Kingston Tools

Ultimate Memory Guide

WHAT IS MEMORY?

INTRODUCTION

These days, no matter how much memory your computer has, it never seems to be quite enough. Not long ago, it was unheard of for a **PC (Personal Computer)**, to have more than 1 or 2 **MB (Megabytes)** of memory. Today, most systems require 128MB to run basic applications. And up to 512MB or more is needed for optimal performance when using graphical and multimedia programs.

As an indication of how much things have changed over the past two decades, consider this: in 1981, referring to computer memory, Bill Gates said, "640K (roughly 1/2 of a megabyte) ought to be enough for anybody."

For some, the memory equation is simple: more is good; less is bad. However, for those who want to know more, this reference guide contains answers to the most common questions, plus much, much more.

THE ROLE OF MEMORY IN THE COMPUTER

People in the computer industry commonly use the term "memory" to refer to RAM (Random Access Memory). A computer uses Ram to hold temporary instructions and data needed to complete tasks. This enables the computer's CPU (Central Processing Unit), to access instructions and data stored in memory very quickly.

A good example of this is when the CPU loads an application program - such as a word processing or page layout program - into memory, thereby allowing the application program to work as quickly and efficiently as possible. In practical terms, having the program loaded into memory means that you can get work done more quickly with less time spent waiting for the computer to perform tasks.

The process begins when you enter a command from your keyboard. The CPU interprets the command and instructs the hard drive to load the command or program into memory. Once the data is loaded into memory, the CPU is able to access it much more quickly than if it had to retrieve it from the hard drive.

This process of putting things the CPU needs in a place where it can get at them more quickly is similar to placing various electronic files and documents you're using on the computer into a single file folder or directory. By doing so, you keep all the files you need handy and avoid searching in several places every time you need them.

THE DIFFERENCE BETWEEN MEMORY AND STORAGE

People often confuse the terms **memory** and **storage**, especially when describing the amount they have of each. The term memory refers to the amount of RAM installed in the computer, whereas the term storage refers to the capacity of the

computer's hard disk. To clarify this common mix-up, it helps to compare your computer to an office that contains a desk and a file cabinet.

The file cabinet represents the computer's hard disk, which provides storage for all the files and information you need in your office. When you come in to work, you take out the files you need from storage and put them on your desk for easy access while you work on them. The desk is like memory in the computer: it holds the information and data you need to have handy while you're working.

Consider the desk-and-file-cabinet metaphor for a moment. Imagine what it would be like if every time you wanted to look at a document or folder you had to retrieve it from the file drawer. It would slow you down tremendously, not to mention drive you crazy. With adequate desk space - our metaphor for memory - you can lay out the documents in use and retrieve information from them immediately, often with just a glance.

Here's another important difference between memory and storage: the information stored on a hard disk remains intact even when the computer is turned off. However, any data held in memory is lost when the computer is turned off. In our desk space metaphor, it's as though any files left on the desk at closing time will be thrown away.

MEMORY AND PERFORMANCE

It's been proven that adding more memory to a computer system increases its performance. If there isn't enough room in memory for all the information the CPU needs, the computer has to set up what's known as a **virtual memory** file. In so doing, the CPU reserves space on the hard disk to simulate additional RAM. This process, referred to as "swapping", slows the system down. In an average computer, it takes the CPU approximately 200**ns (nanoseconds)** to access RAM compared to 12,000,000ns to access the hard drive. To put this into perspective, this is equivalent to what's normally a 3 1/2 minute task taking 4 1/2 months to complete!

Access time comparison between RAM and a hard drive.

MEMORY UPGRADE ON A PC: LIFE IS GOOD

If you've ever had more memory added to your PC, you probably noticed a performance improvement right away. With a memory upgrade, applications respond more quickly, Web pages load faster, and you can have more programs running simultaneously. In short, additional memory can make using your computer a lot more enjoyable.

MEMORY UPGRADE ON A SERVER: LIFE IS EVEN BETTER

Server performance has a huge impact on the performance of a network: if a server is performing poorly, everyone on the network "feels the pain." So, while a memory upgrade on an individual PC makes a big difference only for the person who uses it, a memory upgrade in a server has far-reaching effects and benefits everyone who accesses the server.

To better understand the benefits of increasing memory on a server, take a look at these benchmark results compiled by

Mindcraft, an independent testing organization specializing in quantitative research. Mindcraft's study focused on servers running Windows 2000 Advanced Server, Sun Solaris 9.0 and Red Hat Linux 8.0 on both web and DBMS server platforms. The objective of this research was to reveal how server performance might improve incrementally by adding more memory as an alternative to the installation of additional processors.

In conducting it's research, Mindcraft discovered the performance of Windows 2000, Solaris and Linux servers showed significant improvement through the use of additional memory leading to performance benefits of as much as 40% over servers with an additional processor and no increase in memory. The performance improvement could be as high as 1000% on dual processor DBMS systems with the memory expanded from 512MB to 4GB.

WEB SERVER TESTING

A web server is comprised generally of hardware, operating system, server software, TCP/IP protocols, application software and Web site content. A client or visitor views the content on the web server via HTTP request. Mindcraft used SPECweb99 to test web server performance. SPECweb99 measures a web servers performance while servicing a specific workload of static and dynamic page requests. This workload is intended to simulate the type of requests that an Internet Service Provider (ISP) would encounter on its web servers. The results from the web server tests are shown below.

Table 1: percentage performance improvement with increased memory - 1CPU (Windows 2000 Server)

Base memory size	Increase memory to		
	1GB	2GB	4GB
512MB	31%	46%	48%
1GB		11%	13%

Table 2: percentage performance improvement with increased memory - 2CPUs (Windows 2000 Server)

Base memory size	Increase memory to		
	1GB	2GB	4GB
512MB	37%	76%	90%
1GB		29%	30%

Table 3: percentage performance improvement with increased memory - 1CPU (Sun Solaris)

Base memory size	Increase memory to	
	2GB	4GB
1GB	51%	63%
2GB		63%

Table 4: percentage performance improvement with increased memory - 2CPUs (Sun Solaris)

Base memory size	Increase memory to	
	2GB	4GB
1GB	66%	82%
2GB		10%

Table 5: percentage performance improvement with increased memory - 1CPU (Linux)

Base memory size	Increase memory to		
	1GB	2GB	4GB
512MB	33%	48%	62%
1GB		11%	21%

Table 6: percentage performance improvement with increased memory - 2CPUs (Linux)

Base memory size	Increase memory to		
	1GB	2GB	4GB
512MB	53%	102%	125%
1GB		32%	47%

DBMS SERVER TESTING

Mindcraft used the OSDB (Open Source Database Benchmark) test suite to conduct the DBMS testing. The OSDB is based on the ANSI SQL Standard Scalable and Portable (AS3AP) benchmark for relational databases systems. The OSDB implements much of the AS3AP benchmark. It implements both single-user and multi-user tests for several types of data base management systems (DBMSs). The single-user tests focus on database loading and index creation functions. They also measure basic query and access method performance. The results from the DBMS test are below.

Table 7: percentage performance improvement with increased memory - 1CPU (Windows 2002 Advanced Server)

Base memory size	Increase memory to		
	1GB	2GB	4GB
512MB	9%	837%	1458%
1GB		762%	1334%

Table 8: percentage performance improvement with increased memory - 2CPUs (Windows 2000 Server)

Base memory size	Increase memory to		
	1GB	2GB	4GB
512MB	24%	1510%	1559%
1GB		1202%	1241%

Table 9: percentage performance improvement with increased memory - 1CPU (Sun Solaris)

Base memory size	Increase memory to	
	2GB	4GB
1GB	376%	439%
2GB		439%

Table 10: percentage performance improvement with increased memory - 2CPUs (Sun Solaris)

Base memory size	Increase memory to	
	2GB	4GB
1GB	775%	1010%
2GB		27%

Table 11: percentage performance improvement
with increased memory - 1CPU (Linux)

	Increase memory to		
Base memory size	1GB	2GB	4GB
512MB	50%	981%	1243%
1GB		620%	795%

Table 12: percentage performance improvement
with increased memory - 2CPUs (Linux)

	Increase memory to		
Base memory size	1GB	2GB	4GB
512MB	52%	1357%	1477%
1GB		858%	937%

CONCLUSION: MORE MEMORY EQUALS BETTER PERFORMANCE
The Mindcraft research focused on the performance of Web and DBMS servers, because applications response times in both of these areas are among the most visible to the end user. Nobody likes having to work with a slow web server and where database applications are involved, the benchmark results clearly show that more memory available means that the Operating System and DBMS can keep more data in main memory, eliminating or reducing the frequency of relatively slow disk accesses, which in turn lead to slower applications response times.

Kingston Tools

Ultimate Memory Guide

How Much Memory Do You Need?

Perhaps you already know what it's like to work on a computer that doesn't have quite enough memory. You can hear the hard drive operating more frequently and the "hour glass" or "wrist watch" cursor symbol appears on the screen for longer periods of time. Things can run more slowly at times, memory errors can occur more frequently, and sometimes you can't launch an application or a file without first closing or quitting another.

So, how do you determine if you have enough memory, or if you would benefit from more? And if you do need more, how much more? The fact is, the right amount of memory depends on the type of system you have, the type of work you're doing, and the software applications you're using. Because the right amount of memory is likely to be different for a desktop computer than for a server, we've divided this section into two parts - one for each type of system.

Memory Requirements For A Desktop Computer

If you're using a desktop computer, memory requirements depend on the computer's operating system and the application software you're using. Today's word processing and spreadsheet applications require as little as 32MB of memory to run. However, software and operating system developers continue to extend the capabilities of their products, which usually means greater memory requirements. Today, developers typically assume a minimum memory configuration of 256MB. Systems used for graphic arts, publishing, and multimedia call for at least 512MB of memory and it's common for such systems to require 1GB or more for best performance.

The chart on the next page provides basic guidelines to help you decide how much memory is optimal for your desktop computer. The chart is divided by operating system and by different kinds of work. Find the operating system you're using on your computer, then look for the descriptions of work that most closely match the kind of work you do.

WINDOWS® VISTA				
	Entry	**Mainstream**	**Performance**	**High-end**
Memory Requirements – Desktop	1GB – 3GB	2GB – 4GB	2GB – 8GB*	4GB tp System mas**
Memory Requirements – Notebook	1GB – 3GB	2GB – 4GB	2GB – 8GB*	4GB tp System mas**
At Home	– Email – Picture share – Web Surfing – Basic Office Products – Personal Finance	– Download and Manage Music – Manage and share photos – Interactive web surfing – Personal Finance – Basic creative software – Basic Office products	– Download and Manage Photos, Music, Movies and TV – Photo enhancement – Web Hosting – Games – Personal Finance – Full Suite of office products	– Movie/Photo/Music Editing – Digital Home – Gaming – Web Hosting – Personal Finance – Full Suite of office products Advanced Features
At Work	– Email – Web Surfing – Basic Office Products – Corporate Software (CRM, Accounting, Manufacturing)	– Email – Web Surfing – Full suite of office products (Word, Excel, PowerPoint, Project, etc.) – Corporate Software (CRM, Accounting, Manufacturing)	– Email – Web Surfing and Development – Graphic Design – Full suite of office products Advanced Features (Word, Excel, PowerPoint, Project, etc.) – Corporate Software (CRM, Accounting, Manufacturing)	– Software programming – Design Engineering – Advanced Web and Database Development – Production level editing – Sound Engineering – CAD based engineering programs

WINDOWS® XP PROFESSIONAL and HOME

Windows XP Professional and Home feature built in wireless networking support and excellent power management for laptop / notebook users. With enhanced support of USB and Firewire, and its Media Center applications Windows XP is the best choice for the digital home and office.

Baseline: 128MB - 256MB
Optimal: 512MB - 1GB

Administrative & Service	**Light-** Word processing, email, data-entry	128MB - 256MB
	Medium- Fax/communications, database administration, spreadsheets; >2 applications open at a time	128MB - 256MB
	Heavy- Complex documents, accounting, business graphics, presentation software, network connectivity	256MB - 512MB
Executives & Analysts	**Light-** Proposals, reports, spreadsheets, business graphics, databases, scheduling, presentations	128MB- 256MB
	Medium- Complex presentations, sales/market analysis, project management, Internet access	256MB - 512MB
	Heavy- Statistical applications, large databases, research/technical analysis, complex presentations, video conferencing	512MB - 1GB
Engineers & Designers	**Light-** Page layout, 2 - 4 color line drawings, simple image manipulation, simple graphics	256MB - 512MB
	Medium- 2D CAD, rendering, multimedia presentations, simple photo-editing, Web development	512MB - 1GB
	Heavy- Animation, complex photo-editing, real-time video, 3D CAD, solid modeling, finite element analysis	1GB - 2GB

WINDOWS® 2000 PROFESSIONAL

Windows 2000 Professional runs software applications faster. Notebook-ready and designed with the future in mind, Windows 2000 Professional allows users to take advantage of a full range of features today. Windows 2000 Professional is future-ready and promises to run today's and tomorrow's applications better.

Baseline: 128MB - 256MB
Optimal: 512MB - 1GB

Administrative & Service	**Light-** Word processing, email, data-entry	128MB - 256MB
	Medium- Fax/communications, database administration, spreadsheets; >2 applications open at a time	128MB - 256MB
	Heavy- Complex documents, accounting, business graphics, presentation software, network connectivity	256MB - 512MB
Executives & Analysts	**Light-** Proposals, reports, spreadsheets, business graphics, databases, scheduling, presentations	128MB - 256MB
	Medium- Complex presentations, sales/market analysis, project management, Internet access	256MB - 512MB
	Heavy- Statistical applications, large databases, research/technical analysis, complex presentations, video conferencing	512MB - 1GB
Engineers & Designers	**Light-** Page layout, 2 - 4 color line drawings, simple image manipulation, simple graphics	256MB - 512MB
	Medium- 2D CAD, rendering, multimedia presentations, simple photo-editing, Web development	512MB - 1GB
	Heavy- Animation, complex photo-editing, real-time video, 3D CAD, solid modeling, finite element analysis	1GB - 2GB

WINDOWS® 98

Windows 98 requires 16 - 32MB to run basic applications. Tests show 45 - 65% performance improvements at 64MB and beyond.

Baseline: 64MB - 128MB
Optimal: 128MB - 256MB

Students	**Light-** Word processing, basic financial management, email and other light Internet use.	64MB - 128MB
	Medium- Home office applications, games, Internet surfing, downloading images, spreadsheets,	128MB - 256MB

	presentations	
	Heavy- Multimedia use such as video, graphics, music, voice recognition, design, complex images	256MB - 384MB

Home User	**Light-** Word processing, basic financial management, email and other light Internet use	64MB - 128MB
	Medium- Home office applications, games, Internet surfing, downloading images, spreadsheets, presentations	128MB - 256MB
	Heavy- Multimedia use such as video, graphics, music, voice recognition, design, complex images	256MB - 384MB

LINUX

The Linux operating system is quickly gaining popularity as an alternative to Microsoft Windows. It includes true multitasking, virtual memory, shared libraries, demand loading, proper memory management, TCP/IP networking, and other features consistent with Unix-type systems.

Baseline: 128MB - 256MB
Optimal: 512MB - 1GB

Administrative & Service	**Light-** Word processing, email, data-entry	128MB - 256MB
	Medium- Fax /communications, database administration, spreadsheets; >2 applications open at a time	128MB - 256MB
	Heavy- Complex documents, accounting, business graphics, presentation software, network connectivity	256MB - 512MB

Executives & Analysts	**Light-** Proposals, reports, spreadsheets, business graphics, databases, scheduling, presentations	128MB - 256MB
	Medium- Complex presentations, sales/market analysis, project management, Internet access	256MB - 512MB
	Heavy- Statistical applications, large databases, research/technical analysis, complex presentations, video conferencing	512MB - 1GB

Engineers & Designers	**Light-** Page layout, 2 - 4 color line drawings, simple image manipulation, simple graphics	256MB - 512MB
	Medium- 2D CAD, rendering, multimedia presentations, simple photo-editing, Web development	512MB - 1GB
	Heavy- Animation, complex photo-editing, real-time video, 3D CAD, solid modeling, finite element analysis	1GB - 2GB

MACINTOSH™ OS X

The Macintosh operating system manages memory in substantially different ways than other systems. Still, System X users will find that 48MB is a bare minimum. When using PowerMac $^®$ applications with Internet connectivity, plan on a range between 64 and 128MB as a minimum.

