

The Hardware Guide

Do you know from which parts your computer is made. Or want to invest in new hardware by purchasing a new one or by upgrade. Which one to upgrade to get most performance. How to spent in wisely manner. This document will help you.

 [Bios Guides](#)

 [The Bus Speed Guide](#)

 [Chipset Guides](#)

 [CPU Guide](#)


 [The Graphics Card Guide](#)

 [The Hard Disk Guide](#)

 [Keyboard Guide](#)

 [The Motherboard Guide](#)

 [Mouse and Joystick Guide](#)

 [Modems and Fax Guide](#) [Monitor Guide](#) [Printer Guide](#) [RAM Guide](#) [Scanner Guide](#) [Speakers and Sound Card Guide](#) [Troubleshooting Guide](#) [Upgrading Your Hardware Guide](#) [Video Formats Guide](#) [About Me](#)

THE BIOS GUIDE

BIOS settings are a frequent problem asked about in several hardware related newsgroups. Did you ever experienced a system lock up or poor performance and erratic behavior due to improper BIOS settings? Have you ever been left in the dark by a cryptic 5 pages, badly written motherboard manual? The answer is probably yes.

BIOS

Basic Input Output System. All computer hardware has to work with software through an interface. The BIOS gives the computer a little built-in starter kit to run the rest of softwares from floppy disks (FDD) and hard disks (HDD). The BIOS is responsible for booting the computer by providing a basic set of instructions. It performs all the tasks that need to be done at start-up time: POST (Power-On Self Test, booting an operating system from FDD or HDD). Furthermore, it provides an interface to the underlying hardware for the operating system in the form of a library of interrupt handlers. For instance, each time a key is pressed, the CPU (Central Processing Unit) perform an interrupt to read that key. This is similar for other input/output devices (Serial and parallel ports, video cards, sound cards, hard disk controllers, etc...). Some older PC's cannot co-operate with all the modern hardware because their BIOS doesn't support that hardware. The operating system cannot call a BIOS routine to use it; this problem can be solved by replacing your BIOS with an newer one, that does support your new hardware, or by installing a device driver for the hardware.

CMOS

Complementary Metal Oxide Semiconductor. To perform its tasks, the BIOS need to know various parameters (hardware configuration). These are permanently saved in a little piece (64 bytes) of CMOS RAM (short: CMOS). The CMOS power is supplied by a little battery, so its contents will not be lost after the PC is turned off. Therefore, there is a battery and a small RAM memory on board, which never (should...) loses its information. The memory was in earlier times a part of the clock chip, now it's part of such a highly Integrated Circuit (IC). CMOS is the name of a technology which needs very low power so the computer's battery is not too much in use.

Actually, there is not a battery on new boards, but an accumulator (Ni_Cad in most cases). It is recharged every time the computer is turned on. If your CMOS is powered by external batteries, be sure that they are in good operating condition. Also, be sure that they do not leak. That may damage the motherboard. Otherwise, your CMOS may suddenly "forget" its configuration and you may be looking for a problem elsewhere. In the monolithic PC and PC/XT, this information is supplied by setting the DIP (Dual-In-line Package) switches at the motherboard or peripheral cards. Some new motherboards have a technology named the Dallas Nov-Ram. It eliminates having an on-board battery: There is a 10 year lithium cell epoxyed into the chip.

Chipset

A PC consists of different functional parts installed on its motherboard: ISA (Industry Standard Architecture), EISA (Enhanced Industry Standard Architecture) VESA (Video Enhanced Standards Association) and PCI (Peripheral Component Interface) slots, memory, cache memory, keyboard plug etc... Not all of these are present on every motherboard. The chipset enables a set of instructions so the CPU can work (communicate) with other parts of the motherboard. Nowadays most of the discrete chips; PIC (Programmable Interrupt Controller), DMA (Direct Memory Access), MMU (Memory Management Unit), cache, etc... are packed together on one, two or three chips; the chipset. Since chipsets of a different brand are not the same, for every chipset there is a BIOS version. Now we have fewer and fewer chipsets which do the job. Some chipsets have more features, some less. OPTi is such a commonly used chipset. In some well integrated motherboards, the only components present are the CPU, the two BIOS chips (BIOS and Keyboard BIOS), one chipset IC, cache memory (DRAMs, Dynamic Random Access Memory), memory (SIMMs, Single Inline Memory Module, most of the time) and a clock chip.

Setup

Setup is the set of procedures enabling the configure a computer according to its hardware characteristics. It allows you to change the parameters with which the BIOS configures your chipset. The original IBM PC was configured by means of DIP switches buried on the motherboard. Setting PC and XT DIP switches

properly was something of an arcane art. DIP switches/jumpers are still used for memory configuration and clock speed selection. When the PC-AT was introduced, it included a battery powered CMOS memory which contained configuration information. CMOS was originally set by a program on the Diagnostic Disk, however later clones incorporated routines in the BIOS which allowed the CMOS to be (re)configured if certain magic keystrokes were used.

Unfortunately as the chipsets controlling modern CPUs have become more complex, the variety of parameters specifiable in SETUP has grown. Moreover, there has been little standardization of terminology between the half dozen BIOS vendors, three dozen chipset makers and large number of motherboard vendors. Complaints about poor motherboard documentation of SETUP parameters are very common.

To exacerbate matters, some parameters are defined by BIOS vendors, others by chipset designers, others by motherboard designers, and others by various combinations of the above. Parameters intended for use in Design and Development, are intermixed with parameters intended to be adjusted by technicians -- who are frequently just as baffled by this stuff as everyone else is. No one person or organization seems to understand all the parameters available for any given SETUP.

Hardware Performance

Although computers may have basic similarities (they all look the same on a shelf), performance will differ markedly between them, just the same as it does with cars. The PC contains several processes running at the same time, often at different speeds, so a fair amount of coordination is required to ensure that they don't work against each other.

Most performance problems arise from bottlenecks between components that are not necessarily the best for a particular job, but a result of compromise between price and performance. Usually, price wins out and you have to work around the problems this creates.

The trick to getting the most out of any machine is to make sure that each component is giving of its best, then eliminate potential bottlenecks between them. You can get a bottleneck simply by having an old piece of equipment that

is not designed to work at modern high speed - **a computer is only as fast as its slowest component**, but bottlenecks can also be due to badly written software.

System Timing

The clock is responsible for the speed at which numbers are crunched and instructions executed. It results in an electrical signal that switches constantly between high and low voltage several millions times a second.

The **System Clock**, or CLKIN, is the frequency used by the processor; on 286s and 386s, this will be half the speed of the main crystal on the motherboard (the CPU divides it by two), which is often called CLK2IN. 486 processors run at the same speed, because they use both edges of the timing signal. A clock generator chip (82284 or similar) is used to synchronize timing signals around the computer, and the data bus would be run at slower speed synchronously with the CPU, e.g. CLKIN/4 for an ISA bus with a 33 MHz CPU.

ATCLK is a separate clock for the bus, when it's run asynchronously, or not derived from CLK2IN. There is also a 14.138 MHz crystal which was used for all system timing on XTs. Now it's only used for the colour frequency of the video controller (6845).

Memory Access

The **cycle time** is the time it takes to read from and write to a memory cell, and it consists of two stages; precharge and access. **Precharge** is where the capacitor in the memory cell is able to recover from a previous access and stabilize. **Access** is where a data bit is actually moved between memory and the bus or the CPU. Total **access time** includes the finding of data, data flow and recharge, and parts of the access time can be eliminated or overlapped to improve performance. The combination of precharge and access equals cycle time, which is what you should use to calculate wait states from.

There are ways of making refreshes happen so that the CPU doesn't notice (i.e. Concurrent and Hidden), which is helped by the 486 being able to use its on-board cache and not needing to use memory so often anyway. In addition, you can affect the **Row Access Strobe (RAS)**, or have **Column Access Strobe (CAS)** before RAS (see Advanced Chipset Setup).

The fastest DRAM commonly available is rated at 60ns. As these chips need alternate refresh cycles, under normal circumstances data will actually be obtained every 120ns, giving you an effective speed of around 8 MHz for the whole computer, regardless of the CPU speed, assuming no action is taken to compensate. Memory chips therefore need to be operating at something like 20ns to keep up, assuming that the CPU needs only one clock cycle for each one from the memory bus; one internal cycle for each external one. Intel processors mostly use two for one, so the 33 MHz CPU is actually ready to use memory every 60ns, but you need to allow a little more for overheads, such as data assembly and the like. One way of matching the capacities of components with different speeds includes the use of wait states.

