DRAM (189)

continuity RIMM (CRIMM) (183)

disk thrashing (190)

double data rate SDRAM (DDR SDRAM) (183)

double-sided RAM (187)

dual-channel architecture (183)

dual inline memory module (DIMM) (182)

dynamic random access memory (DRAM) (177)

error correction code RAM (ECC RAM) (188)

general protection fault (GPF) (197)

latency (188)

module (180)

non-maskable interrupt (NMI) (196)

page fault (196)

page file (190)

parity RAM (188)

Rambus DRAM (RDRAM) (182)

RIMM (183)

serial presence detect (SPD) (193)

single inline memory module (SIMM) (180)

single-sided RAM (187)

small-outline DIMM (SO-DIMM) (182)

stick (180)

swap file (190)

synchronous DRAM (SDRAM) (181)

■ Key Term Quiz

Use the Key Terms list to complete the sentences that follow. Not all terms will be used.

1. If your motherboard uses RDRAM, you must fill each slot with either a(n) _____ or a(n) _____.
2. If the LED for your hard drive stays on most of the time, your computer is suffering from _____, a sure sign you need to add more memory.
3. A special kind of memory stick for laptops is called a(n) _____.
4. Memory that makes two data accesses during each clock tick is called _____.
5. Unlike regular DRAM, _____ is tied to the system clock.
6. Unlike regular DRAM, _____ enables error checking and correcting.
7. The time lag between when the memory controller starts to fetch a line of data from RAM and when the RAM actually starts to deliver the data is known as _____.
8. Systems automatically detect new RAM by polling the module’s _____ chip.
9. Motherboards that support more than four sticks of RAM may require _____ to accommodate for the additional electrical hassles caused by the additional sticks.
10. Memory chips are soldered to a small circuit board called a(n) _____ or a(n) _____.

■ Multiple-Choice Quiz

1. What is the correct throughput of DDR-SDRAM, and what is the speed of PC1600 RAM?
 - A. 4 bytes per second, 133 MHz
 - B. 8 bytes per second, 200 MHz
 - C. 4 bits per second, 400 MHz
 - D. 8 bits per second, 200 MHz
2. How many sticks of RAM do you need to fill a bank in a computer that can use 168-pin DIMMs?
 - A. One
 - B. Two
 - C. Four
 - D. Eight
3. If you upgrade your memory but notice that the RAM count does not reflect the additional memory, what should you do?
 - A. Remove the RAM and try to reinstall it.
 - B. Restart the computer.
 - C. Return the memory because it's probably bad.
 - D. Go to Setup and configure the memory to reflect the new amount.
4. What does a non-maskable interrupt cause the CPU to produce?
 - A. The Blue Screen of Death
 - B. A parity error
 - C. Excessive heat
 - D. An incorrect memory count
5. What does the CPU use to access the system's RAM?
 - A. The system bus
 - B. The MMC
 - C. The address bus
 - D. The expansion bus
6. Which of the following statements is true about RDRAM?
 - A. It uses dual-channel architecture.
 - B. It offers speeds ranging from 200 MHz to 600 MHz.
 - C. It is less expensive than SDRAM.
 - D. It is used by AMD but not Intel processors.
7. Which of the following SDRAM speeds would not work with a 100-MHz motherboard?
 - A. 66 MHz
 - B. 100 MHz
 - C. 133 MHz
 - D. 200 MHz
8. Which of the following is a valid package size for DDR3 DIMMs?
 - A. 168-pin
 - B. 172-pin
 - C. 184-pin
 - D. 240-pin
9. If you are running Windows 2000, Windows XP, or Windows Vista, you can use all of the following methods to find out how much RAM is installed in your computer *except* _____.
 - A. From the Control Panel, select System and then the Hardware tab.
 - B. Use the Performance tab under the Task Manager.
 - C. Select Properties from My Computer/Computer.
 - D. With a newer keyboard, press the WINDOWS-PAUSE/BREAK keystroke combination.
10. What package does DDR-SDRAM use for desktop PCs?
 - A. 30-pin
 - B. 72-pin
 - C. 168-pin
 - D. 184-pin
11. What happens if you mix RAM sticks of different speeds?
 - A. Your computer will work fine as long as it uses dual-channel architecture.
 - B. Your computer may slow down.
 - C. Your computer will work fine if all the memory sticks are slower than the speed of the motherboard.
 - D. Your computer may lock up every few seconds or provide corrupted data.

12. Why is SDRAM faster than regular DRAM?
 - A. It makes two processes per clock cycle.
 - B. It runs synchronously with the system clock.
 - C. It uses dual-channel architecture.
 - D. It has fewer pins, resulting in fewer corrupt bits.
13. What is true about a double-sided DIMM?
 - A. It has memory chips on the front and back.
 - B. It can be installed forward or backwards.
 - C. It is twice as fast as a single-sided DIMM.
 - D. It has half the capacity of a quad-sided DIMM.
14. What is the minimum requirement for RAM for Windows XP?
 - A. 128 MB
 - B. 256 MB
 - C. 512 MB
 - D. 1 GB
15. What happens if you add two RAM sticks to your PC, and one has a bad SPD?
 - A. When your system boots, it will recognize both RAM sticks but will not register any special features (such as ECC) of the stick with the bad SPD.
 - B. When your system boots, it will only register the presence of the RAM stick with the good SPD.
 - C. When your system boots, it won't register the presence of the RAM stick with the bad SPD until you configure the RAM settings by using the Setup utility.
 - D. When you try to boot the system, you will get a POST error message and the system will not boot.

■ Essay Quiz

1. Celia tells you she just received a new motherboard for the system she is building. She is confused about the RAM slots. The motherboard has three of them, but two are blue and one is black. What can you tell her about the RAM slots on her motherboard?
2. Your cousin James recently bought a new computer with 512 MB of RAM and a Windows Vista operating system. He's complaining about how slow his new computer is, especially when he's working with his graphics applications. How can you convince him that he needs more memory?
3. Now that you've convinced your cousin James that he needs more memory for his new PC, how will you explain what kind he should buy and how he should install it? You may assume that he has the motherboard book and that his computer has a 166-MHz system bus and can support DDR2 RAM.
4. Your computer is acting funny. Sometimes you get an error message on the screen. Other times data seems to be corrupted. Sometimes the computer just locks up. You suspect that it may be bad memory. How can you find out whether a memory problem or something else is causing your trouble?

Lab Project

• Lab Project 6.1

To learn more about memory, go to the Web site www.kingston.com, select Memory Tools from the buttons at the top of the screen, and examine the “Ultimate Memory Guide.” This resource contains information about all aspects of computer memory. After using this guide, answer the following questions:

- 1 Why do memory prices vary so frequently?

- 2 What are the differences in tin- and gold-edged memory sticks, and how does one know which to choose when upgrading?
- 3 Describe the notches on a 30-pin SIMM, a 72-pin SIMM, a 168-pin DIMM, and a 184-pin DIMM. What function do the notches serve?