Baseline: 128MB - 256MB
Optimal: 512MB - 1GB

Administrative & Service	**Light-** Word processing, email, data- entry	128MB - 256MB
	Medium- Fax /communications, database administration, spreadsheets; >2 applications open at a time	128MB - 256MB
	Heavy- Complex documents, accounting, business graphics, presentation software, network connectivity	128MB - 256MB

Executives & Analysts	**Light-** Proposals, reports, spreadsheets, business graphics, databases, scheduling, presentations	256MB - 512MB
	Medium- Complex presentations, sales/ market analysis, project management, Internet access	256MB - 512MB
	Heavy- Statistical applications, large databases, research/ technical analysis, complex presentations, video conferencing	256MB - 512MB

Engineers & Designers	**Light-** Page layout, 2 - 4 color line drawings, simple image manipulation, simple graphics	512MB - 1GB
	Medium- 2D CAD, rendering, multimedia presentations, simple photo-editing, Web development	512MB - 1GB
	Heavy-Animation, complex photo-editing, real- time video, 3D CAD, solid modeling,	1GB - 2GB

* Please Note: These figures reflect work done in a typical desktop environment. Higher-end workstation tasks may require up to 8GB. Naturally, a chart such as this evolves as memory needs and trends change. Over time, developers of software and operating systems will continue to add features and functionality to their products. This will continue to drive the demand for more memory. More complex character sets, like Kanji, may require more memory than the standard Roman based (English) character sets.

SERVER MEMORY REQUIREMENTS

How can you tell when a server requires more memory? Quite often, the users of the network are good indicators. If network-related activity such as email, shared applications, or printing slows down, they'll probably let their Network Administrator know. Here are a few proactive strategies that can be used to gauge whether or not a server has sufficient memory:

- Monitor server disk activity. If disk swapping is detected, it is usually a result of inadequate memory.

- Most servers have a utility that monitors CPU, memory, and disk utilization. Review this at peak usage times to measure the highest spikes in demand.

Once it's determined that a server does need more memory, there are many factors to consider when deciding on how much is enough:

What functions does the server perform (application, communication, remote access, email, Web, file, multimedia, print, database)?

Some servers hold a large amount of information in memory at once, while others process information sequentially. For example, a typical large database server does a lot of data processing; with more memory, such a server would likely run much faster because more of the records it needs for searches and queries could be held in memory - that is, "at the ready." On the other hand, compared to a database server, a typical file server can perform efficiently with less memory because its primary job is simply to transfer information rather than to process it.

What operating system does the server use?

Each server operating system manages memory differently. For example, a **network operating system (NOS)** such as the Novell operating system handles information much differently than an application-oriented system such as Windows NT. Windows NT's richer interface requires more memory, while the traditional Novell functions of file and print serving require less memory.

How many users access the server at one time?

Most servers are designed and configured to support a certain number of users at one time. Recent tests show that this number is directly proportional to the amount of memory in the server. As soon as the number of users exceeds maximum capacity, the server resorts to using hard disk space as virtual memory, and performance drops sharply. In recent studies with Windows NT, additional memory allowed an application server to increase by several times the number of users supported while maintaining the same level of performance.

What kind and how many processors are installed on the server?

Memory and processors affect server performance differently, but they work hand in hand. Adding memory allows more information to be handled at one time, while adding processors allows the information to be processed faster. So, if you add processing power to a system, additional memory will enable the processors to perform at their full potential.

How critical is the server's response time?

In some servers, such as Web or e-commerce servers, response time directly affects the customer experience and hence revenue. In these cases, some IT Managers choose to install more memory than they think they would ever need in order to accommodate surprise surges in use. Because server configurations involve so many variables, it's difficult to make precise recommendations with regard to memory. The following chart shows two server upgrade scenarios.

SERVER MEMORY MAP

WINDOWS® 2000 SERVER

Designed to help businesses of all sizes run better, Windows 2000 Server offers a manageable, reliable and internet-ready solution for today's growing enterprises. For optimal performance, consider adding more memory to take advantage of Windows 2000 Server's robust feature set. Windows 2000 Server is internet-ready and promises to run today's and tomorrow's applications better.

Baseline: 128MB
Optimal: 256MB - 1GB

Application Server	Houses one or more applications to be accessed over a wide user base	256MB - 4GB
Directory Server	Central Management of network resources	128MB - 1GB
Print Server	Distributes print jobs to appropriate printers	128MB - 512MB
Communication Server	Manages a variety of communications such as PBX, Voicemail, Email, and VPN	512MB - 2GB
Web Server	Internet and intranet solutions	512MB - 2GB
Database Server	Manages simple to complex databases of varying sizes	256MB - 4GB

LINUX®

Linux is a reliable, cost-effective alternative to traditional UNIX servers. Depending on the distribution, the Linux server platform features a variety of utilities, applications, and services.

Baseline: 64MB - 128MB
Optimal: 256MB - 1GB

Application Server	Houses one or more applications to be accessed over a wide user base	64MB - 4GB
Directory Server	Central Management of network resources	128MB - 1GB
Print Server	Distributes print jobs to appropriate printers	128MB - 512MB
Communication Server	Manages a variety of communications such as PBX, Voicemail, Email, and VPN	512MB - 2GB
Web Server	Internet and intranet solutions	512MB - 2GB
Database Server	Manages simple to complex databases of varying sizes	256MB - 4GB

* Please Note: These figures reflect work done in a typical server environment. Higher-end workstation tasks may require up to 4GB. Naturally, a chart such as this evolves as memory needs and trends change. Over time, developers of software and operating systems will continue to add features and functionality to their products. This will continue to drive the demand for more memory. More complex character sets, like Kanji, may require more memory than the standard Roman based (English) character sets.

Top of page

## Ultimate Memory Guide

A CLOSER LOOK

Memory comes in a variety of sizes and shapes. In general, it looks like a flat green stick with little black cubes on it. Obviously, there's a lot more to memory than that. The illustration below shows a typical memory module and points out some of its most important features.

WHAT MEMORY LOOKS LIKE

A closer look at a 168-pin SDRAM DIMM.

PCB(PRINTED CIRCUIT BOARD)
The green board that all the memory chips sit on is actually made up of several layers. Each layer contains traces and circuitry, which facilitate the movement of data. In general, higher quality memory modules use PCBs with more layers. The more layers a PCB has, the more space there is between traces. The more space there is between traces, the lesser the chance of noise interference. This makes the module much more reliable.

DRAM (DYNAMIC RANDOM ACCESS MEMORY)
DRAM is the most common form of RAM. It's called "dynamic" RAM because it can only hold data for a short period of time and must be refreshed periodically. Most memory chips have black or chrome coating, or **packaging**, to protect their circuitry. The following section titled "Chip Packaging" shows pictures of chips housed in different types of chip packages.

CONTACT FINGERS
The contact fingers, sometimes referred to as "connectors" or "leads," plug into the memory socket on the system board, enabling information to travel from the system board to the memory module and back. On some memory modules, these leads are plated with tin while on others, the leads are made of gold.

INTERNAL TRACE LAYER
The magnifying glass shows a layer of the PCB stripped away to reveal the traces etched in the board. Traces are like roads the data travels on. The width and curvature of these traces as well as the distance between them affect both the speed and the reliability of the overall module. Experienced designers arrange, or "lay out", the traces to maximize speed and reliability and minimize interference.

CHIP PACKAGING
The term "chip packaging" refers to the material coating around the actual silicon. Today's most common packaging is called TSOP (Thin Small Outline Package). Some earlier chip designs used DIP (Dual In-line Package) packaging and SOJ (Small Outline J-lead). Newer chips, such as RDRAM use CSP (Chip Scale Package). Take a look at the different chip packages

below, so you can see how they differ.

DIP (DUAL IN-LINE PACKAGE)
When it was common for memory to be installed directly on the computer's system board, the DIP-style DRAM package was extremely popular. DIPs are through-hole components, which means they install in holes extending into the surface of the PCB. They can be soldered in place or installed in sockets.

SOJ (SMALL OUTLINE J-LEAD)
SOJ packages got their name because the pins coming out of the chip are shaped like the letter "J". SOJs are surface-mount components - that is, they mount directly onto the surface of the PCB.

TSOP (THIN SMALL OUTLINE PACKAGE)
TSOP packaging, another surface-mount design, got its name because the package was much thinner than the SOJ design. TSOPs were first used to make thin credit card modules for notebook computers.

CSP (CHIP SCALE PACKAGE)
Unlike DIP, SOJ, and TSOP packaging, CSP packaging doesn't use pins to connect the chip to the board. Instead, electrical connections to the board are through a BGA (Ball Grid Array) on the underside of the package. RDRAM (Rambus DRAM) chips utilize this type of packaging.

CHIP STACKING
For some higher capacity modules, it is necessary to stack chips on top of one another to fit them all on the PCB. Chips can be "stacked" either internally or externally. "Externally" stacked chip arrangements are visible, whereas "internally" stacked chip arrangements are not.

Example of externally stacked chips.

WHERE MEMORY COMES FROM

MAKING THE CHIP

Amazing but true: memory starts out as common beach sand. Sand contains silicon, which is the primary component in the manufacture of semiconductors, or "chips." Silicon is extracted from sand, melted, pulled, cut, ground, and polished into silicon wafers. During the chip-making process, intricate circuit patterns are imprinted on the chips through a variety of techniques. Once this is complete, the chips are tested and die-cut. The good chips are separated out and proceed through a stage called "bonding": this process establishes connections between the chip and the gold or tin leads, or pins. Once the chips are bonded, they're packaged in hermetically sealed plastic or ceramic casings. After inspection, they're ready for sale.

MAKING THE MEMORY MODULE

This is where memory module manufacturers enter the picture. There are three major components that make up a memory module: the memory chips, PCB, and other "on-board" elements such as resistors and capacitors. Design engineers use CAD (computer aided design) programs to design the PCB. Building a high-quality board requires careful consideration of the placement and the trace length of every signal line. The basic process of PCB manufacture is very similar to that of the memory chips. Masking, layering, and etching techniques create copper traces on the surface of the board. After the PCB is produced, the module is ready for assembly. Automated systems perform surface-mount and through-hole assembly of the components onto the PCB. The attachment is made with solder paste, which is then heated and cooled to form a permanent bond. Modules that pass inspection are packaged and shipped for installation into a computer.

Sand → Silicon Wafer → Chip → Memory Module → Computer

WHERE MEMORY GOES IN THE COMPUTER

Originally, memory chips were connected directly to the computer's motherboard or system board. But then space on the board became an issue. The solution was to solder memory chips to a small modular circuit board - that is, a removable module that inserts into a socket on the motherboard. This module design was called a SIMM (single in-line memory module), and it saved a lot of space on the motherboard. For example, a set of four SIMMs might contain a total of 80 memory chips and take up about 9 square inches of surface area on the motherboard. Those same 80 chips installed flat on the motherboard would take up more than 21 square inches on the motherboard.

These days, almost all memory comes in the form of memory modules and is installed in sockets located on the system motherboard. Memory sockets are easy to spot because they are normally the only sockets of their size on the board. Because it's critical to a computer's performance for information to travel quickly between memory and the processor(s), the sockets for memory are typically located near the CPU.

High-end Workstation Mini-Tower PC Notebook Printer

Examples of where memory can be installed.

MEMORY BANKS AND BANK SCHEMAS

Memory in a computer is usually designed and arranged in memory banks. A memory bank is a group of sockets or modules that make up one logical unit. So, memory sockets that are physically arranged in rows may be part of one bank or divided into different banks. Most computer systems have two or more memory banks - usually called bank A, bank B, and so on. And each system has rules or conventions on how memory banks should be filled. For example, some computer systems require all the sockets in one bank to be filled with the same capacity module. Some computers require the first bank to house the highest capacity modules. If the configuration rules aren't followed, the computer may not start up or it may not recognize all the memory in the system.

You can usually find the memory configuration rules specific to your computer system in the computer's system manual. You can also use what's called a memory configurator. Most third-party memory manufacturers offer free memory configu-rators available in printed form, or accessible electronically via the Web. Memory configurators allow you to look up your computer and find the part numbers and special memory configuration rules that apply to your system.

Kingston Technology's memory configurator includes "bank schema" drawings for different computer systems (a bank schema drawing depicts the sockets in the system), along with special instructions that list any unusual configuration rules that apply to the systems.
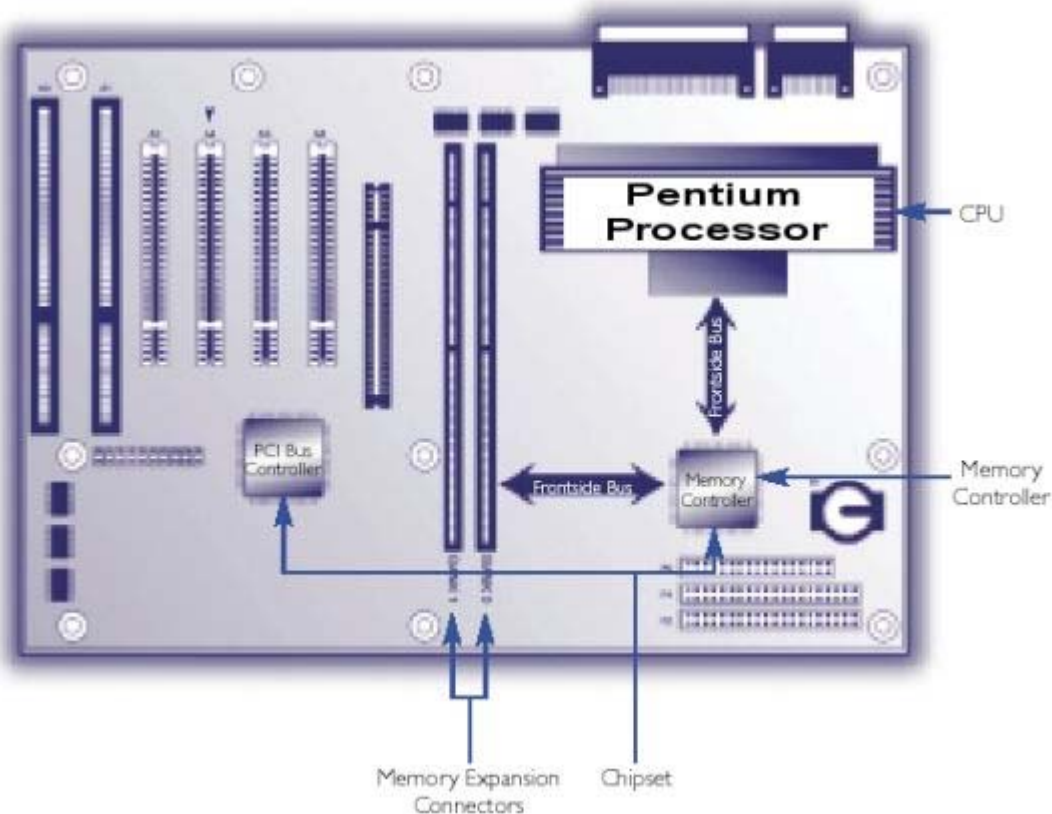
Kingston Tools

Ultimate Memory Guide

HOW MEMORY WORKS

Earlier, we talked about how memory holds information in a place where the CPU can get to it quickly. Let's look at that process in more detail.

HOW MEMORY WORKS WITH THE PROCESSOR

Main components of a computer system.

The **CPU** is often referred to as the brain of the computer. This is where all the actual computing is done.

The **chipset** supports the CPU. It usually contains several "controllers" which govern how information travels between the processor and other components in the system. Some systems have more than one chipset.