Wait States

A wait state indicates how many ticks of the system clock the CPU has to wait for memory to catch up—it will generally be 0 or 1, but can be up to 3 if you're using slower memory chips. Ways of avoiding wait states include:

◆ **Page-mode memory.** This will cut-down address cycles to retrieve information from one general area, based on the fact that the second access to a memory location on the same page takes around half the time as the first; addresses are normally in two halves, with high bits (for row) and low bits (for column) being multiplexed onto one set of address pins. The page address of data is noted, and if the next data is in the same area, a second address cycle is eliminated as a whole row of memory cells can be read in one go; that is, once a row access has been made, you can get to subsequent column addresses in that row in the time available (you should therefore increase row access time for best performance). Otherwise data is retrieved normally, which will take twice as long. **Fast Page Mode** is a quicker version of the same thing; the DRAMs concerned have a faster CAS access speed. Memory capable of running in page mode is different from normal bit-by-bit type, and the two don't mix. It's unlikely that low capacity SIMMs are so capable.

◆ **Interleaved memory,** which divides memory into two or four portions that process data alternately; that is, the CPU sends information to one section while another goes through a refresh cycle; a typical installation will have odd addresses on one side and even on the other (you can have word or block interleave). If memory accesses are sequential, the precharge of one will overlap

the access time of the other. To put interleaved memory to best use, fill every socket you've got (that is, eight 1 Mb SIMMs are better than two 4 Mb ones). The SIMM types must be the same. As an example, a machine in non-interleaved mode (say a 386SX/20) may need 60ns or faster DRAM for 0ws access, where 80ns chip could do if interleaving were enabled.

◆ A processor **RAM cache**, which is a bridge between the CPU and slower main memory; it consists of anywhere between 32-512K of (fast) Static RAM chips and is designed to retain the most frequently accessed code and data from main memory. It can make 1 wait state RAM look like that with 0 wait states, without physical adjustments, assuming that the data the CPU wants is in the cache when required (known as a **cache hit**). To minimize the penalty of a cache miss, cache and memory access are often in parallel, with one being terminated when not required. On a 486, how much cache you need really depends on the amount of memory; Dell say that jumping from 128K to 256K only increases the hit rate by around 5% and Viglen say you only need more than 256K if you have more than 32 Mb RAM. A cache should be fast and capable of holding the contents of several different parts of main memory. Software plays a part as well, since cache operation is based on the assumption that programs access memory where they have done so already, or are likely to next, maybe through looping (where code is reused) or code is organized to be next to other relevant parts. A basic cache design will look up an address for the CPU and return the data inside one clock cycle, or 20ns at 50 MHz. *Asynchronous* SRAM will be used for this. As the round trip from the CPU to cache and back again takes up a certain amount of time, only the remainder is available to retrieve data, which gets smaller as the motherboard speed is increased. *Synchronous* SRAM uses a buffer to keep the whole routine inside one clock cycle, even though it may use two (or more) clock cycles the first time round. The address from the CPU is stored, and while the next is coming in to the buffer, the data for the first is retrieved, and the cycle continues. Pipeline SRAM uses more clock cycles, typically three, the first time round, and Burst SRAM will deliver 4 words (blocks of data) over for consecutive cycles if the request from the CPU is for the first; there will be no waiting for the CPU to request each one individually. Note the level 2 cache can be unreliable, so be prepared to disable it in the interests of reliability. For maximum efficiency, or minimum access time, a cache may be subdivided into smaller blocks that can be separately loaded, so the chances of a different part of memory being requested and the time needed to replace a wrong section are minimized. There are three mapping schemes that assist with this:

◆ **Fully Associative**, where the whole address is kept with each block of data in the cache (in tag RAM), needed because it is assumed there is no relationships between the blocks. This can be inefficient, as an address comparison needs to be made with every entry each time the CPU presents the address for its next instruction.

◆ **Direct Mapped**, where every block can only be in one place in the cache, so only one address comparison is needed to see if the data required is there. Although simple, the cache controller must go to main memory more frequently if program code needs to jump between locations with the same index, which defeats the object somewhat, as alternate references to the same cache cell mean cache misses for other processes. The "index" comes from the lower order addresses presented by the CPU.

◆ **Set Associative**, a compromise between the above two. Here, an index can select several entries, so in a *2 Way Set Associative* cache, 2 entries can have the same index, so two comparisons are needed to see if the data required is in the cache. Also, the tag field is correspondingly wider and needs larger SRAMs to store address information. As there are two locations for each index, the cache controller has to decide which one to update or overwrite, as the case may be. The most common methods used to make these decisions are *Random Replacement*, *First In First Out (FIFO)* and *Least Recently Used (LRU)*. The latter is the most efficient. If the cache is large enough (e.g. 64K), performance from this over direct-mapping may not be much. A *Write Thru Cache* means that every write access is immediately passed on to memory; although it means that cache contents are always identical to main memory, it is slow, as the CPU then has to wait for DRAMs. Buffers can be used to provide a variation on this, where data is written into a temporary buffer so the CPU is released quickly before main memory is updated. A *Write Back Cache*, on the other hand, exists where changed data is temporarily stored in the cache and written to memory when the system is quiet, or when absolutely necessary. This will give better performance when main memory is slower than the cache, or when several writes are made in a very short space of time, but is more expensive. A "dirty bit" is used as a mental note that the cache and main memory contents are different, and that the cache contains the most up to date data. This bit will be checked if the cache needs to be written to, and main memory updated first if this bit is set. Some motherboards don't have the required SRAM for the dirty bit, but it's still faster than Write Thru.

Shadow RAM

ROMs are used by components that need their own instructions to work properly, such as video card or cacheing disk controller. ROMs are 8-bit devices, so only one byte is accessed at a time; also, they typically run between 150-400ns, so using them will be slow relative to 32-bit memory at 60-80ns, which is capable of making four accesses at once.

Shadow RAM is the process of copying the contents of a ROM directly into extended memory which is given the same address as the ROM, from where it will run much faster. The original ROM is then disabled, and the new location write protected. If your applications execute ROM routines often enough, enabling Shadow RAM will make a difference in performance of around 8%, assuming a program spends about 10% of its time using instructions from ROM, but theoretically as high as 300%. The drawback is that the RAM set aside for shadowing cannot be used for anything else, and you will lose a corresponding amount of extended memory, The remainder of Upper Memory, however, can usually be remapped to the end of extended memory and used there.

With some VGA cards, if video shadow is disabled, you might get DMA errors, because of timing when code is fetched from the VGA BIOS, when the CPU cannot accept DMA requests. Some programs don't make use of the video ROM, preferring to directly address the card's registers, so you may want to use extended memory for something else. If your machine hangs during the startup sequence for no apparent reason, check that you haven't shadowed an area of upper memory containing a ROM that doesn't like it-particularly one on a hard disk controller, or that you haven't got two in the same 128K segment.

Bus Types

The expansion bus (where expansion cards go) is an extension of the Central Processor, so when adding cards to it, you are extending the capabilities of the CPU itself. The relevance of this regard to the BIOS is that older cards are less able to cope with modern buses running at higher speeds than the original design of 8 or so MHz. Also, when the bus is accessed, the **whole computer** slows down to the bus speed, so it's often worth altering the speed of the bus or the wait states between it and the CPU to speed things up. The PC actually has four

buses; the **processor bus** connects the CPU to its support chips, the **memory bus** connects it to its memory, the **address bus** is part of both of them, and the **I/O (or expansion) bus** is what concerns us here.

ISA

Industry Standard Architecture. The 8-bit version came on the original PC and the AT, but the latter uses an extension to make it 16-bit. It has a *maximum* data transfer rate of about 8 megabits per second on an AT, which is actually well above the capability of disk drives, or most network and video cards. The *average* data throughput is around a quarter of that. Its design makes it difficult to mix 8- and 16-bit RAM or ROM within the same 128K block of upper memory; an 8-bit VGA card could force all other cards in the same (C000-DFFF) range to use 8 bits as well, which was a common source of inexplicable crashes where 16-bit network cards were involved.

EISA

Extended Industry Standard Architecture. An evolution of ISA and (theoretically) backward compatible with it, including the speed (8 MHz), so the increased data throughput is mainly due to the bus doubling in size-but you must use EISA expansion cards. It has its own DMA arrangements, which can use the complete address space. One advantage of EISA (and Micro Channel) is the ease of setting up expansion cards-plus them in and run the configuration software which will automatically detect their settings.

MCA

Micro Channel Architecture. A proprietary standard established by IBM to take over from ISA, and therefore incompatible with anything else. It comes in two versions, 16- and 32-bit and, in practical terms, is capable of transferring around 20 mbps.

