# IT PRESENTATION COMPUTER MEMORY

Performed by a student group Ktmo1-3
Reuka A.S.
Professor at the Department of foreign languages
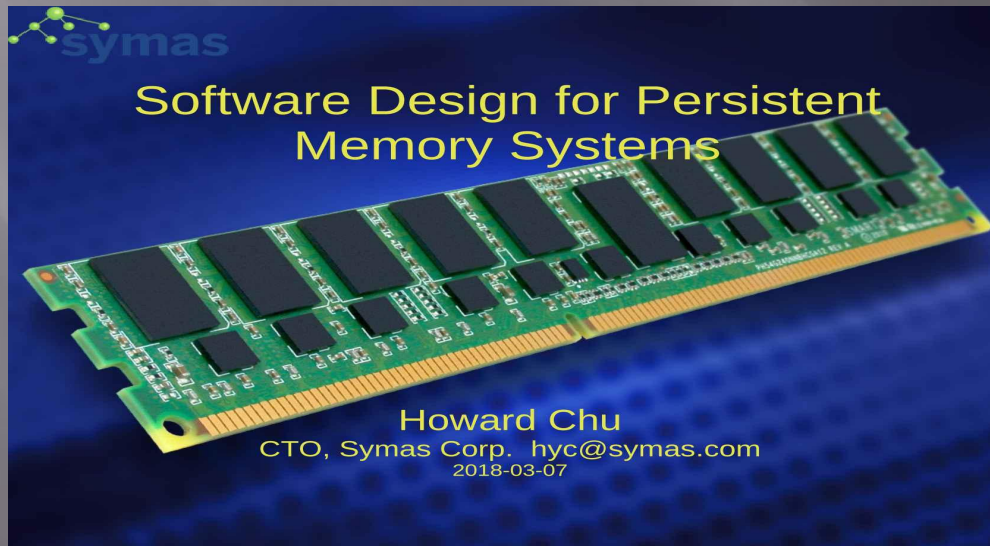Korobka P.L.

Random access memory

- Random access memory (RAM) is the best known form of computer memory. RAM is considered "random access" because you can access any memory cell directly if you know the row and column that intersect at that cell.
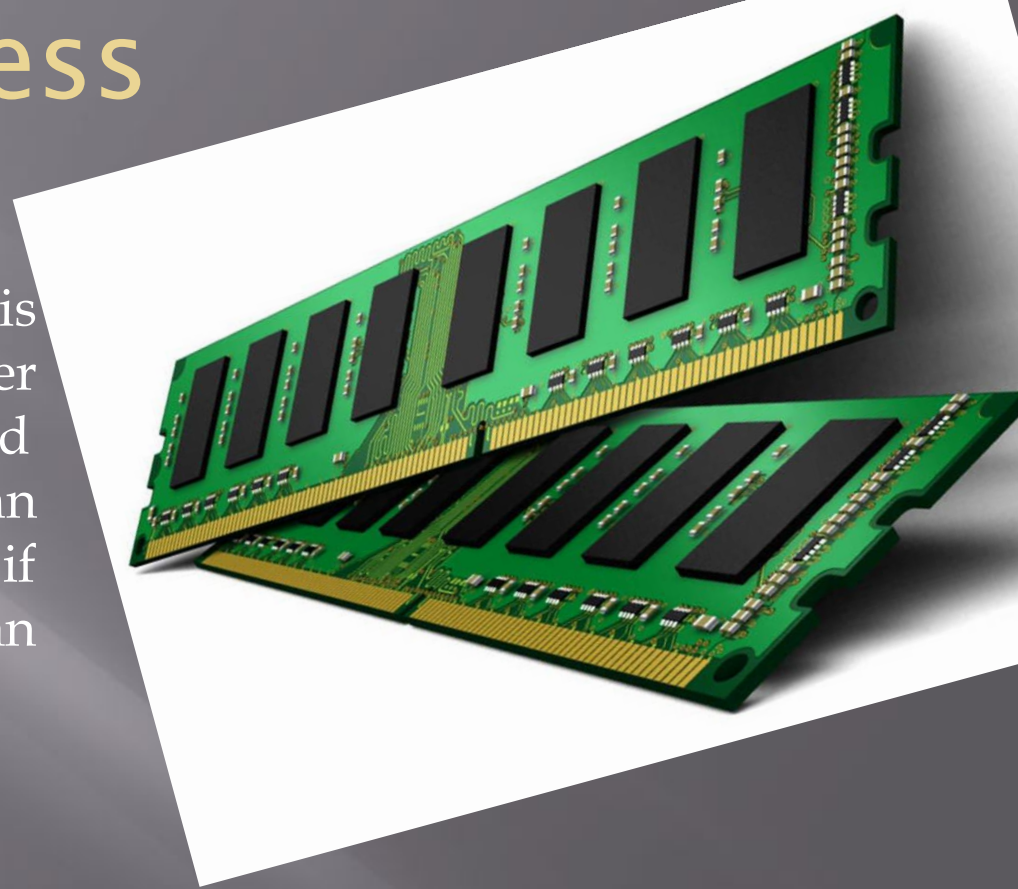
The opposite of RAM is serial access memory (SAM). SAM stores data as a series of memory cells that can only be accessed sequentially (like a cassette tape). If the data is not in the current location, each memory cell is checked until the needed data is found. SAM works very well for memory buffers, where the data is normally stored in the order in which it will be used (a good example is the texture buffer memory on a video card). RAM data, on the other hand, can be accessed in any order.