The **memory controller** is part of the chipset, and this controller establishes the information flow between memory and the CPU.

A **bus** is a data path in a computer, consisting of various parallel wires to which the CPU, memory, and all input/output devices are connected. The design of the bus, or **bus architecture**, determines how much and how fast data can move across the motherboard. There are several different kinds of busses in a system, depending on what speeds are required for those particular components. The **memory bus** runs from the memory controller to the computer's memory sockets. Newer systems have a memory bus architecture in which a **frontside bus (FSB)** runs from the CPU to main memory and a **backside bus (BSB)** which runs from the memory controller to L2 cache.

MEMORY SPEED

When the CPU needs information from memory, it sends out a request that is managed by the memory controller. The memory controller sends the request to memory and reports to the CPU when the information will be available for it to read.

This entire cycle - from CPU to memory controller to memory and back to the CPU - can vary in length according to memory speed as well as other factors, such as bus speed.

Memory speed is sometimes measured in **Megahertz (MHz)**, or in terms of **access time** - the actual time required to deliver data - measured in **nanoseconds (ns)**. Whether measured in Megahertz or nanoseconds, memory speed indicates how quickly the memory module itself can deliver on a request once that request is received.

ACCESS TIME (NANOSECONDS)

Access time measures from when the memory module receives a data request to when that data becomes available. Memory chips and modules used to be marked with access times ranging from 80ns to 50ns. With access time measurements (that is, measurements in nanoseconds), lower numbers indicate faster speeds.

In this example, the memory controller requests data from memory and memory reacts to the request in 70ns.The CPU receives the data in approximately 125ns. So, the total time from when the CPU first requests information to when it actually receives the information can be up to 195ns when using a 70ns memory module. This is because it takes time for the memory controller to manage the information flow, and the information needs to travel from the memory module to the CPU on the bus.

MEGAHERTZ (MHZ)

Beginning with Synchronous DRAM technology, memory chips had the ability to synchronize themselves with the computer's **system clock**, making it easier to measure speed in megahertz, or millions of cycles per second. Because this is the same way speed is measured in the rest of the system, it makes it easier to compare the speeds of different components and synchronize their functions. In order to understand speed better, it's important to understand the system clock.

SYSTEM CLOCK

A computer's **system clock** resides on the motherboard. It sends out a signal to all other computer components in rhythm, like a metronome. This rhythm is typically drawn as a square wave, like this:

In reality, however, the actual clock signal, when viewed with an oscilloscope, looks more like the example shown below.

Each wave in this signal measures one **clock cycle**. If a system clock runs at 100MHz, that means there are 100 million clock cycles in one second. Every action in the computer is timed by these clock cycles, and every action takes a certain number of clock cycles to perform. When processing a memory request, for example, the memory controller can report to the processor that the data requested will arrive in six clock cycles.

It's possible for the CPU and other devices to run faster or slower than the system clock. Components of different speeds simply require a multiplication or division factor to synchronize them. For example, when a 100MHz system clock interacts with a 400MHz CPU, each device understands that every system clock cycle is equal to four clock cycles on the CPU; they use a factor of four to synchronize their actions.

Many people assume that the speed of the processor is the speed of the computer. But most of the time, the system bus and other components run at different speeds.

MAXIMIZING PERFORMANCE

Computer processor speeds have been increasing rapidly over the past several years. Increasing the speed of the processor increases the overall performance of the computer. However, the processor is only one part of the computer, and it still relies on other components in a system to complete functions. Because all the information the CPU will process must be written to or read from memory, the overall performance of a system is dramatically affected by how fast information can travel between the CPU and main memory.

So, faster memory technologies contribute a great deal to overall system performance. But increasing the speed of the memory itself is only part of the solution. The time it takes for information to travel between memory and the processor is typically longer than the time it takes for the processor to perform its functions. The technologies and innovations described in this section are designed to speed up the communication process between memory and the processor.

CACHE MEMORY

Cache memory is a relatively small amount (normally less than 1MB) of high speed memory that resides very close to the CPU. Cache memory is designed to supply the CPU with the most frequently requested data and instructions. Because retrieving data from cache takes a fraction of the time that it takes to access it from main memory, having cache memory can save a lot of time. If the information is not in cache, it still has to be retrieved from main memory, but checking cache memory takes so little time, it's worth it. This is analogous to checking your refrigerator for the food you need before running to the store to get it: it's likely that what you need is there; if not, it only took a moment to check.

The concept behind caching is the "80/20" rule, which states that of all the programs, information, and data on your computer, about 20% of it is used about 80% of the time. (This 20% data might include the code required for sending or deleting email, saving a file onto your hard drive, or simply recognizing which keys you've touched on your keyboard.) Conversely, the remaining 80% of the data in your system gets used about 20% of the time. Cache memory makes sense because there's a good chance that the data and instructions the CPU is using now will be needed again.

HOW CACHE MEMORY WORKS
Cache memory is like a "hot list" of instructions needed by the CPU. The memory controller saves in cache each instruction the CPU requests; each time the CPU gets an instruction it needs from cache - called a "cache hit" - that instruction moves to the top of the "hot list." When cache is full and the CPU calls for a new instruction, the system overwrites the data in cache that hasn't been used for the longest period of time. This way, the high priority information that's used continuously stays in cache, while the less frequently used information drops out.

LEVELS OF CACHE
Today, most cache memory is incorporated into the processor chip itself; however, other configurations are possible. In some cases, a system may have cache located inside the processor, just outside the processor on the motherboard, and/or it may have a memory cache socket near the CPU, which can contain a cache memory module. Whatever the configuration, any cache memory component is assigned a "level" according to its proximity to the processor. For example, the cache that is closest to the processor is called **Level 1 (L1) Cache**, the next level of cache is numbered L2, then L3, and so on. Computers often have other types of caching in addition to cache memory. For example, sometimes the system uses main memory as a cache for the hard drive. While we won't discuss these scenarios here, it's important to note that the term cache can refer specifically to memory and to other storage technologies as well.

Main Memory

CPU with 16KB Internal Cache (Primary Level 1)

256KB External Cache (Secondary Level 2)

You might wonder: if having cache memory near the processor is so beneficial, why isn't cache memory used for all of main memory? For one thing, cache memory typically uses a type of memory chip called SRAM (Static RAM), which is more expensive and requires more space per megabyte than the DRAM typically used for main memory. Also, while cache memory does improve overall system performance, it does so up to a point. The real benefit of cache memory is in storing the most frequently-used instructions. A larger cache would hold more data, but if that data isn't needed frequently, there's little benefit to having it near the processor.

It can take as long as 195ns for main memory to satisfy a memory request from the CPU. External cache can satisfy a memory request from the CPU in as little as 45ns.

SYSTEM BOARD LAYOUT

As you've probably figured out, the placement of memory modules on the system board has a direct effect on system performance. Because local memory must hold all the information the CPU needs to process, the speed at which the data can travel between memory and the CPU is critical to the overall performance of the system. And because exchanges of information between the CPU and memory are so intricately timed, the distance between the processor and the memory becomes another critical factor in performance.

INTERLEAVING

The term **interleaving** refers to a process in which the CPU alternates communication between two or more memory banks. Interleaving technology is typically used in larger systems such as servers and workstations. Here's how it works: every time the CPU addresses a memory bank, the bank needs about one clock cycle to "reset" itself. The CPU can save processing time by addressing a second bank while the first bank is resetting. Interleaving can also function within the memory chips themselves to improve performance. For example, the memory cells inside SDRAM chip are divided into two independent **cell banks**, which can be activated simultaneously. Interleaving between the two cell banks produces a continuous flow of data. This cuts down the length of the memory cycle and results in faster transfer rates.

BURSTING

Bursting is another time-saving technology. The purpose of bursting is to provide the CPU with additional data from memory based on the likelihood that it will be needed. So, instead of the CPU retrieving information from memory one piece of at a time, it grabs a block of information from several consecutive addresses in memory. This saves time because there's a statistical likelihood that the next data address the processor will request will be sequential to the previous one. This way, the CPU gets the instructions it needs without having to send an individual request for each one. Bursting can work with many different types of memory and can function when reading or writing data.

Both bursting and pipelining became popular at about the same time that EDO technology became available. EDO chips that featured these functions were called "Burst EDO" or "Pipeline Burst EDO" chips.

PIPELINING

Pipelining is a computer processing technique where a task is divided into a series of stages with some of the work completed at each stage. Through the division of a larger task into smaller, overlapping tasks, pipelining is used to improve performance beyond what is possible with non-pipelined processing. Once the flow through a pipeline is started, execution rate of the instructions is high, in spite of the number of stages through which they progress.

Kingston Tools

Ultimate Memory Guide

HOW MUCH MEMORY IS ON A MODULE?

Up to now, we've discussed some of the technical attributes of memory and how memory functions in a system. What's left are the technical details - the "bits and bytes," as they say. This section covers the **binary numbering system**, which forms the basis of computing, and calculation of a memory module's capacity.

BITS AND BYTES

Computers speak in a "code" called **machine language**, which uses only two numerals: 0 and 1. Different combinations of 0s and 1s form what are called **binary numbers**. These binary numbers form instructions for the chips and microprocessors that drive computing devices - such as computers, printers, hard disk drives, and so on. You may have heard the terms "bit" and "byte." Both of these are units of information that are important to computing. The term **bit** is short for "binary digit." As the name suggests, a bit represents a single digit in a binary number; a bit is the smallest unit of information used in computing and can have a value of either 1 or a 0. A **byte** consists of 8 bits. Almost all specifications of your computer's capabilities are represented in bytes. For example, memory capacity, data-transfer rates, and data-storage capacity are all measured in bytes or multiples thereof (such as kilobytes, megabytes, or gigabytes).

This discussion of bits and bytes becomes very relevant when it comes to computing devices and components working together. Here, we'll address specifically how bits and bytes form the basis of measuring memory component performance and interaction with other devices like the CPU.

CPU AND MEMORY REQUIREMENTS

A computer's CPU (central processing unit) processes data in 8-bit chunks. Those chunks, as we learned in the previous section, are commonly referred to as bytes. Because a byte is the fundamental unit of processing, the CPU's processing power is often described in terms of the maximum number of bytes it can process at any given time. For example, Pentium and PowerPC microprocessors currently are 64-bit CPUs, which means they can simultaneously process 64 bits, or 8 bytes, at a time.

Each transaction between the CPU and memory is called a **bus cycle**. The number of data bits a CPU can transfer during a single bus cycle affects a computer's performance and dictates what type of memory the computer requires. Most desktop computers today use 168-pin DIMMs, which support 64-bit data paths. Earlier 72-pin SIMMs supported 32-bit data paths, and were originally used with 32-bit CPUs. When 32-bit SIMMs were used with 64-bit processors, they had to be installed in pairs, with each pair of modules making up a **memory bank**. The CPU communicated with the bank of memory as one logical unit.

Interestingly, RIMM modules, which are newer than DIMMs, use smaller 16-bit data paths; however they transmit information very rapidly, sending several packets of data at a time. RIMM modules use pipelining technology to send four 16-bit packets at a time to a 64-bit CPU, so information still gets processed in 64-bit chunks.

CALCULATING THE CAPACITY OF A MODULE

Memory holds the information that the CPU needs to process. The capacity of memory chips and modules are described in **megabits** (millions of bits) and **megabytes** (millions of bytes). When trying to figure out how much memory you have on a module, there are two important things to remember:

A module consists of a group of chips. If you add together the capacities of all the chips on the module, you get the total capacity of the module. Exceptions to this rule are:

- If some of the capacity is being used for another function, such as error checking.

- If some of the capacity is not being used, for example some chips may have extra rows to be used as back-ups. (This isn't common.)

While chip capacity is usually expressed in megabits, module capacity is expressed in megabytes. This can get confusing, especially since many people unknowingly use the word "bit" when they mean "byte" and vice versa. To help make it clear, we'll adopt the following standards in this book:

When we talk about the amount of memory on a module, we'll use the term "module capacity"; when we are referring to chips, we'll use the term "chip density". Module capacity will be measured in megabytes (MB) with both letters capital, and

chip density will be measured in megabits (Mbit), and we'll spell out the word "bit" in small letters.

COMPONENT	CAPACITY EXPRESSION	CAPACITY UNITS	EXAMPLE
Chips	Chip Density	Mbit (megabits)	64Mbit
Memory Modules	Module Capacity	MB (megabytes)	64MB

CHIP DENSITY

Each memory chip is a matrix of tiny cells. Each cell holds one bit of information. Memory chips are often described by how much information they can hold. We call this chip density. You may have encountered examples of chip densities, such as "64Mbit SDRAM" or "8M by 8". A 64Mbit chip has 64 million cells and is capable of holding 64 million bits of data. The expression "8M by 8" describes one kind of 64Mbit chip in more detail.

In the memory industry, DRAM chip densities are often described by their cell organization. The first number in the expression indicates the depth of the chip (in locations) and the second number indicates the width of the chip (in bits). If you multiply the depth by the width, you get the density of the chip. Here are some examples:

CURRENT AVAILABLE CHIP TECHNOLOGY

	CHIP DEPTH IN MILLIONS OF LOCATIONS	CHIP WIDTH IN BITS	CHIP DENSITY = DEPTH x WIDTH
16Mbit Chips			
4Mx4	4	4	16
1Mx16	1	16	16
2Mx8	2	8	16
16Mx1	16	1	16
64Mbit Chips			
4Mx16	4	16	64
8Mx8	8	8	64
16Mx4	16	4	64
128Mbit Chips			
8Mx16	8	16	128
16Mx8	16	8	128
32Mx4	32	4	128
256Mbit Chips			
32Mx8	32	8	256

MODULE CAPACITY

It's easy to calculate the capacity of a memory module if you know the capacities of the chips on it. If there are eight 64Mbit chips, it's a 512Mbit module. However, because the capacity of a module is described in megabytes, not megabits, you have to convert bits to bytes. To do this, divide the number of bits by 8. In the case of the 512Mbit module:

$$\frac{512\text{Mbits}}{8\text{ bits per byte}} = 64\text{MB}$$

You may hear standard memory modules in the industry being described as: "4Mx32" (that is, "4 Meg by 32"), or "16Mx64" ("16 Meg by 64"). In these cases, you can calculate the capacity of the module exactly as if it were a chip:

4Mx32 is 128Mbits.

$$\frac{128\text{Mbits}}{8\text{ bits per byte}} = 16\text{MB module}$$

16Mx64 is 1024Mbits.

$$\frac{1024\text{Mbits}}{8\text{ bits per byte}} = 128\text{MB module}$$

Here are some additional examples:

STANDARD MODULE TYPES

| STANDARD		MODULE DEPTH IN LOCATIONS	MODULE WIDTH IN DATA BITS	CAPACITY IN MBITS = DEPTH X WIDTH	CAPACITY IN MB = MBITS/8
72-Pin	1Mx32	1	32	32	4
	2Mx32	2	32	64	8
	4Mx32	4	32	128	16
	8Mx32	8	32	256	32
	16Mx32	16	32	512	64
	32Mx32	32	32	1024	128
168-Pin	2Mx64	2	64	128	16
	4Mx64	4	64	256	32
	8Mx64	8	64	512	64
	16Mx64	16	64	1024	128
	32Mx64	32	64	2048	256

As we mentioned earlier, there's only room for a certain number of chips on a PCB. Based on an industry standard 168-pin DIMM, the largest capacity module manufacturers can make using 64Mbit chips is 128MB; with 128Mbit chips, the largest module possible is 256MB; and with 256Mbit chips, the largest module possible is 512MB.

STACKING

Many large servers and workstations require higher capacity modules in order to reach system memory capacities of several gigabytes or more. There are two ways to increase the capacity of a module. Manufacturers can stack chips on top of one another, or they can stack boards.

CHIP STACKING
With chip stacking, two chips are stacked together and occupy the space that one chip would normally take up. In some cases, the stacking is done internally at the chip manufacturing plant and can actually appear to be one chip. In other cases the chips are stacked externally. The example below shows two externally stacked chips.