Local Bus

The local bus is one more directly suited to the CPU; it's next door (hence local),

has the same bandwidth and runs at the same speed, so the bottleneck is less (ISA was local in the early days). Data is therefore moved along the bus at processor speeds. There are two varieties:

◆ **VL-BUS**, a 32-bit bus which allows bus mastering, and uses two cycles to transfers a 32-bit word, peaking at 66 Mb/sec. It also supports burst mode, where a single address cycle precedes four data cycles, meaning that 4 32-bit words can move in only 5 cycles, as opposed to 8, giving 105 Mb/sec at 33 MHz. The speed is mainly obtained by allowing VL-Bus adapter cards first choice at intercepting CPU cycles. It's not designed to cope with more than a certain number of cards at particular speeds; e.g. 3 at 33, 2 at 40 and only 1 at 50 MHz, and even that often needs a wait state inserted. VL-Bus 2 is 64-bit, yielding 320 Mb/sec at 50 MHz. There are two types of slot; Master and Slave. Master boards (e.g. SCSI controllers) have their own CPUs which can do their own things; slaves (i.e. video cards) don't. A slave board will work on a master slot, but not vice versa.

◆ **PCI**, which is a mezzanine bus, divorced from the CPU, giving it some independence and the ability to cope with more devices, so it's more suited to cross-platform work. It is time multiplexed, meaning that address and data lines share connections. It has its own burst mode that allows 1 address cycle to be followed by as many data cycles as system overheads allow. At nearly 1 word per cycle, the potential is 264 Mb/sec. It can operate up to 33 MHz, or 66 MHz with PCI 2.1 and can transfer data at 32 bits per clock cycle so you can get up to 132 Mbytes/sec (264 with 2.1). Each PCI card can perform up to 8 functions, and you can have more than one busmastering card on the bus. It should be noted, though, that many functions are not available with PCI, such as sound. Not yet, anyway. It is part of the *plug and play* standard, assuming your operating system and BIOS agree, so it is auto configuring (although some cards use jumpers instead of storing information in a chip); it will also share interrupts under the same circumstances. The PCI chipset handles transactions between cards and the rest of the system, and allows other buses to be bridged to it (typically and ISA bus to allow older cards to be used). Not all of them are equal, though; certain features, such as *byte merging*, may be absent. The connector may vary according to the voltage the card uses (3.3 or 5v; some cards can cope with both).

Basic Optimization Tricks

This section is intended for users who have a limited knowledge of BIOS setup. It provides four fundamental procedures that may help improve the performance of your system.

◆ Make sure that all **standard settings** correspond to the installed components of your system. For instance, you should verify the date, the time, available memory, hard disks and floppy disks. For more information, go to the

Standard CMOS Setup

section.

◆ Make sure that your **cache memory** (internal and external) is enabled. Of course you must have internal (L1) and external (L2) cache memory present which is always the case for recent systems (less than five years old). For more information, go to the Advanced CMOS Setup section. Recently, some motherboards were found having **fake cache memory**. Some unscrupulous manufacturers are using solid plastics chips containing no memory to lure vendors and customers and then gain extra profits in an highly competitive semiconductors market. Beware!

- Make sure that your **Wait States** values are at the minimum possible. You must however be careful because if values are too low, your system may freeze (hang up).
- Make sure to **shadow** your Video and System ROM. On older systems, this may improve performance significantly.
- Make sure to use a **coherent power management strategy**. Choosing the right timing may increase the life expectancy of your hard disk. See the Power Management section.

Hard disk speed is the major bottleneck for a system performance, notably for those with 16 MB of memory and more. You may have the fastest CPU, lots of memory and a comfortable cache, but if you have a crummy hard disk, you may not see improvement in performances.

POST and Entering Setup

When the system is powered on, the BIOS will perform diagnostics and initialize system components, including the video system. (This is self-evident when the screen first flicks before the Video Card header is displayed). This is commonly referred as POST (**Power-On Self Test**). Afterwards, the computer will proceed its final boot-up stage by calling the operating system. Just before that, the user may interrupt to have access to SETUP.

To allow the user to alter the CMOS settings, the BIOS provides a little program, **SETUP**. Usually, setup can be entered by pressing a special key combination (**DEL**, **ESC**, **CTRL-ESC**, or **CTRL-ALT-ESC**) at boot time (Some BIOSes allow you to enter setup at any time by pressing **CTRL-ALT-ESC**). The AMI BIOS is mostly entered by pressing the **DEL** key after resetting (**CTRL-ALT-DEL**) or powering up the computer. You can bypass the extended CMOS settings by holding the **<INS>** key down during boot-up. This is really helpful, especially if you bend the CMOS settings right out of shape and the computer won't boot properly anymore. This is also a handy tip for people who play with the older AMI BIOSes with the XCMOS setup. It allows changes directly to the chip registers with very little technical explanation.

A Typical BIOS POST Sequence

Most BIOS POST sequences occur along four stages:

1. Display some basic information about the video card like its brand, video BIOS version and video memory available.
2. Display the BIOS version and copyright notice in upper middle screen. You will see a large sequence of numbers at the bottom of the screen. This sequence is the BIOS identification line.
3. Display memory count. You will also hear tick sounds if you have enabled it (see Memory Test Tick Sound section).
4. Once the POST have succeeded and the BIOS is ready to call the operating system (DOS, OS/2, NT, WIN95, etc.) you will see a basic table of the system's configurations:

◆ **Main Processor:** The type of CPU identified by the BIOS. Usually Cx386DX, Cx486DX, etc..

◆ **Numeric Processor:** Present if you have a FPU or None on the contrary. If you have a FPU and the BIOS does not recognize it, see section Numeric Processor Test in Advanced CMOS Setup.

◆ **Floppy Drive A:** The drive A type. See section Floppy drive A in Standard CMOS Setup to alter this setting.

◆ **Floppy Drive B:** Idem.

◆ **Display Type:** See section Primary display in Standard CMOS Setup.

◆ **AMI or Award BIOS Date:** The revision date of your BIOS. Useful to mention when you have compatibility problems with adaptor cards (notably fancy ones).

◆ **Base Memory Size:** The number of KB of base memory. Usually 640.

◆ **Ext. Memory Size:** The number of KB of extended memory.

In the majority of cases, the summation of base memory and extended memory **does not equal** the total system memory. For instance in a 4096 KB (4MB) system, you will have 640KB of base memory and 3072KB of extended memory, a total of 3712KB. The missing 384KB is reserved by the BIOS, mainly as shadow memory (see Advanced CMOS Setup).

◆ **Hard Disk C: Type:** The master HDD number. See Hard disk C: type section in Standard CMOS Setup.

◆ **Hard Disk D: Type:** The slave HDD number. See Hard disk D: type section in Standard CMOS Setup.

◆ **Serial Port(s):** The hex numbers of your COM ports. 3F8 and 2F8 for COM1 and COM2.

◆ **Parallel Port(s):** The hex number of your LTP ports. 378 for LPT1.

◆ **Other information:** Right under the table, BIOS usually displays the size of

cache memory. Common sizes are 64KB, 128KB or 256KB. See External Cache Memory section in Advanced CMOS Setup.

AMI BIOS POST Errors

During the **POST routines**, which are performed each time the system is powered on, errors may occur. **Non-fatal errors** are those which, in most cases, allow the system to continue the boot up process. The error messages normally appear on the screen. **Fatal errors** are those which will not allow the system to continue the boot-up procedure. If a fatal error occurs, you should consult with your system manufacturer or dealer for possible repairs. These errors are usually communicated through **a series of audible beeps**. The numbers on the fatal error list correspond to the number of beeps for the corresponding error. All errors listed, with the exception of #8, are fatal errors. All errors found by the BIOS will be forwarded to the I/O port 80h.

◆**1 beep**: DRAM refresh failure. The memory refresh circuitry on the motherboard is faulty.

◆**2 beeps**: Parity Circuit failure. A parity error was detected in the base memory (first 64k Block) of the system.

◆**3 beeps**: Base 64K RAM failure. A memory failure occurred within the first 64k of memory.

◆**4 beeps**: System Timer failure. Timer #1 on the system board has failed to function properly.

◆**5 beeps**: Processor failure. The CPU on the system board has generated an error.

◆**6 beeps**: Keyboard Controller 8042-Gate A20 error. The keyboard controller (8042) contains the gate A20 switch which allows the computer to operate in virtual mode. This error message means that the BIOS is not able to switch the CPU into protected mode.

◆**7 beeps**: Virtual Mode (processor) Exception error. The CPU on the motherboard has generated an Interrupt Failure exception interrupt.