Memory Modules

▫ Memory chips in desktop computers originally used a pin configuration called dual inline package (DIP). This pin configuration could be soldered into holes on the computer's motherboard or plugged into a socket that was soldered on the motherboard. This method worked fine when computers typically operated on a couple of megabytes or less of RAM, but as the need for memory grew, the number of chips needing space on the motherboard increased.


Software Design for Persistent Memory Systems

Howard Chu
CTO, Symas Corp. hyc@symas.com
2018-03-07

Memory Modules

- The solution was to place the memory chips, along with all of the support components, on a separate printed circuit board (PCB) that could then be plugged into a special connector (memory bank) on the motherboard. Most of these chips use a small outline J-lead (SOJ) pin configuration, but quite a few manufacturers use the thin small outline package (TSOP) configuration as well. The key difference between these newer pin types and the original DIP configuration is that SOJ and TSOP chips are surface-mounted to the PCB. In other words, the pins are soldered directly to the surface of the board, not inserted in holes or sockets.

- Memory chips are normally only available as part of a card called a module. You've probably seen memory listed as 8x32 or 4x16. These numbers represent the number of the chips multiplied by the capacity of each individual chip, which is measured in megabits (Mb), or one million bits. Take the result and divide it by eight to get the number of megabytes on that module. For example, 4x32 means that the module has four 32-megabit chips. Multiply 4 by 32 and you get 128 megabits. Since we know that a byte has 8 bits, we need to divide our result of 128 by 8. Our result is 16 megabytes!