Example of externally stacked chips.

BOARD STACKING
As you might expect, board stacking involves putting two memory module printed circuit boards (PCBs) together. With board stacking, a secondary board mounts onto the primary board, which fits into the memory socket on the system motherboard.

Example of a stacked module.

Kingston Tools

Ultimate Memory Guide

DIFFERENT KINDS OF MEMORY

Some people like to know a lot about the computer systems they own - or are considering buying - just because. They're like that. It's what makes them tick. Some people never find out about their systems and like it that way. Still other people - most of us, in fact - find out about their systems when they have to - when something goes wrong, or when they want to upgrade it. It's important to note that making a choice about a computer system - and its memory features - will affect the experience and satisfaction you derive from the system. This chapter is here to make you smarter about memory so that you can get more out of the system you're purchasing or upgrading.

MODULE FORM FACTORS

The easiest way to categorize memory is by form factor. The form factor of any memory module describes its size and pin configuration. Most computer systems have memory sockets that can accept only one form factor. Some computer systems are designed with more than one type of memory socket, allowing a choice between two or more form factors. Such designs are usually a result of transitional periods in the industry when it's not clear which form factors will gain predominance or be more available.

SIMMS

The term **SIMM** stands for **Single In-Line Memory Module**. With SIMMs, memory chips are soldered onto a modular printed circuit board (PCB), which inserts into a socket on the system board.

The first SIMMs transferred 8 bits of data at a time. Later, as CPUs began to read data in 32-bit chunks, a wider SIMM was developed, which could supply 32 bits of data at a time. The easiest way to differentiate between these two different kinds of SIMMs was by the number of pins, or connectors. The earlier modules had 30 pins and the later modules had 72 pins. Thus, they became commonly referred to as 30-pin SIMMs and 72-pin SIMMs.

Another important difference between 30-pin and 72-pin SIMMs is that 72-pin SIMMs are 3/4 of an inch (about 1.9 centimeters) longer than the 30-pin SIMMs and have a notch in the lower middle of the PCB. The graphic below compares the two types of SIMMs and indicates their data widths.

4-1/4" 72-Pin SIMM

3-1/2" 30-Pin SIMM

Comparison of a 30-pin and a 72-pin SIMM

DIMMS

Dual In-line Memory Modules, or **DIMMs**, closely resemble SIMMs. Like SIMMs, most DIMMs install vertically into expansion sockets. The principal difference between the two is that on a SIMM, pins on opposite sides of the board are "tied together" to form one electrical contact; on a DIMM, opposing pins remain electrically isolated to form two separate contacts.

DIMMs come in various form factors and are specific to different DRAM technologies.

168-pin DIMM: EDO and PC66/100/133 SDRAM

184-pin DIMM: DDR 200/266/333/400 DDR SDRAM

240-pin DIMM: DDR2 400/533/667/800 DDR-2 SDRAM

DIMMs transfer 64 bits of data at a time and are typically used in computer configurations that support a 64-bit or wider memory bus. Some of the physical differences between DIMMs and 72-pin SIMMs include: the length of module, the number of notches on the module, and the way the module installs in the socket. Another difference is that many 72-pin SIMMs install at a slight angle, whereas DIMMs install straight into the memory socket and remain completely vertical in relation to the system motherboard. The illustration below compares a 168-pin DIMM to a 72-pin SIMM.

4-1/4" 72-Pin SIMM

5-1/4" 168-Pin DIMM

Comparison of a 72-pin SIMM and a 168-pin DIMM.

SO DIMMs

A type of memory commonly used in notebook computers is called SO DIMM or Small Outline DIMM. The principal difference between a SO DIMM and a DIMM is that the SO DIMM, because it is intended for use in notebook computers, is significantly smaller than the standard DIMM. The 72-pin SO DIMM is 32 bits wide and the 144-pin SO DIMM is 64 bits wide. 144-pin and 200-pin modules are the most common SO DIMMs today.

2.35" 72-pin SO DIMM   2.66" 144-Pin SO DIMM

Comparison of a 72-pin SO DIMM and a 144-pin SO DIMM.

**MicroDIMM** (Micro Dual In-Line Memory Module)

Smaller than an SO DIMM, MicroDIMMs are primarily used in sub-notebook computers. MicroDIMMs are available in 144-pin SDRAM, 172-pin DDR and 214-pin DDR2.

RIMMS AND SO-RIMMS

RIMM is the trademarked name for a Direct Rambus memory module. RIMMs look similar to DIMMs, but have a different pin count. RIMMs transfer data in 16-bit chunks. The faster access and transfer speed generates more heat. An aluminum sheath, called a **heat spreader**, covers the module to protect the chips from overheating.

A 184-pin Direct Rambus RIMM shown with heat spreaders pulled away.

16-Device Rambus RIMM

Heat Spreaders

An **SO-RIMM** looks similar to an SO DIMM, but it uses Rambus technology.

Heat Spreader

A 160-pin SO-RIMM module.

FLASH MEMORY

Flash memory is a solid-state, non-volatile, rewritable memory that functions like RAM and a hard disk drive combined. Flash memory stores bits of electronic data in memory cells, just like DRAM, but it also works like a hard-disk drive in that when the power is turned off, the data remains in memory. Because of its high speed, durability, and low voltage requirements, flash memory is ideal for use in many applications - such as digital cameras, cell phones, printers, handheld computers, pagers, and audio recorders.

Flash memory is available in many different form factors, including: CompactFlash, Secure Digital, SmartMedia, MultiMedia and USB Memory

PC CARD AND CREDIT CARD MEMORY

Before SO DIMMs became popular, most notebook memory was developed using proprietary designs. It is always more cost-effective for a system manufacturer to use standard components, and at one point, it became popular to use the same "credit card" like packaging for memory that is used on PC Cards today. Because the modules looked like PC Cards, many people thought the memory cards were the same as PC Cards, and could fit into PC Card slots. At the time, this memory was described as "Credit Card Memory" because the form factor was the approximate size of a credit card. Because of its compact form factor, credit card memory was ideal for notebook applications where space is limited.

PC Cards use an input/output protocol that used to be referred to as PCMCIA (Personal Computer Memory Card International Association). This standard is designed for attaching input/output devices such as network adapters, fax/modems, or hard drives to notebook computers. Because PC Card memory resembles the types of cards designed for use in a notebook computer's PC Card slot, some people have mistakenly thought that the memory modules could be used in the PC Card slot. To date, RAM has not been packaged on a PCMCIA card because the technology doesn't allow the processor to communicate quickly enough with memory. Currently, the most common type of memory on PC Card modules is Flash memory.

On the surface, credit card memory does not resemble a typical memory module configuration. However, on the inside you will find standard TSOP memory chips.

TECHNOLOGY TIMELINE

This section presents the most common memory technologies used for main memory: This road map offers an overview of the evolution of memory.

YEAR INTRODUCED	TECHNOLOGY	SPEED LIMIT
1987	FPM	50ns
1995	EDO	50ns
1997	PC66 SDRAM	66MHz
1998	PC100 SDRAM	100MHz
1999	RDRAM	800MHz
1999/2000	PC133 SRAM	133MHz (VCM option)
2000	DDR SDRAM	266MHz
2001	DDR SDRAM	333MHz
2002	DDR SDRAM	434MHz
2003	DDR SDRAM	500MHz
2004	DDR2 SDRAM	533MHz
2005	DDR2 SDRAM	800MHz
2006	DDR2 SDRAM	667 - 800MHz
2007	DDR3 SDRAM	1066 - 1333MHz

MAJOR CHIP TECHNOLOGIES

It's usually pretty easy to tell memory module form factors apart because of physical differences. Most module form factors can support various memory technologies so, it's possible for two modules to appear to be the same when, in fact, they're not. For example, a 168-pin DIMM can be used for EDO, Synchronous DRAM, or some other type of memory. The only way to tell precisely what kind of memory a module contains is to interpret the marking on the chips. Each DRAM chip manufacturer has different markings and part numbers to identify the chip technology.

FAST PAGE MODE (FPM)

At one time, FPM was the most common form of DRAM found in computers. In fact, it was so common that people simply called it "DRAM," leaving off the "FPM". FPM offered an advantage over earlier memory technologies because it enabled faster access to data located within the same row.

EXTENDED DATA OUT (EDO)

In 1995, EDO became the next memory innovation. It was similar to FPM, but with a slight modification that allowed consecutive memory accesses to occur much faster. This meant the memory controller could save time by cutting out a few steps in the addressing process. EDO enabled the CPU to access memory 10 to 15% faster than with FPM.

SYNCHRONOUS DRAM (SDRAM)

In late 1996, SDRAM began to appear in systems. Unlike previous technologies, SDRAM is designed to synchronize itself with the timing of the CPU. This enables the memory controller to know the exact clock cycle when the requested data will be ready, so the CPU no longer has to wait between memory accesses. SDRAM chips also take advantage of interleaving and burst mode functions, which make memory retrieval even faster. SDRAM modules come in several different speeds so as to synchronize to the clock speeds of the systems they'll be used in. For example, PC66 SDRAM runs at 66MHz, PC100 SDRAM runs at 100MHz, PC133 SDRAM runs at 133MHz, and so on. Faster SDRAM speeds such as 200MHz and 266MHz are currently in development.

DOUBLE DATA RATE SYNCHRONOUS DRAM (DDR SDRAM)

DDR SDRAM, is a next-generation SDRAM technology. It allows the memory chip to perform transactions on both the rising and falling edges of the clock cycle. For example, with DDR SDRAM, a 100 or 133MHz memory bus clock rate yields an effective data rate of 200MHz or 266MHz. Click here for more on DDR. Click here for more on DDR.

DOUBLE DATA RATE 2 SYNCHRONOUS DRAM (DDR2 SDRAM)

DDR2 is the second generation of Double Data Rate (DDR) SDRAM memory. It is an evolution of DDR memory technology that delivers higher speeds (up to 800 MHz), lower power consumption and heat dissipation. It is an ideal memory solution for bandwidth hungry systems and the lower power consumption is a perfect match for today's mobile users. Click here for more on DDR2.

DOUBLE DATA RATE 3 SYNCHRONOUS DRAM (DDR3 SDRAM)

DDR3 is the third generation of Double Data Rate (DDR) SDRAM memory. Similar to DDR2, it is a continuing evolution of

DDR memory technology that delivers higher speeds (up to 1600 MHz), lower power consumption and heat dissipation. It is an ideal memory solution for bandwidth hungry systems equipped with dual and quad core processors and the lower power consumption is a perfect match for both server and mobile platforms. DDR3 modules will be available in the second half of 2007.

DIRECT RAMBUS

Direct Rambus is a DRAM architecture and interface standard that challenges traditional main memory designs. Direct Rambus technology is extraordinarily fast compared to older memory technologies. It transfers data at speeds up to 800MHz over a narrow 16-bit bus called a Direct Rambus Channel. This high-speed clock rate is possible due to a feature called "double clocked," which allows operations to occur on both the rising and falling edges of the clock cycle. Also, each memory device on an RDRAM module provides up to 1.6 gigabytes per second of bandwidth - twice the bandwidth available with current 100MHz SDRAM.

In addition to chip technologies designed for use in main memory, there are also specialty memory technologies that have been developed for video applications.

MEMORY TECHNOLOGIES FOR VIDEO OR GRAPHICS PROCESSING

VIDEO RAM (VRAM)

VRAM is a video version of FPM technology. VRAM typically has two ports instead of one, which allows the memory to allocate one channel to refreshing the screen while the other is focused on changing the images on the screen. This works much more efficiently than regular DRAM when it comes to video applications. However, since video memory chips are used in much lower quantities than main memory chips, they tend to be more expensive. So, a system designer may choose to use regular DRAM in a video subsystem, depending on whether cost or performance is the design objective.

WINDOW RAM (WRAM)

WRAM is another type of dual-ported memory also used in graphics-intensive systems. It differs slightly from VRAM in that its dedicated display port is smaller and it supports EDO features.

SYNCHRONOUS GRAPHICS RAM (SGRAM)

SGRAM is a video-specific extension of SDRAM that includes graphics-specific read/write features. SGRAM also allows data to be retrieved and modified in blocks, instead of individually. This reduces the number of reads and writes that memory must perform and increases the performance of the graphics controller by making the process more efficient.

BASE RAMBUS AND CONCURRENT RAMBUS

Before it even became a contender for main memory, Rambus technology was actually used in video memory. The current Rambus main memory technology is called Direct Rambus. Two earlier forms of Rambus are Base Rambus and Concurrent Rambus. These forms of Rambus have been used in specialty video applications in some workstations and video game systems like Nintendo 64 for several years.

Top of page

<< PREV   ULTIMATE MEMORY GUIDE HOME   NEXT >>

Ultimate Memory Guide

**OTHER MEMORY TECHNOLOGIES YOU MAY
HAVE HEARD ABOUT**

ENHANCED SDRAM (ESDRAM)
In order to increase the speed and efficiency of standard memory modules, some manufacturers have incorporated a small amount of SRAM directly into the chip, effectively creating an on-chip cache. ESDRAM is essentially SDRAM, plus a small amount of SRAM cache, which allows for burst operations of up to 200MHz. Just as with external cache memory, the goal of cache DRAM is to hold the most frequently used data in the SRAM cache to minimize accesses to the slower DRAM. One advantage of on-chip SRAM is that it enables a wider bus between the SRAM and DRAM, effectively increasing the bandwidth and speed of the DRAM.

FAST CYCLE RAM (FCRAM)
FCRAM, co-developed by Toshiba and Fujitsu, is intended for specialty applications such as high-end servers, printers, and telecommunications switching systems. It includes memory array segmentation and internal pipelining that speed random access and reduce power consumption.

SYNCLINK DRAM (SLDRAM)
Though considered obsolete today, SLDRAM was developed by a consortium of DRAM manufacturers as an alternative to Rambus technology in the late 1990s.

VIRTUAL CHANNEL MEMORY (VCM)
Developed by NEC, VCM allows different "blocks" of memory to interface inde-pendently with the memory controller, each with its own buffer. This way, different system tasks can be assigned their own "virtual channels," and information related to one function does not share buffer space with other tasks occurring at the same time, making operations more efficient.

ERROR CHECKING
Ensuring the integrity of data stored in memory is an important aspect of memory design. Two primary means of accomplishing this are **parity** and **error correction code (ECC)**.

Historically, **parity** has been the most commonly used data integrity method. Parity can detect - but not correct - single-bit errors. **Error Correction Code (ECC)** is a more comprehensive method of data integrity checking that can detect and correct single-bit errors.

Fewer and fewer PC manufacturers are supporting data integrity checking in their designs. This is due to a couple of factors. First, by eliminating support for parity memory, which is more expensive than standard memory, manufacturers can lower the price of their computers. Fortunately, this trend is complemented by the second factor: that is, the increased quality of memory components available from certain manufacturers and, as a result, the relative infrequency of memory errors.

The type of data integrity checking depends on how a given computer system will be used. If the computer is to play a critical role - as a server, for example - then a computer that supports data integrity checking is an ideal choice. In general:

- Most computers designed for use as high-end servers support ECC memory.
- Most low-cost computers designed for use at home or for small businesses support non-parity memory.

PARITY
When parity is in use on a computer system, one parity bit is stored in DRAM along with every 8 bits (1 byte) of data. The two types of parity protocol - odd parity and even parity - function in similar ways.

With normal parity, when 8 bits of data are written to DRAM, a corresponding parity bit is written at the same time. The value of the parity bit (either a 1 or 0) is determined at the time the byte is written to DRAM, based on an odd or even quantity of 1s. Some manufacturers use a less expensive "fake parity" chip. This chip simply generates a 1 or a 0 at the time the data is being sent to the CPU in order to accommodate the memory controller. (For example, if the computer uses odd parity, the fake parity chip will generate a 1 when a byte of data containing an even number of 1s is sent to the CPU. If the byte contains an odd number of 1s, the fake parity chip will generate a 0.) The issue here is that the fake parity chip sends an "OK" signal no matter what. This way, it "fools" a computer that's expecting the parity bit into thinking that parity checking is actually taking place when it is not. The bottom line: fake parity cannot detect an invalid data bit.