◆**8 beeps:** Display Memory R/W Test failure. The system video adapter is either missing or Read/Write Error its memory is faulty. This is not a fatal error.

◆**9 beeps:** ROM-BIOS Checksum failure. The ROM checksum value does not match the value encoded in the BIOS. This is good indication that the BIOS ROMs went bad.

◆**10 beeps:** CMOS Shutdown Register. The shutdown register for the CMOS memory Read/Write Error has failed.

◆**11 beeps:** Cache Error / External Cache Bad. The external cache is faulty.

Other AMI BIOS POST Codes

◆**2 short beeps:** POST failed. This is caused by a failure of one of the hardware testing procedures.

◆**1 long & 2 short beeps:** Video failure. This is caused by one of two possible hardware faults. 1) Video BIOS ROM failure, checksum error encountered. 2) The video adapter installed has a horizontal retrace failure.

◆**1 long & 3 short beeps:** Video failure. This is caused by one of three possible hardware problems. 1) The video DAC has failed. 2) the monitor detection process has failed. 3) The video RAM has failed.

◆**1 long beep: POST successful.** This indicates that all hardware tests were completed without encountering errors.

If you have access to a POST Card reader, (Jameco, etc.) you can watch the system perform each test by the value that's displayed. If/when the system hangs (if there's a problem) the last value displayed will give you a good idea where and what went wrong, or what's bad on the system board. Of course, having a description of those codes would be helpful, and different BIOSes have different meanings for the codes. (could someone point out FTP sites where we could have access to a complete list of error codes for different versions of AMI and Award BIOSes?).

BIOS Error Messages

This is a short list of most frequent on-screen BIOS error messages. Your system may show them in a different manner. When you see any of these, you are in trouble - Doh! (Does someone has any additions or corrections?)

❖ **"8042 Gate - A20 Error"**: Gate A20 on the keyboard controller (8042) is not working.

❖ **"Address Line Short!"**: Error in the address decoding circuitry.

❖ **"Cache Memory Bad, Do Not Enable Cache!"**: Cache memory is defective.

❖ **"CH-2 Timer Error"**: There is an error in timer 2. Several systems have two timers.

❖ **"CMOS Battery State Low"** : The battery power is getting low. It would be a good idea to replace the battery.

❖ **"CMOS Checksum Failure"** : After CMOS RAM values are saved, a checksum value is generated for error checking. The previous value is different from the current value.

❖ **"CMOS System Options Not Set"**: The values stored in CMOS RAM are either corrupt or nonexistent.

❖ **"CMOS Display Type Mismatch"**: The video type in CMOS RAM is not the one detected by the BIOS.

❖ **"CMOS Memory Size Mismatch"**: The physical amount of memory on the motherboard is different than the amount in CMOS RAM.

❖ **"CMOS Time and Date Not Set"**: Self evident.

❖ **"Diskette Boot Failure"**: The boot disk in floppy drive A: is corrupted (virus?). Is an operating system present?

❖ **"Display Switch Not Proper"**: A video switch on the motherboard must be set to either color or monochrome.

❖ **"DMA Error"**: Error in the DMA (Direct Memory Access) controller.

- ◆ **"DMA #1 Error"**: Error in the first DMA channel.
- ◆ **"DMA #2 Error"**: Error in the second DMA channel.
- ◆ **"FDD Controller Failure"**: The BIOS cannot communicate with the floppy disk drive controller.
- ◆ **"HDD Controller Failure"**: The BIOS cannot communicate with the hard disk drive controller.
- ◆ **"INTR #1 Error"**: Interrupt channel 1 failed POST.
- ◆ **"INTR #2 Error"**: Interrupt channel 2 failed POST.
- ◆ **"Keyboard Error"**: There is a timing problem with the keyboard.
- ◆ **"KB/Interface Error"**: There is an error in the keyboard connector.
- ◆ **"Parity Error ????"**: Parity error in system memory at an unknown address.
- ◆ **"Memory Parity Error at xxxxx"**: Memory failed at the xxxxx address.
- ◆ **"I/O Card Parity Error at xxxxx"**: An expansion card failed at the xxxxx address.
- ◆ **"DMA Bus Time-out"**: A device has used the bus signal for more than allocated time (around 8 microseconds).

If you encounter any POST error, there is a good chance that it is an **HARDWARE** related problem. You should at least verify if adaptor cards or other removable components (simms, drams etc...) are properly inserted before calling for help. One common attribute in human nature is to rely on others before investigating the problem yourself. Please don't be a politician (Aide-toi et le ciel t'aidera).

Standard CMOS Setup

You should have your current setup options **written down ON PAPER** somewhere, preferably taped to the inside or the outside of the case. CMOS memory has a tendency to get erased as the battery gets old, or become inaccessible if you forget the password. Especially remember the **hard disk settings**; they are the most important.

If you have warm-booted the computer (via CTRL-ALT-DEL) to go into the CMOS setup, the BIOS routine to handle the "Print Screen" key will probably be installed. You can display each screen of the CMOS setup and press SHIFT-PRINT SCREEN to get a printed copy directly. There are several good CMOS saver programs out on the market, including the PC-Tools and Norton recovery programs. They allow a user to save a copy of the CMOS registers to a file in case the battery dies, or if they messed around with the settings, etc.

◆ **Date (mn/date/year) and Time:** To change the date and time of the system clock. Do not expect your computer to keep track of time as accurately as an atomic clock, or even a wrist watch! Depending on the quality of the motherboard expect to lose (or gain) several seconds per month. On rare occasion you will need to setup the clock in BIOS Setting as all operating systems allow to change these settings within their environments.

◆ **Daylight Saving:** Allows the clock to automatically adapt to the daylight saving scheme which is removing one hour on the last Sunday of October and adding one hour on the last Sunday of April.

◆ **Hard disk C type:** The number of your primary (master) hard drive. Most of the time this number is 47, which means that you must specify the drive specs according to your hard drive manual.

◆ **Cyl#:** The number of cylinders on your hard disk.

◆ **Head:** The number of heads.

◆ **WPcom:** Write Precompensation. Older hard drives have the same number of sectors per track at the innermost tracks as at the outermost tracks. This means that the data density at the innermost tracks is higher and thus the bits are lying closer together. Starting with this Cyl# until the end of Cyl#s the writing starts

earlier on the disk. In modern HDs (all AT-BUS and SCSI, Small Computer Systems Interface) this entry is useless. Set it either to -1 or max Cyl (a common value is 65535). For IDE (Integrated Device Electronics) hard drives it is not necessary to enter a WP cylinder. The IDE HDD will ignore it for it has its own parameters inboard.

◆ **LZone:** The address of the landing zone. Same as WPcom. Used in old HDs without an auto-parking feature (MFM, Modified Frequency Modulated, or RLL, Run Length Limited). Set it to 0 or max Cyl#.

◆ **Sect:** The number of sectors per track. It is often 17 for MFM and 26 for RLL HDD. On other types of drives, it will vary.

◆ **Size :** This is automatically calculated according the number of cylinders, heads and sectors. It is in megabytes and applies this formula: $(Hds * Cyl * Sect * 512) / 1048$.

EIDE specifications. With the growing capacity of hard disks on desktop computers, a redefinition of IDE specifications was necessary. The old IDE specification only supported drives up to 528 megabytes, which is the Normal partition setting. In 1994, the EIDE (Enhanced IDE) protocol was designed and now all new motherboards support it. This new protocol uses the LBA (Logic Block Addressing) system which considers logic blocks instead of heads, cylinders and sectors. If your BIOS does not support LBA, several hard disk manufacturers provide drivers to trick the BIOS. You will also find a Large partition setting that can accommodate drives up to 1024 cylinders, but do not support LBA. Unfortunately, many large implementations don't work correctly for drives of over 1GB (there's no good reason why it wouldn't work for much larger drives though). Note that 1024 cylinders native is 528MB. The 528MB limit is the 1024 cyl / 16 head / 63 sector limit.

◆ **Hard disk D type:** The number of your secondary (slave) hard drive. Same procedure than above. Jumpers must be set for an hard drive to perform as slave. Please refer to your hard drive manual. You might also want to refer to the hard disk data file frequently posted to comp.sys.ibm.pc.hardware.storage

Several of the PCI motherboards can now accommodate up to four IDE drives: Primary Master, Primary Slave, Secondary Master and Secondary Slave.

◆ **Floppy drive A:** The type of floppy drive installed for drive A. Frequent

configurations are 1.44 MB (3 1/2 inches), or 1.2 MB (5 1/4). Newer systems have also a 2.88 MB (3 1/2) setting.