Error Checking

Error Checking Most memory available today is highly reliable. Most systems simply have the memory controller check for errors at start-up and rely on that. Memory chips with built-in error-checking typically use a method known as parity to check for errors. Parity chips have an extra bit for every 8 bits of data. The way parity works is simple. Let's look at even parity first
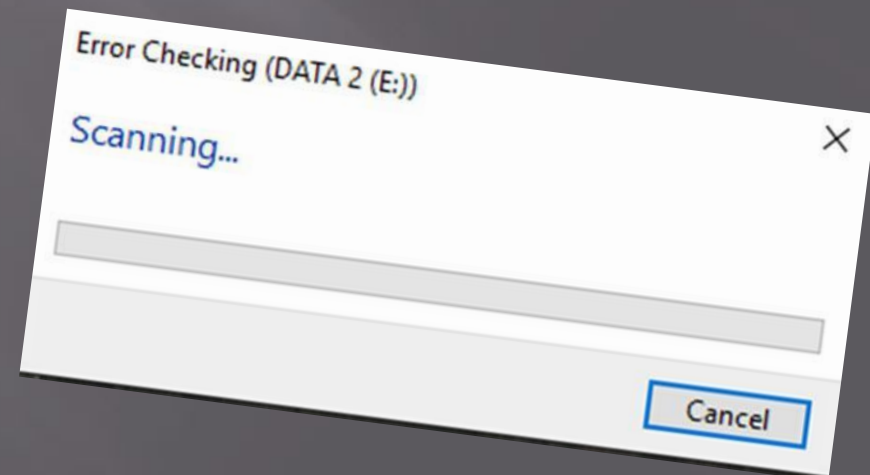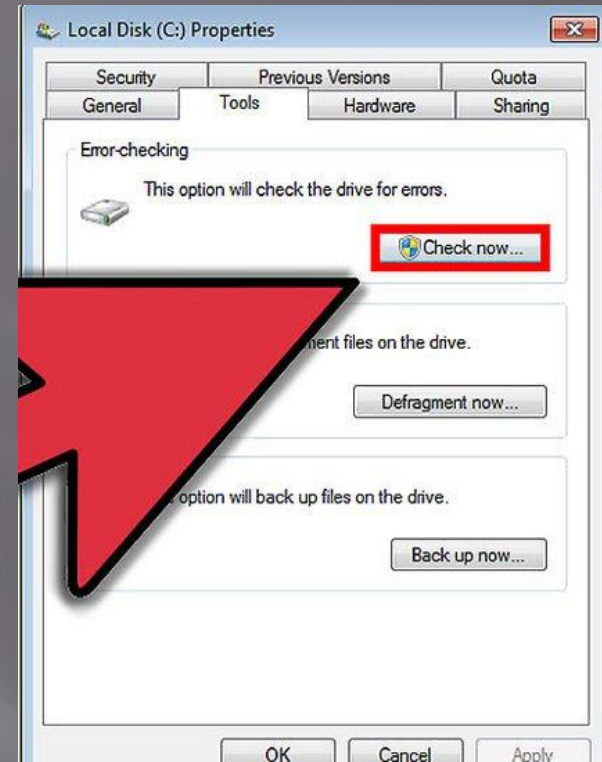
Error Checking

▫ When the 8 bits in a byte receive data, the chip adds up the total number of 1s. If the total number of 1s is odd, the parity bit is set to 1. If the total is even, the parity bit is set to 0. When the data is read back out of the bits, the total is added up again and compared to the parity bit. If the total is odd and the parity bit is 1, then the data is assumed to be valid and is sent to the CPU. But if the total is odd and the parity bit is 0, the chip knows that there is an error somewhere in the 8 bits and dumps the data. Odd parity works the same way, but the parity bit is set to 1 when the total number of 1s in the byte are even.

▫ The majority of computers sold today use nonparity memory chips. These chips do not provide any type of built-in error checking, but instead rely on the memory controller for error detection.

Common RAM Types

- **SRAM** Static random access memory uses multiple transistors, typically four to six, for each memory cell but doesn't have a capacitor in each cell. It is used primarily for cache.

- **DRAM** Dynamic random access memory has memory cells with a paired transistor and capacitor requiring constant refreshing.

- **FPM DRAM** Fast page mode dynamic random access memory was the original form of DRAM. It waits through the entire process of locating a bit of data by column and row and then reading the bit before it starts on the next bit. Maximum transfer rate to L2 cache is approximately 176 MBps.

- **EDO DRAM** Extended data-out dynamic random access memory does not wait for all of the processing of the first bit before continuing to the next one. As soon as the address of the first bit is located, EDO DRAM begins looking for the next bit. It is about five percent faster than FPM. Maximum transfer rate to L2 cache is approximately 264 MBps.
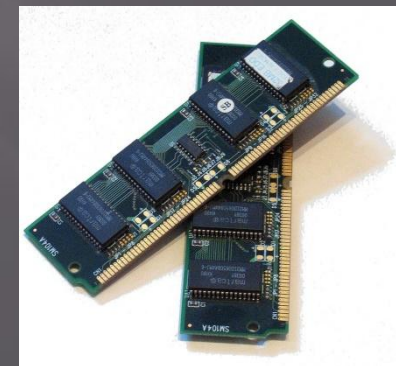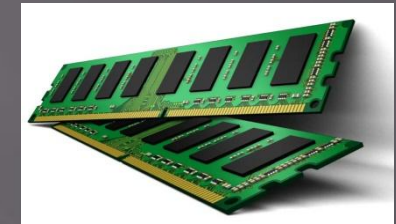
Common RAM Types

- **SDRAM** Synchronous dynamic random access memory takes advantage of the burst mode concept to greatly improve performance. It does this by staying on the row containing the requested bit and moving rapidly through the columns, reading each bit as it goes. The idea is that most of the time the data needed by the CPU will be in sequence. SDRAM is about five percent faster than EDO RAM and is the most common form in desktops today. Maximum transfer rate to L2 cache is approximately 528 MBps.