This table shows how odd parity and even parity work. The processes are identical but with opposite attributes.

	ODD PARITY	EVEN PARITY
Step 1	The parity bit will be forced to 1 (or turned "on") if its corresponding byte of data contains an even number of 1's.	The parity bit is forced to 0 if the byte contains an even number of 1's.
	If the byte contains an odd number of 1's, the parity bit is forced to 0 (or turned "off ").	The parity bit is forced to 1 if its corresponding byte of data contains an odd number of 1's.
Step 2	The parity bit and the corresponding 8 bits of data are written to DRAM.	(Same as for odd parity)
Step 3	Just before the data is sent to the CPU, it is intercepted by the parity circuit.	(Same as for odd parity)
	If the parity circuit sees an odd number of 1's, the data is considered valid. The parity bit is stripped from the data and the 8 data bits are passed on to the CPU.	Data is considered valid if the parity circuit detects an even number of 1's.
	If the parity circuit detects an even number of 1's, the data is considered invalid and a parity error is generated.	Data is invalid if the parity circuit detects an odd number of 1's.

Parity does have its limitations. For example, parity can detect errors but cannot make corrections. This is because the parity technology can't determine which of the 8 data bits are invalid.

Furthermore, if multiple bits are invalid, the parity circuit will not detect the problem if the data matches the odd or even parity condition that the parity circuit is checking for. For example, if a valid 0 becomes an invalid 1 and a valid 1 becomes an invalid 0, the two defective bits cancel each other out and the parity circuit misses the resulting errors. Fortunately, the chances of this happening are extremely remote.

ECC

Error Correction Code is the data integrity checking method used primarily in high-end PCs and file servers. The important difference between ECC and parity is that ECC is capable of detecting and correcting 1-bit errors. With ECC, 1-bit error correction usually takes place without the user even knowing an error has occurred. Depending on the type of memory controller the computer uses, ECC can also detect rare 2 bit memory errors. While ECC can detect a multiple-bit errors, it cannot correct them. However, there are some more complex forms of ECC that can correct multiple bit errors.

Using a special mathematical sequence, algorithm, and working in conjunction with the memory controller, the ECC circuit appends ECC bits to the data bits, which are stored together in memory. When the CPU requests data from memory, the memory controller decodes the ECC bits and determines if one or more of the data bits are corrupted. If there's a single-bit error, the ECC circuit corrects the bit. In the rare case of a multiple-bit error, the ECC circuit reports a parity error.

Memory Scrubbing

Memory scrubbing is a feature initially implemented by most major server OEMs when ECC DIMMs first became available. Memory scrubbing refers to a process that actively reads memory during idle periods to search for and correct errors in memory. Therefore, the entire memory subsystem would be periodically checked and cleansed as the scrubbing process repeated itself over and over. If non-correctable memory errors are found the server management system would be alerted as to which DIMM was causing the errors.

Chipkill

Chipkill is a highly advanced error correction method much more effective than standard ECC correction. Chipkill provides error correction for up to four bits per DIMM. If too many memory errors are detected chipkill technology can take the inoperative chip offline while the server is still running to keep more errors from occurring. Chipkill support is provided in the memory controller and implemented using standard ECC DIMMs, so it is transparent to the OS.

OTHER SPECIFICATIONS

In addition to form factors, memory technologies, and error checking methods, there are several other specifications important to understanding and selecting memory products.

SPEED

The speed of memory components and modules is one of the most important factors in optimizing a memory configuration. In fact, all computer systems specify a memory component speed. Ensuring memory compatibility requires conforming to this specification. This section covers three measurements of memory component and module speed: access time, megahertz, and bytes per second.

ACCESS TIME

Prior to SDRAM, memory speed was expressed by access time, measured in nanoseconds (ns). A memory module's access time indicates the amount of time it takes the module to deliver on a data request. So, smaller numbers indicate faster access times. Typical speeds were 80ns, 70ns, and 60ns. Very often, you can identify the speed of a module by the part number on the chip: such part numbers end in "-6" for 60ns, "-7" for 70ns, and so on.

In most cases you can conform to a computer system's memory specification with a module rated at the required speed or faster. For example, if your system requires 70ns memory, you can use both 70ns and 60ns memory without a problem. However, some older systems check the module ID for the rated speed at system-boot up, and will only boot up if they recognize the exact speed they are looking for. If the system has an 80ns speed specification, for example, it won't accept anything different than 80ns, even if it is faster. In many cases, modules could still be built for these systems with faster memory chips on them, but the ID on the module would be set at the slower speed to insure compatibility with the system. This is why you can't always be sure of the rated speed on a module by looking at the speed markings on the memory chips.

MEGAHERTZ

Beginning with the development of SDRAM technology, memory module speed has been measured in megahertz (MHz). Speed markings on the memory chips them-selves are typically still in nanoseconds. This can be confusing, especially since these nanosecond markings no longer measure access time, but instead measure the number of nanoseconds between clock cycles. For SDRAM chips with speeds of 66MHz, 100MHz, and 133MHz, for example, the corresponding marking on the chips are -15, -10, and -8, respectively.

This table shows the method for determining speed equivalencies between MHz and ns ratings.

STEP 1	STEP 2	STEP 3	STEP 4
MHz = 1 million clock cycles per second	Multiply by 1 million to get **total clock cycles per second**	Constant: 1 billion nanoseconds per second	Divide **nanoseconds per second** (from Step 3) by **clock cycles per second** (from Step 2) to get **nanoseconds per clock cycle**
66	66,000,000	1,000,000,000	15
100	100,000,000	1,000,000,000	10
133	133,000,000	1,000,000,000	7.5

$$\frac{\text{nanoseconds per second}}{\text{clock cycles per second}} = \frac{1,000,000,000\text{ns}}{\text{clock cycles}} = \frac{\text{nanoseconds}}{\text{clock cycle}}$$

As noted in a previous section, the speed of the processor and the speed of the memory bus are normally not the same. The speed of memory is limited by the speed of the memory bus, which is the slowest link in the process.

BYTES PER SECOND

Converting MHz to bytes per second can be confusing at first. The two most important pieces of information you need to make the conversion is the speed (in MHz) and the width (in bits) of the bus.

Bus Width: If you have an 8-bit bus, then 8 bits, or 1 byte of information at a time can travel on the bus. If you have a 64-bit bus, then 64-bits, or 8 bytes of information can travel at a time.

Bus Speed: If the memory bus speed is 100MHz, this measures 100 million clock cycles per second. Typically, one packet of information can travel on each clock cycle. If the 100MHz bus is 1 byte wide, then data can travel at 100 megabytes per second. Data travels on a 100MHz, 64-bit bus at 800 megabytes per second.

Rambus modules are sometimes measured in MHz and sometimes measured in megabytes per second. One type of Rambus module runs on a 400MHz bus, but because Rambus modules can send two pieces of information per clock cycle instead of one, the module is rated at 800MHz. This is sometimes referred to as PC-800. Because the Rambus bus width is 16-bit, or 2 bytes wide, data travels at 1600MB per second, or 1.6GB per second. Using the same logic, PC-600 Rambus transfers data at 1.2 gigabytes per second.

REGISTERS AND BUFFERS

Registers and buffers improve memory operation by "re-driving" control signals in the memory chips. They can be external to the memory module, or they can be located on the module itself. Having registers and buffers placed directly on the memory module, enables a system to support a greater quantity of modules. So, you're likely to find these types of modules in servers and high-end workstations. It is important to note that when upgrading, unbuffered and buffered (or registered) modules cannot be mixed.

Buffering (EDO and FPM): For EDO and fast page modules, the process of re-driving the signals is called buffering. With

buffering there is no loss of performance.

Registering (SDRAM): For SDRAM, the signal driving process is called registering. Registering is similar to buffering, except that in registering, the data is clocked in and out of the register by the system clock. Registered modules are slightly slower than non-registered modules, because the registering process takes one clock cycle.

One of Two 16-Bit Drivers/Buffers

.925"   Buffered Module

An example of a buffered and a non-buffered module. They are keyed differently, to ensure that they can't be used in place of one another.

.964"   Non-Buffered Module

MULTIPLE-BANKED MODULES

A multiple-banked module allows more flexibility in type of chips used. Multiple banking allows a memory designer to divide the module into banks, which means it can appear to the computer system to be more than one module. This design is equivalent to the banks of memory sockets in a computer: the system accesses one bank of memory at a time, regardless of how many actual memory sockets comprise a bank.

Some people confuse the terms "double-sided" and "dual-banked". To clarify: **Double- sided** is a physical term meaning that chips are arranged on two sides of the memory module. **Dual-banked** is an electrical term meaning that the module is divided electrically into two memory banks.

TIN VERSUS GOLD

Memory modules are manufactured with either tin leads (connectors) or gold leads. Gold is a better conductor than tin. However, because tin is much less expensive than gold, computer manufacturers began using tin sockets on system boards in the early 1990s to reduce their costs. If you're buying memory and you have a choice - that is, compatible modules come in both tin and gold - it's best to match the metal of the module to the metal of the socket it will be going into. Matching metals can help avoid corrosion.

Kingston's policy has always been to match metals, so Kingston part numbers assigned to each computer system take the socket metal into account.

REFRESH RATES

Refresh is the process of recharging, or re-energizing, the "memory cells" in a memory chip. Internally, computer memory is arranged as a matrix of memory cells in rows and columns - like the squares on a checkerboard - with each column being further divided by the I/O width of the memory chip. The entire organization of rows and columns is called a **DRAM array**. DRAM is called "dynamic" RAM because it must be refreshed, or re-energized, thousands of times each second in order to retain data. It has to be refreshed because its memory cells are designed around tiny capacitors that store electrical charges. These capacitors work like very tiny batteries that lose their stored charges if they are not re-energized. Also, the process of reading data from the memory array drains these charges, so the memory cells must also be pre-charged before reading the data.

Cells are refreshed one row at a time (usually one row per refresh cycle).The term **refresh rate** refers not to the time it takes to refresh the memory but to the total number of rows that it takes to refresh the entire DRAM array. For example, a refresh rate of 2K indicates that it takes 2,048 rows to refresh the array; likewise, a 4K rate indicates 4,096 rows.

Normally, the system's memory controller initiates the refresh operation. But some chips are able to "self refresh." This means that the DRAM chip has its own refresh circuitry and does not require intervention from the CPU or external memory controller. Self-refresh modules dramatically reduce power consumption and are often used in portable computers.

CAS LATENCY

The term **CAS latency** refers to the number of clock cycles it takes before a column can be addressed on the DRAM chip. Latency is a measure of delay, so a "CL2" CAS latency factor indicates a two-clock cycle delay, and a "CL3" latency factor indicates a three-clock cycle delay. When SDRAM chips first came out, it was difficult to produce chips with a CAS latency factor as low as CL2. And although some specifications called for CL2, many modules worked fine at a CAS latency factor of CL3.

HEAT SPREADERS AND HEAT SINKS

As memory components get faster, chips become more dense and more circuits get squeezed onto smaller boards. Dissipation of excess heat becomes more of an issue. For several years now processors have incorporated fans. Newer memory module designs use heat sinks or heat spreaders to maintain safe operating temperatures.

SERIAL PRESENCE DETECT (SPD) AND PARALLEL PRESENCE DETECT (PPD)

When a computer system boots up, it must "detect" the configuration of the memory modules in order to run properly. Parallel Presence Detect is the traditional method of relaying the required information by using a number of resistors. PPD is the method SIMMs and some DIMMs use to identify themselves. Serial Presence Detect uses an **EEPROM (Electrically Erasable Programmable Read-Only Memory)** to store information about the module.

An EEPROM chip (also known as an E2PROM) differs from an EPROM in that it does not need to be removed from the computer to be modified. However, it does have to be erased and reprogrammed in its entirety, not selectively. It also has a limited lifespan - that is, the number of times it can be reprogrammed is limited.

NUMBER OF CLOCK LINES (2-CLOCK VERSUS 4-CLOCK)

SDRAM memory requires that clock lines run from the system clock to the memory module. "2-clock" means there are two clock lines running to the module, and "4-clock" means there are four clock lines running to the module. The first Intel designs were 2-clock because there were only eight chips on the module. Later, 4-clock designs were developed, which allowed for fewer chips per clock line, thereby decreasing the load on each line and enabling a quicker data interface.

VOLTAGES

Voltages on memory modules keep decreasing as memory cells in DRAMs get closer together and heat becomes more of an issue. Most computer systems used to operate at a standard of 5 volts. Compact notebook computers were the first to use 3.3-volt chips. This was not only because of heat issues; since lower voltage chips use less power, using them made it easier to prolong battery life. Now most desktops are standardized on 3.3-volt memory as well, but this is quickly being replaced by 2.5 voltage chips as products continue to get smaller and components closer together.

COMPOSITE VERSUS NON-COMPOSITE

Composite and non-composite were terms first used by Apple Computer to explain the difference between modules of the same capacity that used a different number of chips. To illustrate: when the industry is transitioning from one chip density to another, there is normally a stage where you can build, for example, a memory module with 8 of the new density chips, or 32 of the old density chips. Apple referred to the module using the latest technology and fewer chips as "non-composite", and the version using earlier technology and greater number of chips as "composite". Because 32 chips on a module can cause heat and spacing problems, Apple would often advise customers to buy non-composite modules.

Ultimate Memory Guide

WHAT TO CONSIDER WHEN BUYING MEMORY

One of the quickest and easiest ways to identify the right memory modules and expansion options for your computer is to use a memory configurator. Like some of the other major memory manufacturers, Kingston makes a memory configurator available. You can access Kingston's configurator through their home page at www.kingston.com.

The most important thing to ensure when buying memory is compatibility with your system. In addition, you'll need to decide how much memory you need and beyond that lie considerations of price, quality, availability, service, and warranty. This section helps you address these important decision factors and helps you answer questions like these:

- How much memory do I need?
- How much memory will my system recognize?
- What kind of memory is compatible with my system?
- How many sockets are open and how should I fill them?
- How do I determine the quality of memory?
- What should I know about memory prices?
- What other issues should I consider?

COMPATIBILITY

As mentioned earlier, compatibility of memory components with your computer system is arguably the most important factor to consider when upgrading memory. This section can get you started; it also makes frequent mention of the advantages of using a memory configurator.

WHAT KIND OF MEMORY IS COMPATIBLE WITH MY SYSTEM?

The easiest way to determine what type of memory goes with your system is to consult with your system documentation. If you need further assistance, consult a memory configurator available from many sources, including Kingston. Kingston and other brand-name memory companies offer such a tool to help you find the right memory configuration for your system.

With Kingston's configurator, you can search by five different criteria:

- System manufacturer/model
- Computer model name
- Memory module part number (Kingston, distributor, manufacturer)
- Specification
- Generic memory

To access Kingston's Memory Configurator, click here

WHAT IF I CAN'T FIND MY SYSTEM IN A MEMORY CONFIGURATOR?

If you can't find your system in the memory configuration programs, you can still find out what kind of memory you need by consulting the manual that came with your system. In most cases, the manual will provide basic specifications such as the speed and technology of the memory you need. This information is usually enough to choose a module by specification. If you don't feel you have enough information, you can call your system manufacturer or Kingston's toll-free technical support number for assistance.

HOW MANY SOCKETS DO I HAVE OPEN?

You may or may not have an idea what the inside of your computer looks like and how memory is configured. You may have opened up your computer when you bought it to see the configuration inside, or you may have looked at a configuration diagram in your user's manual. Even if you have no idea of the memory configuration of your system, you can use Kingston's memory configuration tools to find out. For each system, the configuration includes a diagram, called a **bank schema**, which indicates how the memory sockets are arranged in your system and what the basic configuration rules are. The simple tutorial on the next page outlines how to use a bank schema diagram to determine the number of sockets in your system and how to fill them.