◆**Floppy drive B:** The type of floppy drive installed for drive B.

◆**Primary display:** The type of displaying standard you are using, and in case of systems with two video adapters the primary one. The most frequent is VGA/PGA/EGA. Modern computers have VGA (Video Graphics Array). If you have an older black/white display select Mono or Hercules, if your Video adapter card is text only, select MDA.

◆**Keyboard:** Installed. If "not installed" this option sets the BIOS to pass the keyboard test in the POST, allowing to reset a PC without a keyboard (file server, printer server, etc.), without the BIOS producing a keyboard error. As a system administrator, you can uninstall the keyboard as a supplementary security procedure to prevent people messing up with the server.

Changing Your Password

Enable you to change the active password. The default is no password.

Remember your password!!! Write it down somewhere!!! Ask yourself: Do I really need to set a password to access my system and/or the BIOS? (is your brother / sister / kid / employee / colleague that dangerous?) If security is of some minor concern to you, **disabled recommended**. Why not only password protect (or encrypt) some critical files (personal finances - things the IRS should not see, juicy love letters, pornographic images (the thing that Internet is most used for), customer information databases, etc...)? If you **lose your password**, you will have to **erase your CMOS memory** (see the FAQ). Some systems allow you to choose when the password is needed to change the CMOS settings, to boot the machine, etc.

Auto Configuration

All recent motherboards have now an auto-configuration setting leaving much of BIOS setup problems out of the user's hands, such as Bus Clock Speed and Wait States. On the majority of cases it will do just fine. But you must remember, it is **not an optimization** of your system's performances, but a set of **efficient**

settings that will insure a good result. You will have to **Disable** this setting if you want to alter the BIOS yourself, otherwise your settings **will be ignored**. On some systems, you may get **supplementary performances** by improving over auto configuration settings, but on others auto configuration is **all you will ever need**.

Auto Configuration with BIOS Defaults

The BIOS defaults may not be tuned for your motherboard/chipset, but give a reasonable chance of getting into POST. Usually these settings are **a good start to fine tune** your system. If you did **something wrong** and don't know what, **select this**. It will replace your BIOS settings by default values. You will have to start all over again. Be sure to know your system's configuration. This option does **NOT alter the date, hard disk and floppy disk configurations** in the Standard CMOS setup, so in general you can expect your system to boot without problems after selecting this.

Auto Configuration with Power-on Defaults

When powering on, the BIOS puts the system in the **most conservative state** you can think of. Turbo off, all caches disabled, all wait states to maximum, etc... This is to make sure that you can always enter BIOS setup. Useful if the settings obtained by selecting AUTO CONFIGURATION WITH BIOS DEFAULTS fail. If the system **does not work** with these values, **it's time to panic**: the problem may be hardware-related (DIP switches, cards not inserted properly or worst, something broken).

Exiting BIOS

There are two ways to exit BIOS settings.

◆ **Write to CMOS and Exit**: Save the changes you made in the CMOS. You must do that to permanently keep your configuration. Several users say they changed the CMOS setup but forgot to exit with this one! A common source of error.

◆ **Do Not Write to CMOS and Exit**: If you are not sure of the changes you made in the CMOS settings, use this option to exit safely.

Advanced CMOS Setup

May vary according to your system, BIOS version and brand. Some functions may **not be present or the order and name may be different** (particularly for different BIOS brand). Know **EXACTLY** what you are doing. Some configurations may keep your computer **off from booting**. If that's the case: Switch the power off. Turn your computer on **WHILE** keeping the **DEL** key pressed. This is supposed to erase the BIOS memory. If it still doesn't boot, consult your motherboard manual. Look for a "forget CMOS RAM" jumper. Set it. Try it again. If it still doesn't boot, ask a friend or post to a computer hardware newsgroup. You are permitted to panic.

◆ **Typematic Rate Programming:** Disabled recommended. It enables the typematic rate programming of the keyboard. Not all keyboards support this! The following two entries specify how the keyboard is programmed if enabled.

◆ **Typematic Rate Delay (msec):** 500 ns recommended. The initial delay before key auto-repeat starts, that is how long you've got to press a key before it starts repeating.

◆ **Typematic Rate (Chars/Sec):** 15. It is the frequency of the auto-repeat i.e. how fast a key repeats.

◆ **Above 1 MB Memory Test:** If you want the system to check the memory above 1 MB for errors. **Disabled recommended** for faster boot sequence. The HIMEM.SYS driver for DOS 6.2 now verifies the XMS (Extended Memory Specification), so this test is redundant. It is thus preferable to use the XMS test provided by HIMEM.SYS since it is operating in the real environment (where user wait states and other are operational).

◆ **Memory Test Tick Sound:** Enabled recommended. It gives an audio record that the boot sequence is working properly. Plus, it is an aural confirmation of your CPU clock speed/Turbo switch setting. An experimented user can hear if something is wrong with the system just by the memory test tick sound.

◆ **Memory Parity Error Check: Enabled recommended.** Additional feature to test bit errors in the memory. All (or almost all) PCs are checking their memory

during operation. Every byte in memory has another ninth bit, that with every write access is set in such way that the parity of all bytes is odd. With every read access the parity of a byte is checked for this odd parity. If a parity error occurs, the NMI (Non Maskable Interrupt), an interrupt you mostly cannot switch off, so the computer stops his work and displays a RAM failure) becomes active and forces the CPU to enter an interrupt handler, mostly writing something like this on the screen: PARITY ERROR AT 0AB5:00BE SYSTEM HALTED. On some motherboards you can disable parity checking with standard memory. Enabled to be sure data from memory are correct. Disable only if you have 8-bit RAM, which some vendors use because it is 10% cheaper. If you own a Gravis Ultrasound Soundcard (GUS), it's imperative that this is enabled; otherwise the Sound Blaster emulation won't work(!).

Note: Be sure to have memory chips of the same speed installed. It is not uncommon to have system crashes simply because memory SIMMS are of different speed. Faster memory will not adapt itself to the speed of slower memory. 60 ns and 80 ns SIMMS will surely make your system crash and yourself wonder what is the problem (I know).

◆ **Hard Disk Type 47 RAM Area:** The BIOS has to place the HD type 47 data somewhere in memory. You can choose between DOS memory or PC BIOS (or peripheral card) memory area 0:300. DOS memory is valuable, you only have 640KB of it. So you should try to use 0:300 memory area instead. There may be some peripheral card which needs this area too (sound card, network card, whatever). So if there are some fancy cards in your PC, check the manuals if they're using the 0:300 area. But in most cases this will work without checking. This is redundant if BIOS is shadowed (maybe not in very old BIOSes). The RAM area can be verified by checking address of int41h and int46h. These are fixed disk parameters blocks. If they point to the BIOS area, BIOS made modification of parameters before mapping RAM there.

◆ **Wait for <F1> If Any Error:** When the boot sequence encounter an error it asks you to press F1. Only at 'non-fatal' errors. If disabled, the system prints a warning and continues to boot without waiting for you to press any keys. Enabled recommended. Disabled if you want the system to operate as a server without a keyboard.

◆ **System Boot Up Num Lock:** Specify if you want the Num Lock key to be activated at boot up. Some like it, some do not. MS-DOS (starting with 6.0,

maybe earlier) allows a "NUMLOCK=" directive in config.sys, too; if someone turns the BIOS flag off but has NUMLOCK=ON in their configuration file, they may be a bit perturbed.

◆ **Numeric Processor Test:** Enabled if you have a math coprocessor (built in for the 486DX, 486DX2, 486DX3 and Pentium - 586 - family). Disabled if you don't (386SX, 386DX, 486SX, 486SLC and 486DLC). If disabled, your FPU (Floating Point Unit, if present) isn't recognized as present by the system and will therefore significantly decrease the performance of your system.

◆ **Weitek Coprocessor:** If you have Weitek FPU, enable. If you have not, disable. This high performance FPU has 2-3 times the performance of the Intel FPU. Weitek uses some RAM address space, so memory from this region must be remapped somewhere else.

◆ **Floppy Drive Seek at Boot:** Power up your A: floppy drive at boot. Disabled recommended for faster boot sequence and for reduced damage to heads. Disabling the floppy drive, changing the system boot sequence and setting a BIOS password are good techniques for adding some security to a PC.