- **DDR SDRAM** Double data rate synchronous dynamic RAM is just like SDRAM except that is has higher bandwidth, meaning greater speed. Maximum transfer rate to L2 cache is approximately 1,064 MBps (for DDR SDRAM 133 MHZ).

- **RDRAM** Rambus dynamic random access memory is a radical departure from the previous DRAM architecture. Designed by Rambus, RDRAM uses a Rambus in-line memory module (RIMM), which is similar in size and pin configuration to a standard DIMM. What makes RDRAM so different is its use of a special high-speed data bus called the Rambus channel.
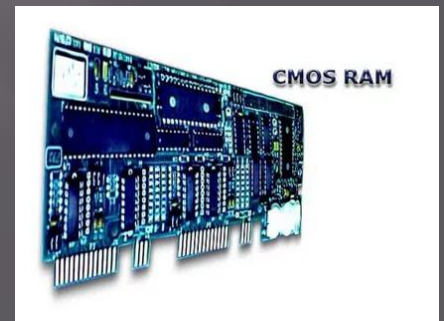
Common RAM Types

- **Credit Card Memory** Credit card memory is a proprietary self-contained DRAM memory module that plugs into a special slot for use in notebook computers.

- **PCMCIA** Memory Card Another self-contained DRAM module for notebooks, cards of this type are not proprietary and should work with any notebook computer whose system bus matches the memory card's configuration.

- **CMOS RAM.** CMOS RAM is a term for the small amount of memory used by your computer and some other devices to remember things like hard disk settings -- see Why does my computer need a battery? for details. This memory uses a small battery to provide it with the power it needs to maintain the memory contents.

Common RAM Types

- **VRAM** VideoRAM, also known as multiport dynamic random access memory (MPDRAM), is a type of RAM used specifically for video adapters or 3-D accelerators. The "multiport" part comes from the fact that VRAM normally has two independent access ports instead of one, allowing the CPU and graphics processor to access the RAM simultaneously. VRAM is located on the graphics card and comes in a variety of formats, many of which are proprietary. The amount of VRAM is a determining factor in the resolution and color depth of the display. VRAM is also used to hold graphics-specific information such as 3-D geometry data and texture maps. True multiport VRAM tends to be expensive, so today, many graphics cards use SGRAM (synchronous graphics RAM) instead. Performance is nearly the same, but SGRAM is cheaper.
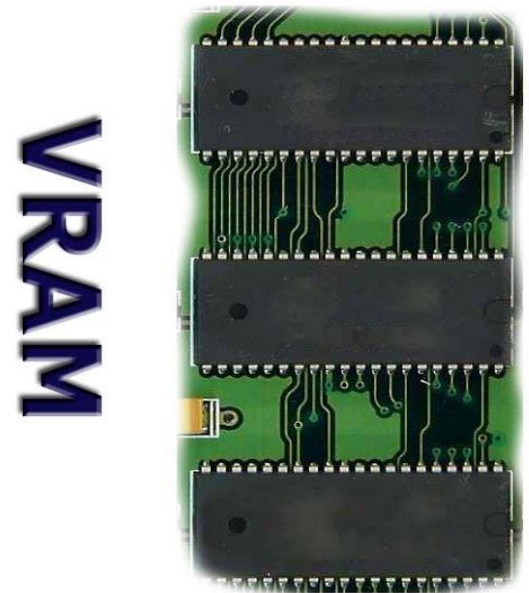
ROM Types

- There are five basic ROM types:
- · ROM
- · PROM
- · EPROM
- · EEPROM
- · Flash memory Each type has unique characteristics, which you'll learn about in this article, but they are all types of memory with two things in common:
- · Data stored in these chips is nonvolatile -- it is not lost when power is removed. · Data stored in these chips is either unchangeable or requires a special operation to change (unlike RAM, which can be changed as easily as it is read).