HOW TO READ A BANK SCHEMA

A bank schema is a diagram of rows and columns that shows the number of memory sockets in your system. This diagram is a theoretical bank layout and not an actual system board layout; it is designed to help you quickly determine configuration requirements when adding memory modules.

In a bank schema, each [__] represents a **memory socket**:

Example: [__][__][__][__] = 4 memory sockets

Each column in the diagram represents a **memory bank**. The number of "[__]" symbols in a column represents the number of memory sockets in a bank. Upgrading is performed one bank at a time. For example, if there are four columns with two [__] in each column, upgrading is done two modules at a time. However, if there is just a single row of [__], upgrading may be performed one module at a time.

Examples:

8 sockets = [__][__][__][__][__][__][__][__]
(Modules may be installed one at a time in any combination)

8 sockets (4 banks of 2) = [___] [___] [___] [___]
[___] [___] [___] [___]
(Modules must be installed two at a time)

4 sockets (1 bank of 4) =
[___]
[___]   (Modules must be installed four
[___]   at a time)
[___]

The **standard memory** (base amount that the system was shipped with) appears in the diagram as either **removable** or **non-removable**.

Removable memory comes in the form of modules that fit into memory sockets, and, if desired, can be removed and replaced with modules of higher capacity. Removable memory is represented by a "▬▬" symbol with a number next to it: ▬▬ 4 [___] indicates that a 4MB module is in the first socket and that the second socket is empty.

Non-removable memory usually comes in the form of memory chips soldered directly onto the system board. It is represented in the bank schema in brackets: [_4MB_] indicates 4MB of non-removable memory soldered onto the board and two available memory sockets.

If your system is not included in the configurator, you may be able to find out how many sockets are in the system and how many are filled by pressing the F1 key during system startup. If your system supports this, a screen will appear that indicates how many memory sockets are in the system, which ones are filled and which are open, and what capacity modules are in each socket. If pressing the F1 key during startup doesn't produce this result, check your computer's system manual for more information.

As a last resort, you can open your computer and take a look at the sockets. (Important Note: Before removing the cover of your computer, refer to the computer's system manual and warranty information for instructions and other relevant information.) If you do open the computer, you may be able to identify "bank labels" that indicate whether memory are installed in pairs. Bank numbering typically begins with 0 instead of 1. So, if you have two banks, the first bank will be labeled "bank 0", and the second bank will be labeled "bank 1."

HOW SHOULD I FILL THE SOCKETS?

In most cases, it's best to plan your memory upgrade so you won't have to remove and discard the memory that came with the computer. The best way to manage this is to consider the memory configuration when you first buy the computer. Because lower-capacity modules are less expensive and more readily available, system manufacturers may achieve a base configuration by filling more sockets with lower-capacity modules. By way of illustration, consider this scenario: a computer system with 64MB standard memory comes with either two (2) 32MB modules or one (1) 64MB module. In this case, the second configuration is the better choice because it leaves more room for growth and reduces the chance that you'll have to remove and discard lower-capacity modules later. Unless you insist on the (1) 64MB module configuration, you may find yourself with only one socket left open for upgrading later.

Once you have purchased a computer and are planning your first upgrade, plan to buy the highest-capacity module you think you may need, especially if you only have one or two sockets available for upgrading. Keep in mind that, in general, minimum memory requirements for software applications double every 12 to 18 months, so a memory configuration that's considered large today will seem much less so a year from now.

QUALITY

As with any type of product, memory can vary in the quality from one manufacturer to another. In general, the larger, more established brand name companies are more consistent in adhering to tight design specifications, using high-quality components, and establishing certified quality control processes for manufacturing and thorough testing. That's not to say that lower-quality modules don't work fine - they may be the right solution, depending on how hard you work your system. In deciding the level of quality you require, consider the following:

1. If the memory you buy doesn't perform well, will you be comfortable returning it for a replacement? Would you have the time to deal with removing the memory and waiting a couple of days to a week to get the situation resolved?

2. When memory is of low quality, you often experience intermittent problems, such as the computer "freezing" unexpectedly, or having frequent errors. How often do you save your work, and if you were to lose your work, how much would that cost you? If you use your computer to play games, read email, and surf the Internet, such interruptions and losses may not be a big problem. But if you're running a business, losing a few hours of work could be a serious matter.

3. The biggest risk with unreliable memory is data corruption: that is, some bits of data may change or be read incorrectly. The result of data corruption could be as harmless as a syntax error in a document, or as potentially serious as a miscalculation in a spreadsheet. How important is the accuracy of the work you do on your computer? Again, if you use your computer for gaming, writing letters, and the Internet, it may not be a problem. But if you're managing your finances, you may want to do all you can to assure the reliability of your data.

4. Just like all products, the quality and durability you require depends on how you use it. Computer applications that require a lot of memory usually work the memory very hard. These applications often work better with memory that exceeds the system's speed and reliability specifications. If you're working in multimedia or using heavy number-crunching programs, the chance of a lower-quality memory module failing is greater than if you're only doing light work, such as simple word processing.

ASSESSING MEMORY QUALITY

Here are some important factors to keep in mind when assessing the quality of a brand of memory:

DESIGN

Designers of memory modules can follow strict specifications or take short cuts to save money on components. In general, manufacturers who do in-house design have more control over the quality of the module than those who farm this work out.

COMPONENTS

The quality of the DRAM chips, PC boards, and other components used on the module are critical to the overall quality of the module. Premium memory chips can cost up to 30% more than low-grade chips and high-quality PC boards cost about 50% more than lower quality alternatives.

ASSEMBLY

Many factors in module assembly can affect the overall quality of the module. In addition to proper handling of components, solder quality affects how reliably information can travel from the chip to the module and back. The temperature and humidity in the assembly and storage areas must be regulated to prevent warping, expansion, and contraction of components during assembly.

PROPER HANDLING

Electro-Static Discharge (ESD) is one of the most common causes of damage to a memory module. ESD damage can result from excessive and inappropriate handling. Memory modules should only be handled by workers who are properly "grounded" and modules should be properly packaged to protect against ESD during shipping.

TESTING

The more thoroughly memory has been tested before it is shipped, the less chance of problems during operation. In addition to standard production tests to ensure that the modules have been built correctly, memory can be tested for compatibility in the systems in which it will be used. The DRAM core can be tested for chip reliability, and modules can be tested "at speed" to make sure they will work in high-use situations. Some companies perform testing at all levels, and some do less testing.

PRICING AND AVAILABILITY

This section contains information that helps make sense of the fluctuations that can occur in the memory market.

THE DRAM CHIP MARKET

Memory modules are made with DRAM chips, which are manufactured in mass quantities in enormous fabrication plants

(often referred to as "fabs"). Fabs can take up to two years to build and require substantial capital investment: approximately $3 billion per plant. These time and cost factors directly affect on the ability of the memory market to adjust quickly to fluctuations in supply and demand. When demand for memory chips increases, chip manufacturers typically do not respond immediately because the investment required to add more capacity is substantial and may not pay off, especially if all the competitors are doing so at the same time. Therefore, the immediate effect is that prices rise as manufacturers assess whether the increase on demand is temporary or substantial enough to warrant investment. By the same token, when there is an oversupply situation in the market, chip manufacturers are willing to sustain losses for a long time while prices fall to below breakeven levels. This is because in many cases it costs more money to shut a plant down than to continue to produce and sell chips at below cost. Also, the longer a manufacturer can hold on, the greater the chance of "being there" to reap the rewards when competitors reduce capacity and the market turns around again.

WHY MEMORY PRICES FLUCTUATE

There are several factors that can affect memory prices. A few of these include: demand, DRAM manufacturing levels, inventory in the marketplace, time of the year, new operating system releases, and computer sales. All these things can affect memory prices at different times, either separately or simultaneously.

The most important thing to keep in mind when buying memory is that the price of 256MB today will most likely be different than the price of 256MB next quarter. The best rule of thumb is to compare memory prices close to your time of purchase. When doing price comparisons, it's more important to make sure you are making equivalent comparisons on module types than how the actual price per MB varies over time. If there are shortages in the market it's most important to be sure that what appears to be a "great deal" isn't a "short-cut" module built from off-spec components. In an oversupply market, you're much more likely to get a great price, but keep in mind that many manufacturers are losing money and may take shortcuts on testing and other expensive production quality measures to compensate. Refer to the quality section above for more details on this.

Kingston Tools

Ultimate Memory Guide

HOW TO INSTALL MEMORY

Congratulations! You have new memory for your computer. Now you have to install it. This chapter will guide you through the basics of memory module installation and refer you to resources that can help with problems.

BEFORE YOU INSTALL

Before you start, make sure you have the following:

1. Your computer manual. To install memory, you must open the computer box (chassis) and locate the memory sockets. You may need to unplug cables and peripherals, and re-install them afterward. The manual will most likely provide instructions specific to your computer.

2. A small screwdriver. Most computer chassis assemble with screws. The screwdriver also comes in handy if the notches on memory sockets are too tiny for your fingers.

IMPORTANT THINGS TO KEEP IN MIND

ESD DAMAGE

Electro-Static Discharge (ESD) is a frequent causes of damage to the memory module. ESD is the result of handling the module without first properly grounding yourself and thereby dissipating static electricity from your body or clothing. If you have a grounded wrist strap, wear it. If you don't, before touching electronic components - especially your new memory module - make sure you first touch an unpainted, grounded metal object. Most convenient is the metal frame inside the computer. In addition, always handle the module by the edges. If ESD damages memory, problems may not show up immediately and may be difficult to diagnose.

Wearing a grounded wrist strap can prevent ESD damage.

SWITCHING OFF THE POWER

Before opening the chassis, always power-off your computer and all attached peripherals. Leaving power on can cause permanent electrical damage to your computer and its components.

INSTALLING THE MEMORY

The vast majority of computers today have memory sockets that accept the following industry-standard memory modules:

Desktops, Workstations and Servers

-   72-pin SIMM
-   168-pin DIMM
-   184-pin RIMM

Notebooks and Mobile Computers

- 144-pin SO DIMM

Although sockets may be in different places on different computers, installation is the same. Consult the computer owner's manual to find out whether the memory sits on an expansion card or on the motherboard, and whether internal computer components must be moved to gain access.

In the section below are installation instructions for the standard modules listed above, followed by installation instructions for some of the more popular proprietary memory modules. If the computer requires proprietary memory, or the instructions below don't seem to apply to your situation, phone Kingston Technology's Technical Support Group at (800) 435-0640.

INSTALLING A 72-PIN SIMM

Installation of
a 72-pin SIMM.

1. Place your computer's power switch in the off position and disconnect the AC power cord.

2. Follow the instructions in your owner's manual that describe how to locate your computer's memory expansion sockets.

3. Before touching any electronic components or opening the package containing your new module(s), make sure you first touch an unpainted, grounded metal object to discharge any static electricity you may have stored on your body or clothing.

4. Handle your new module(s) carefully; do not flex or bend the module(s). Always grasp the module by its edges.

5. As shown in the illustration, the module and the expansion socket are keyed. A small plastic bridge in the socket must align with the curved notch in the module. The bridge ensures the module can only be plugged into the socket one way.

6. Insert the module into the socket at a slight angle. Make sure the module is completely seated in the socket. If you're having problems inserting the module into the socket, stop and examine both the module and the socket; make sure the notch in the module is properly aligned with the keyed plastic bridge in the socket. Do not force the module into the socket. If too much force is used, both the socket and module could be damaged.

7. Once you are satisfied the module is seated properly in the socket, rotate the module upward until the clips at each end of the expansion socket click into place.

8. After all modules have been installed, close the computer, plug in the AC power cord, and reinstall any other cables that may have been disconnected during the installation process.

INSTALLING A 168-PIN DIMM

1. Locate the memory expansion sockets on the computer's motherboard. If all the sockets are full, you will need to remove smaller capacity modules to allow room for higher capacity modules.

2. For some installations, DIMM memory can be installed in any available expansion slot. Other installations may require the memory to be installed in a particular sequence based on the module's capacity. Check your owner's manual to determine the correct installation sequence for your configuration.

3. Insert the module into an available expansion socket as shown in the illustration. Note how the module is keyed to the socket. This ensures the module can be plugged into the socket one way only. Firmly press the module into position, making certain the module is completely seated in the socket. Repeat this procedure for any additional modules you are installing.

4. Most 168-pin DIMM modules have ejector tabs similar to those shown in the illustration. The ejector tabs are used only when you need to remove a module. By pressing down on the ejector tabs, the module will pop up from the socket and it can be removed.

INSTALLING A 184-PIN RIMM

Installation of a 184-pin RIMM. RIMMs install into **RIMM connectors**. Currently, there are two connectors on the computer board, and each must contain a RIMM or a **C-RIMM (continuity RIMM)**. C-RIMMs do not contain memory devices. They are inexpensive pass-through modules that provide a continuous channel for the signal.

1. Turn off the computer and disconnect the AC power cord.

2. Locate your computer's memory expansion sockets by following the instructions in your owner's manual.

3. Before touching any electronic components, make sure you first touch an unpainted, grounded metal object to discharge any static electricity stored on your clothing or body.

4. If all the sockets are full, you will need to remove smaller capacity modules to allow room for higher capacity modules.

5. The ejector tabs shown in the illustration are used to remove a module. By pushing outward on the ejector tabs, the module will pop up from the socket and it can then be removed.

6. For most installations, Rambus modules can be installed in any available expansion socket, but any empty sockets must contain a continuity module as shown in the illustration. Note that some modes may use a specific installation sequence for Rambus modules (e.g. Rambus dual-channel configurations); see your owner's manual for more details.

7. Insert the module into an available expansion socket as shown in the illustration. Note how the module is keyed to the socket. This ensures the module can be plugged into the socket one way only. Firmly press the module into position, making certain the module is completely seated in the socket. The ejector tabs at each end of the socket will automatically snap into the locked position. Repeat this procedure for any additional modules are installing.

8. Once the module or modules have been installed, close the computer.

EXAMPLES OF NOTEBOOK MEMORY INSTALLATION

Although more and more notebooks accept standard SO DIMM memory modules, many still require proprietary modules having unique form factors. No standard dictates where the memory should be installed in the notebook. Because notebooks differ greatly in how memory is installed, you should consult your computer manual for specific instructions.This example shows a typical 144-pin SO DIM installation.

1. Always turn off your computer and remove the rechargeable battery pack before installing the memory.

2. Insert the module into the socket at a slight angle (approximately 30 degrees). Note that the socket and module are both keyed, which means the module can be installed one way only.

3. To seat the module into the socket, apply firm, even pressure to each end of the module (see the arrows) until you feel it slip down into the socket. If you are having problems getting the module to seat properly, try rocking the module up and down slightly, while continuing to apply pressure. When properly seated, the contact fingers on the edge of the module will almost completely disappear inside the socket.

4. With the module properly seated in the socket, rotate the module downward, as indicated in the illustration. Continue pressing downward until the clips at each end of the socket lock into position. With most sockets, you will hear a distinctive click, indicating the module is correctly locked into position.

Ultimate Memory Guide

TROUBLESHOOTING MEMORY PROBLEMS

COMMON MEMORY PROBLEMS

When you have a problem with memory, the cause is usually one of three things:

Improper Configuration: You have the wrong part for your computer or did not follow the configuration rules.

Improper Installation: The memory may not be seated correctly, a socket is bad, or the socket may need cleaning.

Defective Hardware: The memory module itself is defective.

The fact that many computer problems manifest themselves as memory problems makes troubleshooting difficult. For example, a problem with the motherboard or software may produce a memory error message.

This chapter is designed to help you figure out if you have a memory problem, and if so, what kind of problem it is, so you can get to a solution as quickly as possible.

BASIC TROUBLESHOOTING

The following basic steps apply to almost all situations:

1. **Make sure you have the right memory part for your computer.**
   At the manufacturer's Web site you can look up the part number. Many memory manufacturers have configurators, which indicate the compatibilities of your module. If not, phone the memory manufacturer, consult your computer manual, or phone the computer manufacturer.