◆ **System Boot Sequence:** What drive the system checks first for an operating system. C:, A: recommended for faster boot sequence, or to not allow any user to enter your system by booting from the FDD if your autoexec.bat starts with a login procedure. A:, C: if the person who uses the computer is someone who don't knows how to setup CMOS. Because if something fails and a boot floppy won't work, many users won't know what to do next. However, be careful. You had better know this setting is turned on and be prepared to turn it off if your hard disk boot track becomes corrupted, but not obviously absent, since you otherwise won't be able to boot from floppy. Also, it's easy to fool yourself into thinking you booted from a known virus-free floppy when it actually booted from the (virus-infested) hard drive.

◆ **System Boot Up CPU speed:** Specify at what processor speed the system will boot from. Usual settings are HIGH and LOW. HIGH recommended. If you encounter booting problems, you may try LOW. You may also change the CPU speed with Ctrl-Alt +.

◆ **External Cache Memory:** Enabled if you have external cache memory (better known as L2 cache memory). This is a frequent error in CMOS setup as

if Disabled when you have cache memory, the system performance decreases significantly. Most systems have from 64K to 512K of external cache. It is a cache between the CPU and the system bus. Different operating systems may address different levels of cache memory. For instance, DOS and Windows can address up to 64K at one time while Windows 95, OS/2 and Windows NT can address larger memory spaces. So, don't buy 256K of cache if you are using a DOS environment with less than 8MB of memory. It will not improve much the performance of your system. If Enabled when the system does not have cache memory, the system will **freeze** most of the time.

◆ **Internal Cache Memory:** Enable or disable the internal cache memory of the CPU (better known as L1 cache memory). Disabled for 386 and Enabled for 486 (1 to 8KB of internal CPU cache). If the CPU does not have internal cache, the system may freeze if enabled.

In many AMI and AWARD BIOSes, the two previous options are implemented either as separate Internal and External Enable/Disable options, or as a single option (Cache Memory : Disabled/Internal/Both).

◆ **CPU Internal Cache:** same as above.

◆ **Fast Gate A20 Option:** A20 refers to the first 64KB of extended memory (A0 to A19) known as the "high memory area". This option uses the fast gate A20 line, supported in some chipsets, to access memory above 1 MB. Normally all RAM access above 1 MB is handled through the keyboard controller chip (8042 or 8742). Using this option will make the access faster than the normal method. This option is very useful in networking and multitasking operating systems.

◆ **Turbo Switch Function:** Enables or disables the turbo switch. Disabled recommended.

◆ **Shadow Memory Cacheable:** You increase speed by copying ROM to RAM. Do you want to increase it by cacheing it? Yes or no - see Video BIOS Area cacheable. Yes recommended for MS-DOS and OS/2. Linux and other Unix-like operating systems will not use the cached ROMs and will benefit from the additional available memory if they are not cached.

◆ **Password Checking Option:** Setup password to have access to the system and / or to the setup menu. Good if the computer is to be shared with several persons and you don't want anyone (friends, sister, etc.) to mess up with the

BIOS. Default password: AMI (if you have AMI BIOS). Award: BIOSTAR or AWARD_SW for newer versions (Note: I even know a computer store that kept standard AWARD BIOS configuration with their systems because they didn't know what the default password was!).

◆ **Video ROM Shadow C000, 32K:** Memory hidden under the "I/O hole" from 0x0A0000 to 0x0FFFFFFF may be used to "shadow" ROM (Read-Only Memory). Doing so, the contents of the ROM are copied into the RAM and the RAM is used instead, which is obviously faster. Video BIOS is stored in slow EPROM (Erasable Programmable Read-Only Memory) chips (120 to 150ns of access time). Also, ROM is 8 or 16 bit while RAM 32 bit wide access. With Shadow on, the EPROM content is copied to RAM (60 to 80ns of access time with 32 bit wide access). Therefore performance increases significantly. Only sensible on EGA/VGA systems. **Enabled recommended.** If you have flash BIOS (EEPROM), you can disable it. Flash BIOS enables access at speeds similar to memory access so you can use the memory elsewhere. However, flash BIOS is still only accessing it at the speed of the bus (ISA, EISA or VLB). On systems where the BIOS automatically steals 384K of RAM anyway, it shouldn't hurt to enable shadowing even on flash ROM. One side effect is that you will not be able to modify the contents of flash ROM when the chip is shadowed. If you reconfigure an adapter which you think might have flash ROM, and your changes are ignored, or of course if it gives you an error message when you try to change them, you'll need to temporarily disable shadowing that adapter. On (S)VGA you should enable both video shadows. Some video cards maybe using different addresses than C000 and C400. If it is the case, you should use supplied utilities that will shadow the video BIOS, in which case you should disable this setting in the CMOS. Video BIOS shadowing can cause software like XFree86 (the free X Window System) to hang. They should be probably be disabled if you run any of the 386 unices.

Some cards map BIOS or other memory not only in the usual a0000-ffff address range, but also just below the 16MB border or at other places. The BIOS (for PCI buses only?) now allows to create a hole in the address range where the card sits. The hole may be enabled by giving an address, then a size is requested in power of 2, 64k - 1MB.

◆ **Adaptor ROM Shadow C800,16K:** Disabled. Those addresses (C800 to EC00) are for special cards, e.g. network and controllers. Enable only if you've got an adapter card with ROM in one of these areas. It is a BAD idea to use

shadow RAM for memory areas that aren't really ROM, e.g. network card buffers and other memory-mapped devices. This may interfere with the card's operation. To intelligently set these options you need to know what cards use what addresses. Most secondary display cards (like MDA and Hercules) use the ROM C800 address. Since they are slow, shadowing this address would improve their performance. An advanced tip: in some setups it is possible to enable shadow RAM **without** write-protecting it; with a small driver (UMM) it is then possible to use this 'shadow RAM' as UMB (Upper Memory Block) space. This has speed advantages over UMB space provided by EMM386.

◆ **Adaptor ROM Shadow CC00,16K:** Disabled. Some hard drive adapters use that address.

◆ **Adaptor ROM Shadow D000,16K:** Disabled. D000 is the default Address for most Network Interface Cards.

◆ **Adaptor ROM Shadow D400,16K:** Disabled. Some special controllers for four floppy drives have a BIOS ROM at D400..D7FF.

◆ **Adaptor ROM Shadow D800,16K:** Disabled

◆ **Adaptor ROM Shadow DC00,16K:** Disabled

◆ **Adaptor ROM Shadow E000,16K:** Disabled. E000 is a good "out of the way" place to put the EMS page frame. If necessary.

◆ **Adaptor ROM Shadow E400,16K:** Disabled

◆ **Adaptor ROM Shadow E800,16K:** Disabled

◆ **Adaptor ROM Shadow EC00,16K:** Disabled. SCSI controller cards with their own BIOS could be accelerated by using Shadow RAM. Some SCSI controllers do have some RAM areas too, so it depends on the brand.

Some SCSI adapters do not use I/O-Addresses. The BIOS address range contains writable addresses, which in fact are the I/O-ports. This means: this address must not be shadowed and even not be cached.

◆ **System ROM Shadow F000, 64K:** Same thing as Video shadow, but according to the system BIOS (main computer BIOS). **Enabled recommended**

for improved performance. System BIOS shadowing and caching should be disabled to run anything but DOS (Windows).

On older BIOS versions the shadow choices are in 400(hex)-byte increments. For instance, instead of one Video ROM Shadow segment of 32K, you will have two 16K segments (C400 and C800). Same thing for Adaptor ROM Shadow segments.

◆ **BootSector Virus Protection:** It is **not** exactly a **virus protection**. All it does is whenever your boot sector is accessed for writing, it gives a warning to the screen allowing you to disable the access or to continue. Extremely annoying if you use something like OS/2 Boot Manager that needs to write to it. It is completely useless for SCSI or ESDI (Enhanced Small Device Interface) drives as they use their own BIOS on the controller. **Disabled recommended.** If you want virus protection, use a TSR (Terminate and Stay Resident) virus detection (Norton, Central Point, etc...). Scan by Macfee is also a good idea. Available on most FTP servers, it is a shareware.

Advanced Chipset Setup

Configurations may vary according to your system, BIOS version and brand. So, some setting may be present on your computer, some may not. Be sure of what you are doing!

◆ **Automatic Configuration:** Allows the BIOS to set automatically several important settings (e.g. Clock divider, wait states, etc.). Very useful for newbies. Disabled recommended if you want to play around with the settings. If you have some special adapter cards, you will also have to disable this option.

◆ **Keyboard Reset Control:** Enable Ctrl-Alt-Del warm reboot. Enabled recommended for more control over your system.