ROM at Work

3.1 ROM

License key:
XXXXX-XXXXX-XXXXX-XXXXX

Registration and legal: cloanto.com/genuine
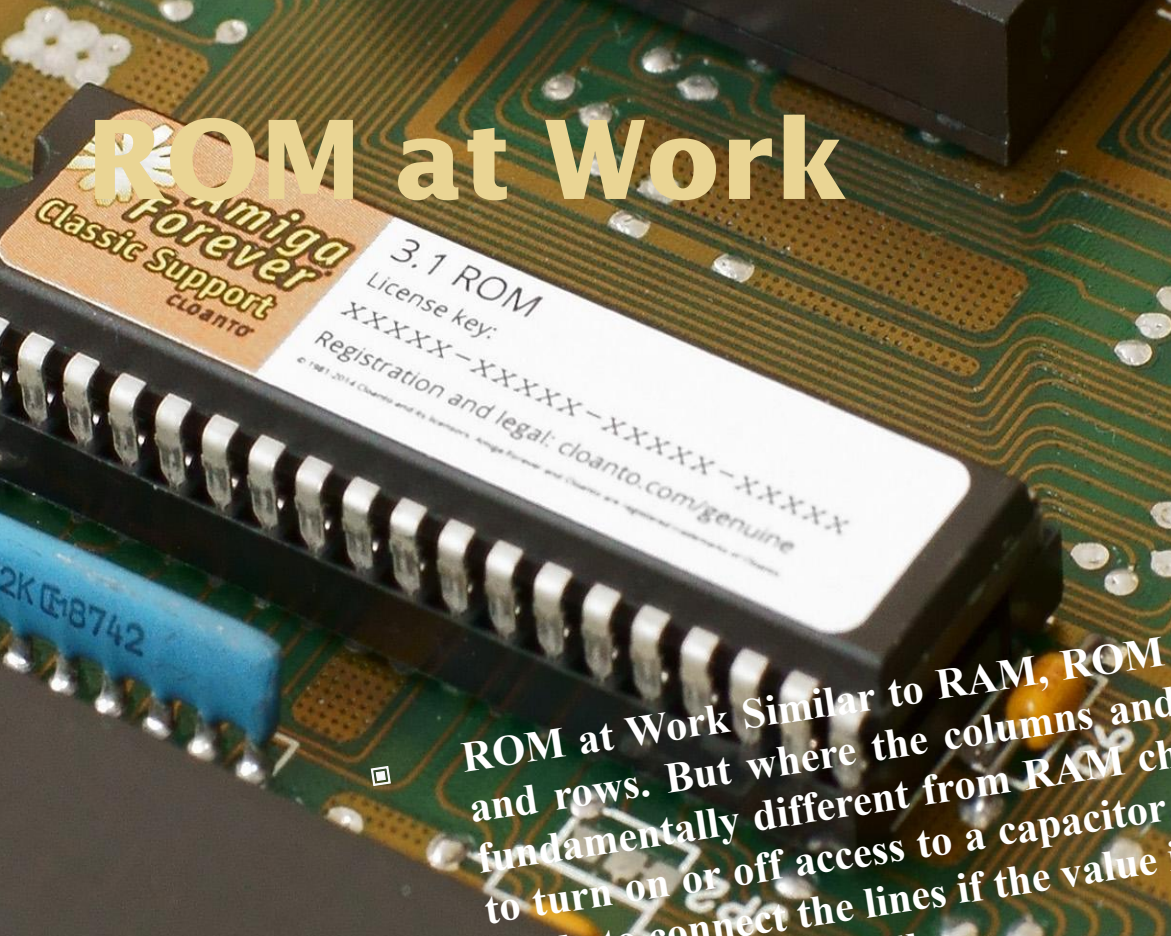
- ROM at Work Similar to RAM, ROM chips contain a grid of columns and rows. But where the columns and rows intersect, ROM chips are fundamentally different from RAM chips. While RAM uses transistors to turn on or off access to a capacitor at each intersection, ROM uses a diode to connect the lines if the value is 1. If the value is 0, then the lines are not connected at all.