2. **Confirm that you configured the memory correctly.**
   Many computers require module installation in banks of equal-capacity modules. Some computers require the highest capacity module to be in the lowest labeled bank. Other computers require that all sockets be filled; still others require single-banked memory. These are only a few examples of special configuration requirements. If you have a name-brand computer, visit Kingston's Web site (www.kingston.com) or use our upgrade manual to look up configuration rules specific to your computer. You can also contact technical support for your memory or computer manufacturer.

3. **Re-install the module.**
   Push the module firmly into the socket. In most cases you hear a click when the module is in position. To make sure you have a module all the way in the socket, compare the height of the module to the height of other modules in neighboring sockets.

4. **Swap modules.**
   Remove the new memory and see whether the problem disappears. Remove the old memory, reinstall the new, and see whether the problem persists. Try the memory in different sockets. Swapping reveals whether the problem is a particular memory module or socket, or whether two types of memory aren't compatible.

5. **Clean the socket and pins on the memory module.**
   Use a soft cloth to wipe the pins on the module. Use a PC vacuum or compressed air to blow dust off the socket. Do NOT use solvent, which may corrode the metal or prevent the leads from making full contact. Flux Off is a cleaner used specifically for contacts. You can purchase it at electronics or computer equipment stores.

6. **Update the BIOS.**
   Computer manufacturers update BIOS information frequently and post revisions on their Web sites. Make sure you have the most recent BIOS for your computer. This applies especially when you have recently installed new software or you are significantly upgrading memory.

WHEN THE PROBLEM OCCURS

When the problem occurs is a clue as to the cause.

For example, your response to a memory error message depends on whether:

1. You have just bought a new computer.
2. You have just installed new memory.
3. You have just installed new software or a new operating system.
4. You have just installed or removed hardware.
5. Your computer has been running fine and you've made no other recent changes.

Here are rules to get started:

YOU'VE JUST BOUGHT A NEW COMPUTER

If you have just purchased a new computer and it is producing memory errors, the problem could be related to anything, including a bad computer board. In this case, you need to troubleshoot the entire computer, including memory. Because the computer dealer will have configured memory and run system tests before shipping, they can best help.

YOU'VE JUST INSTALLED NEW MEMORY

If you have just installed new memory, the first possibility is that you installed incorrect parts. Double-check the part numbers. Confirm that you have configured and installed the memory correctly.

YOU'VE INSTALLED NEW SOFTWARE OR OPERATING SYSTEM

Newer software or operating systems tend to push memory harder than older operating systems. Sometimes memory that worked fine prior to a software installation begins producing errors once it runs memory-intensive software. New software also has bugs, and beta versions are notorious for producing memory errors. In these cases, your first step should be to ensure you have the latest BIOS and service patches for your software. Otherwise contact the memory vendor. A technical support representative may have experience with other software incidents and can walk you through more-detailed troubleshooting.

YOU'VE INSTALLED OR REMOVED HARDWARE

If you have just installed or removed hardware and suddenly receive memory error messages, the first place to look is in the computer itself. A connection may have come loose during the installation or the new hardware may be defective; in either case the errors are manifesting themselves as memory problems. Make sure you have the latest drivers and firmware. Most hardware manufacturers will post updates on their Web sites.

UNEXPECTED PROBLEMS

If your system has been running fine, but suddenly starts to produce memory errors, and crash or lock up frequently, the chance of a hardware failure is more likely, since configuration and installation problems show up as soon as the computer turns on. Sometimes you can get memory problems if your computer is overheating, if you are having a problem with your power supply, or if corrosion has developed between the memory module and the socket, weakening the connection.

HANDLING SPECIFIC PROBLEMS

Here is a list of the most common ways the computer informs you of a memory problem.

1. The computer won't boot, merely beeps.

2. Computer boots but doesn't recognize all the installed memory.

3. The computer boots but the screen is blank.

4. The computer reports a memory error.
   a. Memory mismatch error
   b. Memory parity interrupt at xxxxx
   c. Memory address error at xxxxx
   d. Memory failure at xxxxx, read xxxxx, expecting xxxxx
   e. Memory verify error at xxxxx

5. The computer has other problems caused by memory.
   a. The computer intermittently reports errors, crashes frequently, or spontaneously reboots.
   b. Registry Errors
   c. General-protection faults, page faults, and exception errors

6. The server system manager reports a memory error.

The following translations help you understand what the computer means when it gives you one of these signals.

1. **Computer won't boot, merely beeps.**
   Every time the computer starts, it takes inventory of hardware. Inventory consists of the computer BIOS recognizing, acknowledging, and in some cases, assigning addresses to, the components in the computer. If the computer won't boot, the CPU is unable to communicate with hardware. The cause can be improper installation or failure of the BIOS to recognize hardware. Follow basic troubleshooting, paying special attention to whether the memory module is completely installed and that you have the latest version of the BIOS.

2. **Computer boots but doesn't recognize all the installed memory.**
   When the computer boots, a part of the process is counting memory. On some machines the count appears on the screen and on others is masked. If the count is masked, from the computer set-up menu see how much memory the computer thinks it has. If the computer counts to or lists a number less than the memory you installed, the computer hasn't recognized all the memory.

   Sometimes the computer will recognize only part of a module. This is almost always due to using the wrong kind of memory. For example, if your computer accepts only single-banked memory and you have installed dual-banked, the computer will read only half the memory on the module. Sometimes the computer will accept only modules containing memory chips with specific organizations. For example, the VX chipset doesn't work well with 64 Mbit chips.

   In many computers the maximum amount of memory the computer can recognize is lower than the maximum amount you can physically install. For example, your computer may have three sockets, each of which can hold a 128MB module. If you filled every socket with 128MB, you would have 384MB of memory. However, your computer may recognize a maximum of 256MB. In most cases you can avoid this problem by consulting your computer manual or a memory configuration Web site before purchasing memory. Or visit the Kingston Web site.

3. **The computer boots but the screen is blank.**
   The most common reason for a blank screen is a dislodged card, memory not fully seated, or memory the computer doesn't support. Confirm that the memory is installed properly and that other components in the computer were not accidentally disconnected or dislodged while you installed memory.

   Double-check that you have the right part number for the computer. If you have nonparity memory in a computer that requires error-checking memory, or SDRAM memory in a computer that supports only EDO, the screen may be blank at boot up.

4. **The computer reports a memory error.**
   Memory mismatch error: This is not actually an error. Some computers require you to tell them that it's OK to have a new amount of memory. Use the set-up menu to tell the computer. Follow the prompts, enter the new amount, select Save, and exit.

   Computer memory or address errors:
   All of the following errors, and those similar to them, indicate that the computer has a problem with memory:

   - Memory parity interrupt at xxxxx
   - Memory address error at xxxxx
   - Memory failure at xxxxx, read xxxxx, expecting xxxxx
   - Memory verification error at xxxxx

   Typically the computer will perform a simple memory test as it boots. The computer will write information to memory and read it back. If the computer doesn't get what it was expecting, then it will report an error and sometimes give the address where the error occurred.

   Such errors normally indicate a problem with a memory module but can sometimes indicate a defective motherboard or incompatibility between old and new memory. To verify that the new memory is causing the problem, remove the new memory and see whether the problem goes away. Then remove the old memory and install only the new memory. If the error persists, phone the memory manufacturer and ask for a replacement.

5. **The computer has other problems caused by memory.**
   The Computer Intermittently Reports Errors, Crashes Frequently, or Spontaneously Reboots: Because of the large number of causes, these problems are difficult to diagnose. Possible causes are ESD (Electro-static Discharge), overheating, corrosion, or a faulty power supply. If you suspect ESD damage, contact the memory manufacturer and ask for a replacement. Before you install new memory, see page 85 for information on how to prevent ESD. If you suspect corrosion, clean the memory contacts and the sockets as explained on page 96. If you suspect the power supply, you will have to do overall computer troubleshooting with a focus on the power supply.

   Registry Errors: Windows writes a large portion of the registry to RAM. Sometimes defective memory will cause registry errors. Windows reports a registry error and prompts you to restart and restore. If the prompts repeat, remove your newly installed memory and restart the computer. If the errors disappear, ask the memory manufacturer for replacement modules.

General-Protection Faults, Page Faults, and Exception Errors: The most common cause is software. For example, one application may not have released the memory after quitting or occupies the same memory addresses as another. In these cases, rebooting should solve the problem. If the computer suddenly displays general-protection faults, exception errors, or page faults after you have installed new memory, remove the new memory and see whether the errors stop. If they occur only when the new memory is installed, contact the memory manufacturer for assistance.

6. **The server system manager reports a memory error.**
   Most servers ship with system managers that monitor component utilization and test for abnormalities. Some of these system managers count soft errors in memory. Soft errors have been corrected by ECC memory. If the rate of soft errors is higher than specifications, however, the system manager issues a pre-failure warning. This warning enables the network administrator to replace the memory and prevent system downtime.

   If the system manager on your server issues a pre-failure warning or other memory error, ask your memory manufacturer for a replacement. If the system manager continues to issue errors after memory replacement, make sure you have the latest BIOS, software service patches, and firmware. The chance of receiving two bad memory modules in a row is low. Contact the memory manufacturer for compatibility troubleshooting. Sometimes the server does not work well with certain types of memory chips or certain memory designs.

STILL NEED HELP?

Most memory manufacturers have FAQ or Q-and-A sections on their Web sites. There are also troubleshooting areas on the Web site of the computer manufacturer. If you can't find anything online, phone technical support for the computer or memory manufacturer. The next section gives information about Kingston and how to reach us.

Kingston Tools

Ultimate Memory Guide

MORE ABOUT KINGSTON

COMPANY OVERVIEW

Founded in 1987, Kingston Technology is the world's largest independent manufacturer of memory products. At the close of 2003, Kingston had sales in excess of $1.8 billion. Worldwide, the company presently staffs four manufacturing sites and six international sales and marketing facilities on four continents.

Kingston manufactures memory products that consistently meet or exceed the industry's stringent standards, possessing both the skill and equipment necessary to create the best modules in the marketplace.

The personal philosophy of company founders John Tu and David Sun, has been permanently imprinted into the corporate culture. From the company founders to the employees and ultimately, to the consumer, a special "magic" is infused into every product. At Kingston, it is the people that make the difference.

The Kingston spirit reflects an individual commitment...

CORE VALUES

Respect for one another in our culturally diverse environment

Loyalty to our long-term partnerships

Flexibility and **adaptability** in responding to our customers' needs

Investing in our employees to continuously improve our most valuable resource

Having fun working in the company of friends

WHY KINGSTON?

ISO 9001 CERTIFICATION

Kingston is ISO 9001 certified as part of a continuing global effort to assure customers the highest quality service and products.

EXCEPTIONAL QUALITY

In addition to our ISO 9001 certification, we have established an extensive DCAT (Design, Components, Assembly, and Test) quality-control process to ensure total product reliability.

PROVEN INDUSTRY LEADER

Kingston has been making memory modules since 1987.

100% PRODUCT TESTING

Kingston tests every memory module before shipping. Our customized test equipment is so comprehensive that it can test every cell on every chip on every module-on a 64MB module, that's 512 million cells!

PREMIUM COMPONENTS

Kingston buys 100% of its memory chips from top-tier DRAM suppliers.

FREE TECHNICAL SUPPORT

Technical support professionals are available to answer your call 24 hours a day, 7 days a week at (800) 435-0640. Average wait time is about 2 minutes.

KINGSTON'S LIFETIME WARRANTY

With a lifetime warranty on all Kingston memory upgrades, network adapters, and processor upgrades, you have the most comprehensive warranty in the industry.

FREE EVALUATION UNITS

Kingston offers a free 30-day product evaluation program to qualified corporate customers.

100% COMPATIBILITY GUARANTEE

Kingston guarantees all memory products to be compatible with the system or family of systems for which they're designed.

OVERNIGHT PRODUCT REPLACEMENT

In the unlikely event your Kingston product needs repair or replacement, we will cross-ship a replacement product to you overnight (in most cases).

FREE LITERATURE

Kingston is committed to educating customers on memory and related technology. White papers, memory bits, and other informative tools are available free of charge in electronic and printed formats. Send requests to webmaster@kingston.com.

CORPORATE ALLIANCE PROGRAM (CAP)

As an extension of Kingston's service commitment, we offer Corporate Alliance membership to qualified IT professionals. Benefits include a dedicated Kingston representative, use of the Partner Purchase Program, on site tech support training, and special pricing.

FIELD SUPPORT

Kingston's field sales team is available for regional and on-site customer support. To find a representative in your area, go to www.kingston.com/contact/field.asp.

Top of page

Kingston Tools

Ultimate Memory Guide

THE GLOSSARY

Accelerated Graphics Port

(AGP) - An Intel-developed interface that enables high-speed graphics. Graphics data move between the PC's graphics controller and computer memory directly, instead of being cached in video memory.

Access Time

The average time (in nanoseconds) for RAM to complete one access. Access Time is composed of address setup time and latency (the time it takes to initiate a Request for data and prepare access).

ANSI

(American National Standards Institute) - The U.S. organization responsible for setting information technology standards.

ASCII

(American Standard Code for Information Interchange) - A method of encoding text as binary values. The ASCII coding system contains 256 combinations of 7-bit or 8-bit binary numbers to represent every possible keystroke.

Backside Bus

(BSB) - The data path that runs between the CPU and L2 cache.

Bandwidth

The amount of data moved on electronic lines, such as a bus, per second. bandwidth is usually measured in bits per second, bytes per second, or cycles per Second (Hertz).

Bank

See memory bank.

Bank Schema

A method of diagramming memory configurations. The bank schema consists of rows and/or columns that represent memory sockets on a computer board. Rows indicate independent sockets; columns represent banks.

Base Rambus

The first generation of Rambus technology, first shipped in 1995.

BGA

(Ball Grid Array) - A chip package having solder balls on the underside for mounting. BGA allows for a reduction in die package size, better heat dissipation, and greater module densities.

Binary

A numbering system that uses combinations of 0 and 1 to represent data. Also known as Base 2.

BIOS

(Basic Input-Output System) - Startup routines that prepare the computer for operation.

Bit

The smallest unit of information a computer processes. A bit is 1 or 0.

Buffer

A holding area for data shared by devices that operate at different speeds or have different priorities. A buffer allows a device to operate without the delays that other devices impose.

Buffered Memory

A memory module that contains buffers. Buffers re-drive the signals through the memory chips and allow the module to include more memory chips. Buffered and unbuffered memory cannot be mixed. The design of the computer memory controller dictates whether memory must be buffered or unbuffered.

Burst EDO RAM

(BEDO) - EDO memory that can process four memory addresses in one burst. Bus speeds range from 50MHz to 66MHz (compared to 33MHz for EDO and 25MHz for Fast Page Mode).

Burst Mode

High-speed transmission of a block of data (a series of consecutive addresses) when the processor requests a single address.

Bus

A data path in a computer, consisting of various parallel wires to which the CPU, memory, and all input/output devices are connected.

Bus Cycle

A single transaction between main memory and the CPU.

Byte

Eight bits of information. The byte is the fundamental unit of computer processing; almost all specifications and measures of computer performance are in bytes or multiples thereof. See kilobytes and megabytes.

Cache Memory

A small amount (normally less than 1MB) of high-speed memory residing on or close to the CPU. Cache memory supplies the processor with the most frequently requested data and instructions. Level 1 cache (primary cache) is the cache closest to the processor. Level 2 cache (secondary cache) is the cache second closest to the processor and is usually on the motherboard.

CAS

(Column Address Strobe) - A memory chip signal that latches the column address of a particular location in a row-column matrix.

CAS Latency

The ratio between column access time and clock cycle time. CAS Latency 2 (CL2) offers a slight performance increase over CAS Latency 3 (CL3).

ccNUMA

(Cache-Coherent, Non-uniform Memory Access) - A flexible architecture that uses modular, low-cost components and offers multidimensional scaling potential to high-end servers.

Chipset

Microchips that support the CPU. The chipset usually contains several controllers that govern how information travels between the processor and other components.

