

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.



EXAM TIP 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!

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.


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 4-11
144-pin micro-DIMM
(photo courtesy of
Micron Technology,
Inc.)



(Figure 4-11), or 172-pin *micro-DIMM* packages; and the 72-pin, 144-pin, or 200-pin *small outline DIMM (SO DIMM)* form factors (Figure 4-12). 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.

Figure 4-12
A 168-pin DIMM
above a 144-pin
SO-DIMM



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 it had a clock speed that 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 4-13). 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.

Figure 4-13
RDRAM




NOTE 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.

RDRAM was greatly anticipated by the industry for years, but industry support for RDRAM proved less than enthusiastic due to significant delays in development, plus a price many times that of SDRAM. Despite this grudging support, almost 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 4-14 shows a CRIMM.

Figure 4-14
CRIMM



RDRAM offered dramatic possibilities for high-speed PCs, but ran into three roadblocks that betamaxed 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.



NOTE *Betamaxed* 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.

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.



NOTE 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.

DDR SDRAM for desktops comes in 184-pin DIMMs. These DIMMs match 168-pin DIMMs in physical size, but not in pin compatibility (Figure 4-15). The slots for the two types of RAM appear similar as well, but have different guide notches, making it impossible to insert either type of RAM into the other's slot. DDR SDRAM for laptops comes in either 200-pin SO-DIMMs or 172-pin micro-DIMMs (Figure 4-16).





Figure 4-15 DDR SDRAM

Figure 4-16
172-pin DDR
SDRAM micro-
DIMM (photo
courtesy of
Kingston/Joint
Harvest)



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

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 bytes

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 as both terms are commonly used interchangeably. Table 4-1 shows all the speeds for DDR—not all of these are commonly used.

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

Table 4-1 DDR Speeds

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

There’s one sure thing 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 DDR SDRAM. Dual-channel DDR motherboards use regular DDR sticks, although manufacturers often sell RAM in matched pairs, branding them as dual-channel RAM.

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 (Figure 4-17). If you populate the third slot, the system will use the full capacity of RAM installed, but turns off the dual-channel feature—and no, it doesn’t tell you! The dual slots are blue; the third slot is black, which you could clearly see if this weren’t a black-and-white photo.





Figure 4-17 An nForce motherboard showing the three RAM slots. The two slots bracketing the slim space can run as dual channel as long as you don't populate the third slot.

DDR2

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 now. 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) makes DDR2 runs much faster than regular DDR. DDR2 uses a 240-pin DIMM that's not compatible with DDR. You'll find motherboards running both single-channel and dual-channel DDR2 (Figure 4-18).

Figure 4-18
240-pin DDR2
DIMM



The following table shows some of the common 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


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 DDR2 SDRAM, comes in one of two types: single-sided and double-sided. As their name implies, single-sided sticks only have chips on one side of the stick. Double-sided sticks have chips on both sides (Figure 4-19). The vast majority of RAM sticks are single-sided, but there are plenty of double-sided sticks 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 4-19
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 DDR400 (Figure 4-20).

Figure 4-20

Why is one more expensive than the other?



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

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 5, "BIOS and CMOS," to learn how to make these adjustments (and how to recover if you make a mistake!).

From a tech's standpoint, you need to get the proper RAM for the system on which you're working. 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.




NOTE 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 require electricity to charge up before they can do their job. This is one of the aspects of latency.

Parity and ECC

Given the high speeds and phenomenal amount of data moved by the typical DRAM chip, it's possible that a RAM chip might occasionally give bad data to the memory controller. This doesn't necessarily mean that the RAM has gone bad. This could be an occasional 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 4-21). *Parity* RAM stored an extra bit of data (called the parity bit) that the MCC used to verify if 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.

Figure 4-21
Ancient parity
RAM stick



Today's PCs that need to watch for RAM errors use a special type of RAM called *error correction code (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, as ECC RAM is always slower than non-ECC RAM.




NOTE 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.

ECC DRAM comes in every DIMM package type. 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. There are plenty of techs out there with years of experience who've never even seen ECC RAM.

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 DRAMs are called *buffered* or *registered* DRAM (Figure 4-22). The DDR2 version is called *fully buffered*. (See the “Beyond A+” section for more details.)

Figure 4-22
Buffered RAM



Like ECC, you must have a motherboard with an MCC designed to use this type of DRAM. Rest assured that such a motherboard will have 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.