- A diode normally allows current to flow in only one direction and has a certain threshold, known as the forward breakover, that determines how much current is required before the diode will pass it on. In silicon-based items such as processors and memory chips, the forward breakover voltage is approximately 0.6 volts. By taking advantage of the unique properties of a diode, a ROM chip can send a charge that is above the forward breakover down the appropriate column with the selected row grounded to connect at a specific cell.

A Simple Example

- Caching is a technology based on the memory subsystem of your computer. The main purpose of a cache is to accelerate your computer while keeping the price of the computer low. Caching allows you to do your computer tasks more rapidly.

- The first customer arrives. He asks for the book Moby Dick. The librarian goes into the storeroom, gets the book, returns to the counter and gives the book to the customer. Later, the client comes back to return the book. The librarian takes the book and returns it to the storeroom. He then returns to his counter waiting for another customer. Let's say the next customer asks for Moby Dick (you saw it coming...). The librarian then has to return to the storeroom to get the book he recently handled and give it to the client. Under this model, the librarian has to make a complete round trip to fetch every book -- even very popular ones that are requested frequently. Is there a way to improve the performance of the librarian?

A Simple Example

- Yes, there's a way we can put a cache on the librarian. Let's give the librarian a backpack into which he will be able to store 10 books (in computer terms, the librarian now has a 10-book cache). In this backpack, he will put the books the clients return to him, up to a maximum of 10. Let's use the prior example, but now with our new-and-improved caching librarian.

Caching

CPU

Cache(fast)

Relatively Slow

RAM
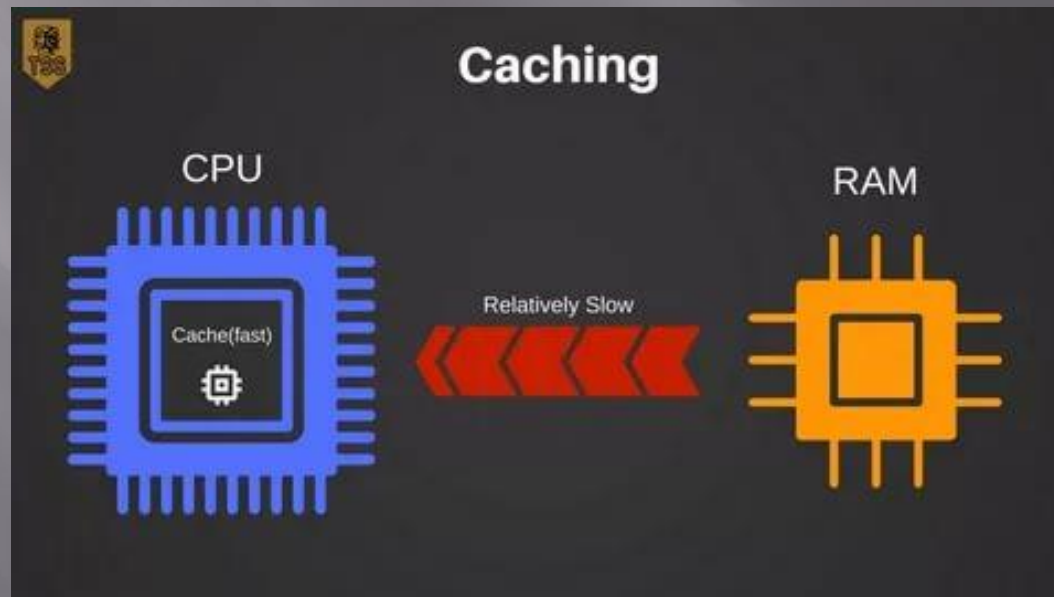
SmartMedia

- The solid-state floppy-disk card (SSFDC), better known as SmartMedia, was originally developed by Toshiba.

- SmartMedia cards are available in capacities ranging from 2 MB to 128 MB. The card itself is quite small, approximately 45 mm long, 37 mm wide and less than 1 mm thick. This is amazing when you consider what is packed into such a tiny package!

CompactFlash

▫ CompactFlash cards were developed by Sandisk in 1994, and they are different from SmartMedia cards in two important ways:

▫ · They are thicker.

▫ · They utilize a controller chip.

▫ CompactFlash consists of a small circuit board with Flash-memory chips and a dedicated controller chip, all encased in a rugged shell that is several times thicker than a SmartMedia card.