Refresh

◆ **Hidden Refresh:** Allows the RAM refresh memory cycles to take place in memory banks not used by your CPU at this time, instead or together with the normal refresh cycles, which are executed every time a certain interrupt (DRQ0 every 15 ms) is called by a certain timer (OUT1). Every time it takes 2 to 4 ms for the refresh. One refresh cycle every ~16 us refreshes 256 rows in ~ 4ms. Each refresh cycle only takes the equivalent of one memory read or less, as CAS (Column Address Strobe) is not needed for a refresh cycle. Some RAM can do it, some not. Try. If the computer fails, turn it off. Enabled recommended. There are typically 3 types of refresh schemes: cycle steal, cycle stretch, or hidden refresh. Cycle steal actually steals a clock cycle from the CPU to do the refresh. Cycle stretch actually delays a cycle from the processor to do the refresh. Since it only occurs every say 4ms or so, it's an improvement from cycle steal. We're not really stealing a cycle, only stretching one. Hidden refresh typically doesn't stretch or steal anything. It's usually tied to DTACK (Data acknowledge) or ALE (Address Latch Enable) or some other signal relating to memory access. Since memory is accessed ALL of the time it is easy to synchronize the refresh on the falling edge of this event. Of course, the system performance is at its optimum efficiency, refresh wise since we're not taking clock cycles away from the CPU.

◆ **Slow Refresh:** Causes RAM refresh to happen less often than usual, around four times. This increases the performance slightly due to the reduced contention

between the CPU and refresh circuitry, but not all DRAM memories necessarily support these reduced refresh rates (in which case you will get parity errors and crashes). It also saves power, a good opportunity for laptop computers. **Enabled recommended**

◆ **Concurrent Refresh:** Both the processor and the refresh hardware have access to the memory at the same time. If you switch this off, the processor has to wait until the refresh hardware has finished (it's a lot slower). **Enabled recommended.**

◆ **Burst Refresh:** Performs several refresh cycles at once. Increase the system performance.

◆ **DRAM Burst at 4 Refresh:** Refresh is occurring at Bursts of four, increasing the system performance.

◆ **Hi-speed Refresh:** Refreshes are occurring at a higher frequency, which is improving the system performance. Of course, not all types of memory can support it and Slow Refresh is preferred.

◆ **Staggered Refresh:** Refresh is performed on memory banks sequentially. The advantages are related to less power consumption and less interference between memory banks.

◆ **Slow Memory Refresh Divider:** The AT refresh cycle occurs normally every 16 ns, straining the CPU. If you can select a higher value, such as 64 ns, you will increase the performance of your system.

◆ **Decoupled Refresh Option:** Enables the ISA bus and the RAM to refresh separately. Because refreshing the ISA bus is more slow, this causes less strain on the CPU.

◆ **Refresh Value:** The lower this value is, the best the performance.

◆ **Refresh RAS Active Time:** The amount of active time needed for Row Address Strobe during refresh. The lower the better.

Data Bus

◆ **Single ALE Enable:** Address Latch Enable (ALE) is an ISA Bus Signal (Pin

B28) that indicates that a valid address is posted on the bus. The bus is used to communicate with 8 and 16 bit peripheral cards. Some chipsets have the capability to support an enhanced mode in which multiple ALE assertions may be made during a single Bus Cycle. Single ALE Enable apparently enables or disables that capability. May slow the video bus speed if enabled. Disabled (No) recommended.

◆ **AT BUS Clock Selection (or AT Bus Clock Source):** Gives a division of the CPU clock (or System Clock) so it can reach the ISA - EISA bus clock. An improper setting may cause significant decrease in performance. The settings are in terms of CLK/x, (or CLKIN/x and CLK2/x) where x may have values like 2, 3, 4, 5, etc. CLK represents your processor speed, with the exception that clock-multiple processors need to use the EXTERNAL clock rate, so a 486DX33, 486DX2/66, and 486DX3/99 all count as 33 and should have a divider value of 4. For 286 and 386 processors, CLK is half the speed of the CPU. You should try to reach 8.33 Mhz (that's the old bus clock of IBM AT; there may be cards which could do higher, but it's not highly recommended). On some motherboards, the AT bus speed is 7.15 Mhz. On new BIOS versions, there is an AUTO setting that will look at the clock frequency and determine the proper divider.

CLK /3	SX / DX16, DX 20, DX 25, DX2 / 50
CLK /4	SX / DX33, DX2 / 66, DX3 / 99
CLK /5	DX 40, DX2 / 80
CLK /6	DX 50, DX2 / 100

Selecting the right clock divider.

You can try other clock settings to **increase performance**. If you choose a too small divider (CLK/2 for a DX33) your system may hang. For a too big divider (CLK/5 for a DX33) the performance of ISA cards will decrease. This setting is for data exchange with ISA cards, NOT VL bus and PCI cards which run at CPU bus clock speeds: 25Mhz, 33Mhz and higher. If your ISA cards are fast enough to keep up, it is possible to run the bus at 12 Mhz. Note that if you switch crystals to overclock your CPU, you are also overclocking the ISA bus unless you change settings to compensate. Just because you can overclock the CPU doesn't mean you can get away with overclocking the ISA bus. It might just be one card that causes trouble, but one is enough. It might cause trouble even if you aren't using it by responding when it shouldn't.

◆ **ISA Bus Speed:** As above, but related for PCI.

◆ **Bus Mode:** It can be set in synchronous and asynchronous modes. In synchronous mode, the CPU clock is used, while in asynchronous mode the ATCLK is used.

◆ **AT Cycle Wait State:** Whenever an operation is performed with the AT bus, it indicates the number of wait states inserted. You may need some wait states if old ISA cards are used, notably if they are in operation with fast adapter cards.

◆ **16-bit Memory, I/O Wait State:** The number of wait states before 16-bit memory and I/O operations.

◆ **8-bit Memory, I/O Wait State:** As above, except this setting is for 8-bit operations.

◆ **16-bit I/O Recovery Time:** The additional delay time inserted after every 16-bit operations. This value is added to the minimum delay inserted after every AT cycles.

◆ **Fast AT Cycle:** If enabled, may speed up transfer rates with ISA cards, notably video.

◆ **ISA IRQ:** Inform the PCI cards of the IRQs used by ISA cards, so they be discarded.

◆ **DMA Wait States:** The number of wait states inserted before direct memory access (DMA). The lower the better.

◆ **DMA Clock Source:** The source of the DMA clock for which some peripheral controllers, like floppy, tape, network and SCSI adapters use to address memory, which is 5 MHz maximum.

◆ **E0000 ROM belongs to ATBUS:** Tells if the E0000 area (upper memory) belongs to the MB DRAM or to the AT bus. Yes recommended.

◆ **Memory Remapping:** Remaps the memory used by the BIOS (A0000 to FFFF - 384 k) above the 1 Mb limit. If enabled you cannot shadow Video and System BIOS. Disabled recommended.

◆ **Fast Decode Enable:** Refers to some hardware that monitors the commands sent to the keyboard controller chip. The original AT used special codes not processed by the keyboard itself to control the switching of the 286 processor back from protected mode to real mode. The 286 had no hardware to do this, so they actually have to reset the CPU to switch back. This was not a speedy operation in the original AT, since IBM never expected that an OS might need to jump back and forth between real and protected modes. Clone makers added a few PLD chips to monitor the commands sent to the keyboard controller chip, and when the "reset CPU" code was seen, the PLD chips did an immediate reset, rather than waiting for the keyboard controller chip to poll its input, recognize the reset code, and then shut down the CPU for a short period. This "fast decode" of the keyboard reset command allowed OS/2 and Windows to switch between real and protected mode faster, and gave much better performance. (early 286 clones with Phoenix 286 BIOS had this setting to enable/disable the fast decode logic.) On 386 and newer processors, the fast decode is probably not used, since these CPUs have hardware instructions for switching between modes. There is another possible definition of the "Fast Decode Enable" command. The design of the original AT bus made it very difficult to mix 8-bit and 16-bit RAM or ROM within the same 128K block of high address space. Thus, an 8-bit BIOS ROM on a VGA card forced all other peripherals using the C000-Dfff range to also use 8 bits. By doing an "early decode" of the high address lines along with the 8/16 bit select flag, the I/O bus could then use mixed 8 and 16 bit peripherals. It is possible that on later systems, this BIOS flag controls the "fast decode" on these address lines.

◆ **Extended I/O Decode:** The normal range of I/O addresses is 0-0x3ff; 10 bits of I/O address space. Extended I/O-decode enables wider I/O-address bus. The CPU support a 64K I/O space, 16 address lines. Most motherboards or I/O adapters can be decoded only by 10 address bits.