Chip-Scale Package

(CSP) - Thin chip packaging whereby electrical connections are typically through a ball grid array. Chip-scale packaging is used in RDRAM and flash memory.

CompactFlash

A small, lightweight form factor for removable storage cards. CompactFlash cards are durable, operate at low voltages, and retain data when power is off. Uses include digital cameras, cell phones, printers, handheld computers, pagers, and audio recorders.

Composite

An Apple Computer, Inc. term for a memory module that used an older technology and contained more but lower-density chips.

Concurrent Rambus

The second generation of Rambus technology. Concurrent Rambus has been used in graphics-based computers, digital TVs, and video game applications (such as Nintendo 64 since 1997).

Continuity RIMM

(C-RIMM) - A Direct Rambus memory module that does not contain memory chips. C-RIMM provides a continuous channel for the signal. In a Direct Rambus system, open connectors must be populated with C-RIMMs.

CPU

(Central Processing Unit) - The computer chip that has primary responsibility for interpreting commands and running programs. The CPU is also known as the processor or microprocessor.

Credit Card Memory
A type of memory typically in laptop and notebook computers. Credit card memory is the size of a credit card.

DDR SDRAM
(Double Data Rate Synchronous Dynamic Random-Access Memory) - The latest generation of SDRAM technology. Data is read on both the rising and the falling edge of the computer clock, thereby delivering twice the bandwidth of standard SDRAM. With DDR SDRAM, memory speed doubles without increasing the clock frequency.

DIMM
(Dual In-line Memory Module) - A printed circuit board with gold contacts and memory devices. A DIMM is similar to a SIMM, but with this primary difference: unlike the metal leads on either side of a SIMM, which are "tied together" electrically, the leads on either side of a DIMM are electrically independent.

Direct Rambus
Rambus technology's third generation, which offers a completely new DRAM architecture for high-performance PCs. Data transfers at speeds up to 800MHz over a narrow 16-bit channel, compared to current SDRAM, which operates at 100MHz on a wide 64-bit bus.

DIP
(Dual In-line Package) - A DRAM component packaging. DIPs can be installed in sockets or permanently soldered into holes on the printed circuit board. The DIP package was extremely popular when memory was installed directly on the motherboard.

DRAM
(Dynamic Random-Access Memory) - The most common form of RAM. DRAM can hold data for only a short time. To retain data, DRAM must be refreshed periodically. If the cell is not refreshed, the data will disappear.

Dual-Banked
A memory module having two banks.

Dual Channel
Based on the dual channel system board designs, memory modules are installed on two separate channels, each with it's own access route to the memory controller. The dual memory channel design doubles the peak bandwidth of single memory channels. Best results occur when a matched-pair memory kit is used ensuring that the modules are identical.

Dual Independent Bus
(DIB) - An Intel-developed bus architecture that offers greater bandwidth by having two separate buses (frontside and backside) access the processor. Pentium II computers have DIBs.

ECC
(Error Correction Code) - A method of checking the integrity of data in DRAM. ECC provides more elaborate error detection than parity; ECC can detect multiple-bit errors and can locate and correct single-bit errors.

EDO
(Extended Data-Out) - A DRAM technology that shortens the read cycle between memory and CPU. On computers that support it, EDO memory allows a CPU to access memory 10 to 20 percent faster than comparable fast-page mode memory.

EDRAM
(Enhanced DRAM) - Enhanced Memory Systems, Inc. DRAM that contains a small amount of SRAM.

EEPROM
(Electrically Erasable Programmable Read-Only Memory) - A memory chip that retains data content after power has been removed. EEPROM can be erased and reprogrammed within the computer or externally.

EISA
(Extended ISA) - A bus architecture that extended the 16-bit ISA bus to 32 bits. EISA operates at 8MHz and has a peak data transfer rate of 33MB per second. EISA was introduced in 1988 as an open alternative to IBM's proprietary Micro Channel bus.

EOS
(ECC on SIMM) - An IBM data-integrity-checking technology that features ECC data-integrity-checking on a SIMM.

EPROM
(Erasable Programmable Read-Only Memory) - A programmable and reusable chip that retains content until erasure under ultraviolet light. Special equipment erases and reprograms EPROMs.

ESDRAM
(Enhanced Synchronous DRAM) - A type of SDRAM developed by Enhanced Memory Systems, Inc. ESDRAM replaces expensive SRAM in embedded systems and offers comparable speed with less power consumption and lower cost.

Even Parity
A type of data integrity checking whereby the parity bit checks for an even number of 1s.

Fast-Cycle RAM
(FCRAM) - FCRAM is a memory technology currently being developed by Toshiba and Fujitsu. FCRAM is not intended for PC main memory but will be used in specialty applications such as high-end servers, printers, and telecommunications switching systems.

Fast-Page Mode
An early form of DRAM, fast-page mode's advantage over previous page mode memory technologies was faster access to data in the same row.

Flash Memory
A solid-state, nonvolatile, rewritable memory that functions like a combination of RAM and hard disk. Flash memory is durable, operates at low voltages, and retains data when power is off. Flash memory cards are used in digital cameras, cell phones, printers, handheld computers, pagers, and audio recorders.

Form Factor
The size, configuration, and other specifications used to describe hardware. Examples of memory form factors are: SIMM, DIMM, RIMM, 30-pin, 72-pin, and 168-pin.

Frontside Bus
(FSB) - The data path that runs between the CPU and main memory (RAM).

Gigabit
Approximately 1 billion bits, or exactly 1 bit x 1,024^3 (1,073,741,824) bits.

Gigabyte
Approximately 1 billion bytes, or exactly 1 byte x 1,024^3 (1,073,741,824) bytes.

Heat Spreader
A sheath, usually aluminum, that covers an electronic device and dissipates heat.

Heat Sink
A component, typically zinc alloy, that dissipates heat. CPUs require heat sinks.

IC
(Integrated Circuit) - An electronic circuit on a semiconductor chip. The circuit includes components and connectors. A semiconductor chip is usually molded in a plastic or ceramic case and has external connector pins.

Interleaving
Techniques for increasing memory speed. For example, with separate memory banks for odd and even addresses, the next byte of memory can be accessed while the current byte refreshes.

JEDEC

    (Joint Electron Device Engineering Council) - An Electronic Industries Alliance (EIA) body that sets semiconductor engineering standards.

Kilobit

    Approximately one thousand bits, or exactly 1 bit x 2^{10} (1,024) bits.

Kilobyte

    Approximately one thousand bytes, or exactly 1 byte x 2^{10} (1,024) bytes.

Level 1 Cache

    (L1) - Also known as primary cache, L1 Cache is a small amount of high-speed memory that resides on or very close to the processor. L1 Cache supplies the processor with the most frequently requested data and instructions.

Level 2 Cache

    (L2) - Also known as secondary cache, L2 Cache is a small amount of high-speed memory close to the CPU and usually on the motherboard. L2 Cache supplies the processor with the most frequently requested data and instructions. Depending on the motherboard, Level 2 cache may be upgraded.

Logic Board

    See Motherboard.

Megabit

    Approximately one million bits, or exactly 1 bit x $1,024^2$ (1,048,576) bits.

Megabyte

    The most common term used to denote the capacity of a memory module. One megabyte equals approximately one million bytes, or exactly 1 byte x $1,024^2$ (1,048,576) bytes.

Memory

    A computer's random-access memory. Memory temporarily holds data and instructions for the CPU. See RAM.

Memory Bank

    A logical unit of memory in a computer, the size of which the CPU determines. For example, a 32-bit CPU requires memory banks that provide 32 bits of information at a time. A bank can consist of one or more memory modules.

Memory Bus

    The bus that runs from the CPU to the memory expansion slots.

Memory Controller Hub

    (MCH) - The interface between the processor, Accelerated Graphics Port, and RDRAM on motherboards that use Intel's 820 or 840 chipsets.

Memory Translator Hub

    (MTH) - The interface that allows SDRAM memory to be supported on a Direct Rambus Channel for motherboards using Intel's 820 chipset.

Micro BGA

    (µBGA) - Tessera, Inc. BGA chip packaging technique allows for a reduction in diepackage size, improved heat dissipation, and greater module densities.

MicroDIMM

    (Micro Dual In-Line Memory Module) - Smaller than an SODIMM, MicroDIMMs are primarily used in sub-notebook computers. MicroDIMMs are available in 144-pim SDRAM and 172-pin DDR.

Motherboard

    Also known as the logic board, main board, or computer board, the motherboard is the computer's main board and in most cases holds all CPU, memory, and I/O functions or has expansion slots for them.

Nanosecond

(ns) - One billionth of a second. Memory data access times are in nanoseconds. For example, memory access times for typical 30- and 72-pin SIMM modules range from 60 to 100 nanoseconds.

Nibble

Half of an 8-bit byte, or 4 bits.

Non-Composite

An Apple Computer, Inc. term for a memory module that used a new technology and contained fewer but higher-density chips. Non-composite modules were more reliable and more expensive than composite modules.

Return to Index

Odd Parity

Data integrity checking in which the parity bit checks for an odd number of 1s.

Return to Index

Parity

Data integrity checking that adds a single bit to each byte of data. The parity bit is used to detect errors in the other 8 bits.

PCB

(Printed Circuit Board) - Generally flat, multi-layer boards made of fiberglass with electrical traces. The surface and sublayers use copper traces to provide electrical connections for chips and other components. Examples of PCBs include: mother-boards, SIMMs, and credit card memory.

PC Card

(PCMCIA: Personal Computer Memory Card International Association) - A standard that allows interchangeability of various computing components on the same connector. The PCMCIA standard supports input-output devices, including memory, fax/modem, SCSI, and networking products.

PCI

(Peripheral Component Interconnect) - A peripheral bus that can send 32 or 64 bits of data simultaneously. PCI offers plug-and-play capability.

Pipeline Burst Cache

Cache that reduces wait states and accelerates memory access by using pipelining and bursting functions.

Pipelining

A technique in which memory loads the requested memory contents into a small cache composed of SRAM, then immediately begins fetching the next memory contents. This creates a two-stage pipeline, where data is read from or written to SRAM in one stage, and data is read from or written to memory in the other stage.

Proprietary Memory

Memory custom designed for a specific computer.

Return to Index

RAM

(Random-Access Memory) - A memory cell configuration that holds data for processing by a central processing unit (CPU). Random means the CPU can retrieve data from any address within RAM. See also Memory.

Rambus

(1) Rambus, Inc. develops and licenses high-performance memory logic and circuit design technology and provides licensees with product design, layout, and testing information. (2) Direct Rambus is a high-speed memory technology that uses a narrow 16-bit bus (Rambus channel) to transmit data at speeds up to 800MHz. See Rambus Channel.

Rambus Channel

The data path of Rambus systems. Because of the narrow data width (two bytes), Rambus modules transfer data at up to 800MHz.

RAS

A memory chip signal that latches the row address of a particular location in a row-column matrix.

Refresh
    Refreshing maintains data stored in DRAM. The process of refreshing electrical cells on a DRAM component is similar to recharging batteries. Different DRAM components require different refresh methods.

Refresh Rate
    The number of DRAM component rows that must be refreshed. Three common refresh rates are 2K, 4K and 8K.

Registered Memory
    SDRAM memory that contains registers directly on the module. The registers re-drive the signals through the memory chips and allow the module to be built with more memory chips. Registered and unbuffered memory cannot be mixed. The design of the computer memory controller dictates which type of memory the computer requires.

RIMM™
    The trademarked name for a Direct Rambus memory module. A RIMM™ conforms to the DIMM form factor and transfers data 16 bits at a time.

RIMM Connector
    A Direct Rambus memory socket.

SDRAM
    (Synchronous DRAM) - A DRAM technology that uses a clock to synchronize signal input and output on a memory chip. The clock is coordinated with the CPU clock so the timing of the memory chips and the timing of the CPU are in synch. Synchronous DRAM saves time in executing commands and transmitting data, thereby increasing the overall performance of the computer. SDRAM allows the CPU to access memory approximately 25 percent faster than EDO memory.

Self-Refresh
    A memory technology that enables DRAM to refresh on its own and independent of the CPU or external refresh circuitry. Self-Refresh technology is built into the DRAM chip itself and reduces power consumption dramatically. Notebook and laptop computers use this technology.

Serial Presence Detect
    An EEPROM chip that contains information about size and speed, as well as other specifications and manufacturer information of a memory module.

SGRAM
    (Synchronous Graphics Random-Access Memory) - Video memory that includes graphics-specific read/write features. SGRAM allows data to be retrieved and modified in blocks instead of individually. Blocking reduces the number of reads and writes the memory must perform and increases the performance of the graphics controller.

SIMM
    (Single In-line Memory Module) - A printed circuit board having memory devices and gold or tin/lead contacts. A SIMM plugs into a computer memory expansion socket. SIMMs offer two main advantages: ease of installation and minimal consumption of board surface. A vertically mounted SIMM requires only a fraction of the space required by a horizontally mounted DRAM. A SIMM may have as few as 30 or as many as 200 pins. On a SIMM, the metal leads on either side of the board are electrically tied together.

SIMM Socket
    A motherboard component that holds a single SIMM.

Single-Banked
    A module that has only one bank or row.

SLDRAM
    (Synclink) - Although obsolete today, SLDRAM was a main memory technology developed by a consortium of twelve DRAM manufacturers as an alternative to Direct Rambus technology.

SMART CARD
    An electronic device, similar in size to a credit card, that can store data and programs while enhancing security. Applications include identification, mass transit, and banking.

SO DIMM
    (Small-Outline Dual In-line Memory Module) - An enhanced version of a standard DIMM. A 72-pin small-outline DIMM is about half the length of a 72-pin SIMM. 144-pin and 200-pin modules are the most common SO DIMMs today.

SO-RIMM™
    The trademarked name for a Direct Rambus memory module in notebook computers. SO-RIMM™s provide memory

bandwidth comparable to desktop memory configurations.

SOJ

(Small-Outline J-lead) - A common form of surface-mounted DRAM packaging. An SOJ is a rectangular package with J-shaped leads on the two long sides.

Static RAM

(SRAM) - A memory chip that requires power to retain content. SRAM is faster than DRAM but more expensive and bulky. A typical use for SRAM is cache memory.

Storage

A data-holding device, such as a hard disk or CD-ROM.

Swapping

Using part of the hard drive as memory when RAM is full. See Virtual Memory.

System Board

See Motherboard.

Transmission Line Technology

A technology that supports the backside bus in Direct Rambus systems. Information is quickly pipelined in simultaneous packets. The memory controller reassembles the packets for frontside bus transfer and communication to the processor.

TSOP

(Thin Small-Outline Package) - A DRAM package that uses gull-wing leads on both sides. TSOP DRAM mounts directly on the surface of the printed circuit board. The TSOP package is one-third the thickness of an SOJ. TSOP components commonly occur in small-outline DIMMs and credit card memory.

Unbuffered Memory

Memory that does not contain buffers or registers located on the module. Instead, these devices are located on the motherboard.

VESA Local Bus

(VL-Bus) - A 32-bit local bus that runs between the CPU and peripheral devices at speeds up to 40MHz.

Virtual Channel Memory

(VCM) - VCM is a memory architecture developed by NEC. VCM allows different blocks of memory-each with its own buffer-to interface separately with the controller. This way, system tasks can be assigned their own virtual channels. Information related to one function does not share buffer space with other tasks running simultaneously, thereby making overall operations much more efficient.

Virtual Memory

Simulated memory. When RAM is full, the computer swaps data to the hard disk and back as needed. See Swapping.

VRAM

(Video Random-Access Memory) - Dual-ported (two separate data ports) memory typically on a video or graphics card. One port is dedicated to the CRT and refreshes and updates the image. The second port is for the CPU or graphics controller and changes the image data in memory.

Wait State

An inactive period for the processor. Wait states result from the different clock speeds of the processor and memory, the latter being typically slower.

Window Random Access Memory

(WRAM) - Samsung Electronics' dual-ported (two separate data ports) memory typically on a video or graphics card. WRAM has a 25% higher bandwidth than VRAM but costs less.