▫ As shown below, CompactFlash cards are 43 mm wide and 36 mm long, and come in two thicknesses: Type I cards are 3.3 mm thick, and Type II cards are 5.5 mm thick.
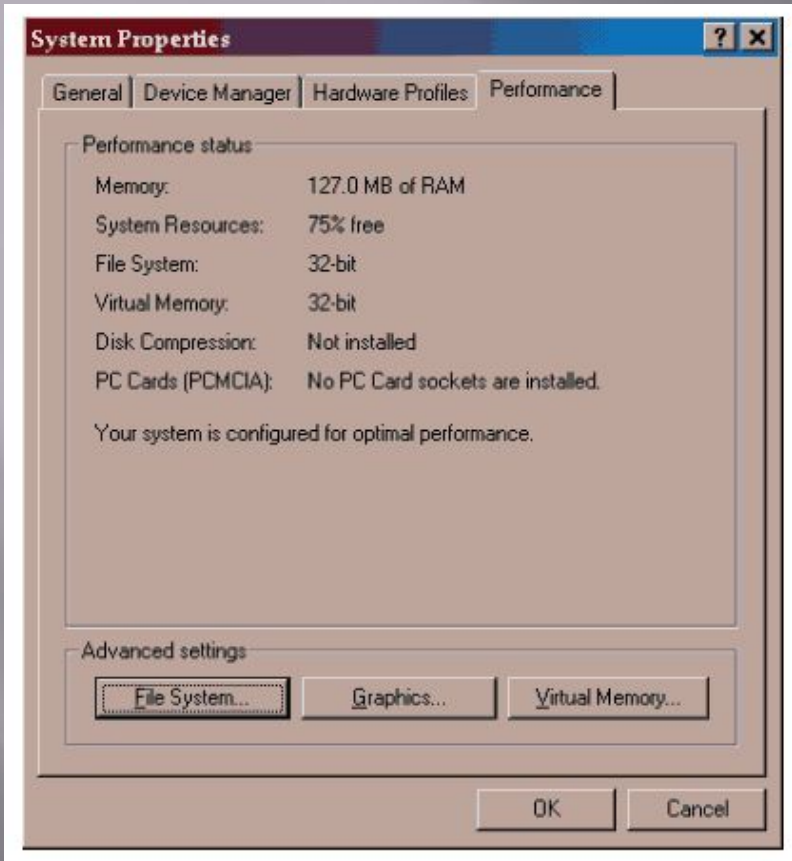
What is Virtual Memory?

- Most computers today have something like 32 or 64 megabytes of RAM available for the CPU to use (see How RAM Works for details on RAM). Unfortunately, that amount of RAM is not enough to run all of the programs that most users expect to run at once. For example, if you load the operating system, an e-mail program, a Web browser and word processor into RAM simultaneously, 32 megabytes is not enough to hold it all. If there were no such thing as virtual memory, then once you filled up the available RAM your computer would have to say, "Sorry, you can not load any more applications. Please close another application to load a new one." With virtual memory, what the computer can do is look at RAM for areas that have not been used recently and copy them onto the hard disk. This frees up space in RAM to load the new application
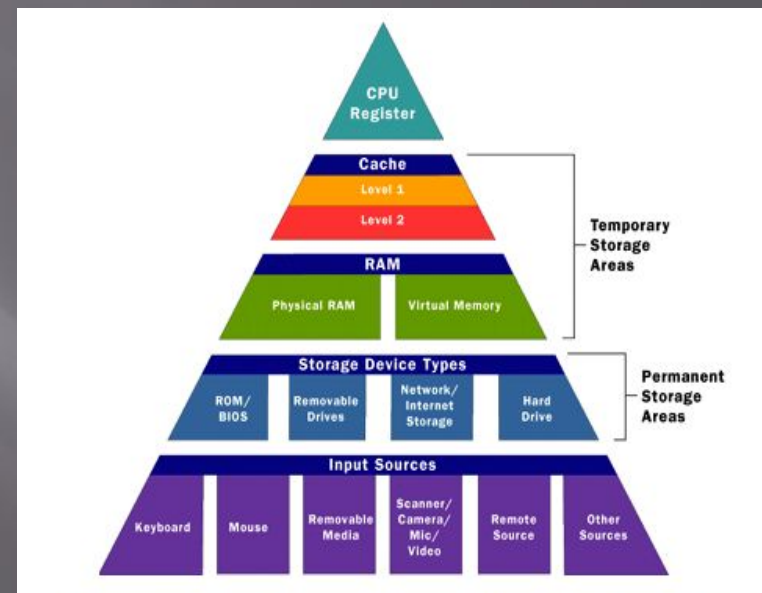
Speed Concerns

- The read/write speed of a hard drive is much slower than RAM, and the technology of a hard drive is not geared toward accessing small pieces of data at a time. If your system has to rely too heavily on virtual memory, you will notice a significant performance drop. The key is to have enough RAM to handle everything you tend to work on simultaneously -- then, the only time you "feel" the slowness of virtual memory is is when there's a slight pause when you're changing tasks. When that's the case, virtual memory is perfect.

- When it is not the case, the operating system has to constantly swap information back and forth between RAM and the hard disk. This is called thrashing, and it can make your computer feel incredibly slow.

- The area of the hard disk that stores the RAM image is called a page file. It holds pages of RAM on the hard disk, and the operating system mo ves data back and forth between the page file and RAM. On a Windows machine, page files have a .SWP extension.

Configuring Virtual Memory

Take Windows 98 as an example of a typical operating system that has virtual memory. Windows 98 has an intelligent virtual memory manager that uses a default setting to help Windows allocate hard drive space for virtual memory as needed. For most circumstances, this should meet your needs, but you may want to manually configure virtual memory, especially if you have more than one physical hard drive or speed-critical applications.

Configuring Virtual Memory
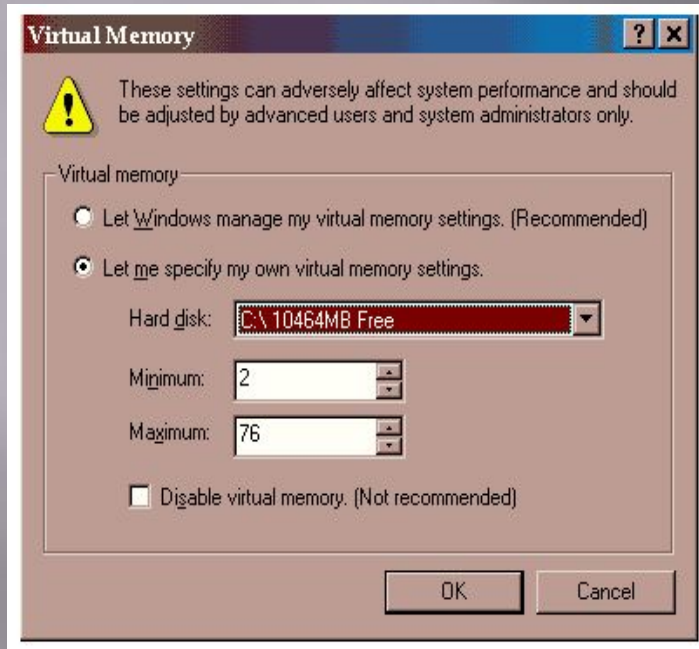
Click on the option that says, "Let me specify my own virtual memory settings." This will make the options below that statement become active. Click on the drop-down list beside "Hard disk:" to select the hard drive that you wish to configure virtual memory for. Remember that a good rule of thumb is to equally split virtual memory between the physical hard disks you have.

Configuring Virtual Memory

- The amount of hard drive space you allocate for virtual memory is important. If you allocate too little, you will get "Out of Memory" errors. If you find that you need to keep increasing the size of the virtual memory, you probably are also finding that your system is sluggish and accesses the hard drive constantly. In that case, you should consider buying more RAM to keep the ratio between RAM and virtual memory about 2:1. Some applications enjoy having lots of virtual memory space but do not access it very much. In that case, large paging files work well.

- Another factor in the performance of virtual memory is the location of the pagefile. If your system has multiple physical hard drives (not multiple drive letters, but actual drives), you can spread the work among them by making smaller pagefiles on each drive. This simple modification will significantly speed up any system that makes heavy use of virtual memory.

Thank you for listening. End