◆ **I/O Recovery Time:** I/O recovery time is the number of wait states to be inserted between two consecutive I/O operations. It is generally specified as a two number pair -- e.g. 5/3. The first number is the number of wait states to insert on an 8 bit operation, the second the number of waits on a 16 bit operation. A few BIOSes specify an I/O Setup time (AT Bus (I/O) Command Delay). It is specified similarly to IO Recovery Time, but is a delay before STARTING an I/O operation rather than a delay BETWEEN I/O operations. 5/3 has been recommended as a value which will often yield a good combination of performance and reliability. When enabled, more I/O wait states are inserted. A

transfer from IDE hard drive to memory happens without any handshaking, meaning the data has to be present (in the cache of the hard disk) when the CPU wants to read them from an I/O Port. This is called PIO (Programmed I/O) and works with a REP INSW assembler instruction. Now I/O Recovery Time enabled adds some wait states to this instruction. When disabled, the hard drive is a lot faster. Note that there is a connection between I/O Recovery Time and AT BUS Clock Selection. For example, if the AT BUS Clock is set to 8 MHz and you have a normal hard disk, I/O Recovery Time can be turned off, resulting in a higher transfer rate from hard disk.

◆ **IDE Multi Block Mode:** Enable IDE drives to transfer several sectors per interrupt. According to the hard drive cache size, six modes are possible. Mode 0 (standard mode transferring a single sector at a time), Mode 1 (no interrupts), Mode 2 (Sectors are transferred in a single burst), Mode 3 (32-bit instructions with speeds up to 11.1 Mb/sec. In BIOSes usually abbreviated as "32-bit mode". Not to be confused with 32-bit protected mode instructions (!) or Windows' 32-bit disk access.), Mode 4 (up to 16,7 Mb/sec.) and Mode 5 (up to 20 Mb/sec.). The so-called "PIO mode 5" is completely bogus. It was launched by some controller manufacturers but was never accepted, never absorbed into the standards and you will not find any disk drives supporting it. Nor will you find any such drives in the future. The relevant parameter for block mode is the number of sectors per interrupt. The maximum number of sectors per interrupt is often (but not always) related to the drive's buffer size. If this setting is not set properly, communication with COM ports may not work properly. If the block size (sectors/interrupt) is set to too large a value, you may experience serial port overruns and CRC errors. To fix this, decrease the block size (preferred) or disable block mode altogether.

◆ **IDE DMA Transfer Mode:** Settings are Disabled, Type B (for EISA) and Standard (for PCI). Standard is the fastest but may cause problems with IDE CD ROM.

◆ **IDE Multiple Sector Mode:** When IDE DMA Transfer Mode is enabled, this sets the number of sectors per burst, with a maximum of 64. Problems may occur with COM ports.

- **IDE Block Mode:** Enables multi-sectors transfers.

Warning. This setting is known to cause crashes in Win95. Disabled recommended. Extremely annoying.

◆ **IDE 32-bit Transfer:** When enabled, the read / write rate of the hard disk is faster. When disabled only 16-bit data transfers is possible. The read/write rate of the harddisk stays the same, but the transfers over the host bus are maybe faster. So, don't expect anything really dramatic. Actually, you should ordinarily expect no difference at all, since even with 16-bit transfers, the local bus is fast enough to accomodate just about any disk drive. However, some interface hardware uses faster timing on the ATA (IDE) bus when 32-bit transfers are used. In those cases you may notice a speedup. Note that ATA (IDE) is a 16-bit bus. The 32-bit transfers referred to here are strictly the transfers between CPU and interface chip.

◆ **Extended DMA Registers:** Within a AT, DMA occurs for 16 Mb. When enabled, DMA covers the whole 4 Gb of a 32-bit processor.

Cacheing

◆ **Cache Read Option:** Often referred as "**SRAM Read wait state**" or "**Cache Read Hit Burst**" (SRAM: Static Random Access Memory). A specification of the number of clocks needed to load four 32-bit words into a CPU internal cache. Typically specified as clocks per word. 2-1-1-1 indicates 5 clocks to load the four words and is the theoretical minimum for current high end CPUs (486DX, 486SX, 486DX2, DX4, Pentium). Conceptually, the m-n-n-n notation is narrowly limited to CPUs supporting burst mode and with caches organized as 4 word "lines". However it would not be a surprise to see it extended to other CPU architectures. It takes simple integer values, like 2-1-1-1, 3-1-1-1 or 3-2-2-2. This determines the number of wait states for the cache RAM in normal and burst transfers (the latter for 486 only). The lower you computer can support, the better. 4-1-1-1 is usually recommended.

◆ **Cache Write Option:** Same thing as memory wait states, but according to cache ram.

◆ **Fast Cache Read/Write:** Enable if you have two banks of cache, 64K or 256K.

◆ **Cache Wait State:** Like conventional memory, the lower wait states for your cache, the better. 0 will give the optimal performance, but 1 wait state may be required for bus speed higher than 33 MHz.

◆ **Tag Ram Includes Dirty:** Enabling will cause an increase in performance,

because the cache is not replaced during cycles, simply written over. It will usually cut the maximum cachable range in half, as one bit is taken off the address tag in order to be used as a dirty tag bit. So, if you have a lot of memory, you might be better off without dirty tag bit.

◆ **Non-Cacheable Block-1 Size:** Disabled. The Non-Cacheable region is intended for a memory-mapped I/O device that isn't supposed to be cached. For example, some video cards can present all video memory at 15 Mb - 16 Mb so software doesn't have to bank-switch. If the non-cacheable region covers actual RAM memory you are using, expect a significant performance decrease for accesses to that area. If the non-cacheable region covers only non-existent memory addresses, don't worry about it. If you don't want to cache some memory you can exclude 2 regions of memory. There are good reasons not to cache some memory areas. For example, if the memory area corresponds to some kind of buffer memory on a card so that the card may alter the contents of this buffer without warning the cache to invalidate the corresponding cache lines. Some BIOSes take more options than enabled /disabled, namely Nonlocal /Noncache /Disabled (VLB only?).

◆ **Non-Cacheable Block-1 Base:** 0KB. Enter the base address of the area you don't want to cache. It must be a multiple of the Non-Cacheable Block-1 Size selected.

◆ **Non-Cacheable Block-2 Size:** Disabled.

◆ **Non-Cacheable Block-2 Base:** 0KB.

◆ **Cacheable RAM Address Range:** Usually chipsets allow memory to be cached just up to 16 or 32 MB. This is to limit the number of bits of a memory address that need to be saved in the cache together with its contents. If you only have 4MB of RAM, select 4MB here. The lower the better, don't enter 16MB if you only have 8MB installed!

◆ **Video BIOS Area Cacheable:** To cache or not to cache video BIOS, a good question. You should try what is better - video access is faster with 'enabled', but cache has its size. With an "accelerated" video card it may be necessary to make the video RAM region non-cacheable so the CPU can see any changes the drawing engine makes in the frame buffer.

Memory

◆ **Memory Read Wait State:** (often referred as **DRAM Wait States**) The CPU is often much faster than the memory access time. On a 486, 1 or more wait states are often required for RAM with 80ns or higher access time. And, depending on the processor and motherboard, also for lower than 80ns access time. The less wait states, the better. Consult your manual. If wait states are too low, a parity error will occur. For 386 or 486 non-burst memory access cycle takes 2 clock ticks. A **rough** indication of RAM speed necessary for 0 wait states is $2000/\text{Clock}[\text{MHz}] - 10$ [ns]. For a 33Mhz processor, this would give 50ns of access time required. The number of wait states necessary is **approximately** $(\text{RamSpeed}[\text{ns}] + 10) * \text{Clock}[\text{MHz}] / 1000 - 2$. For 70ns RAM and a 33Mhz processor (very standard configuration), this would give roughly 1 wait state. But this really is dependent on chipset, motherboard and cache design, CPU type and whether we talk about reads or writes. Take these formulas with a large grain of salt. You can find out the access time of your RAM chips by looking at their product numbers. Mostly at the end there is a 70, 80, 90, or even 60. If 10 stands there, it means 100 ns. Some RAM chips also have an explicitly written speed in ns. The RAM you buy these days mostly have 70ns or 60ns. Each wait states adds 30 ns of RAM access speed.

◆ **Memory Write Wait State:** Same as above.

In some BIOSes, these two options are combined as 'DRAM Wait State'. In that case, the number of read and write wait states is necessarily equal.

◆ **DRAM CAS Timing Delay:** The default is no CAS delay. DRAM is organized by rows and columns and accessed through strobcs. Then a memory read/write is performed, the CPU activates RAS (Row Access Strobe) to find the row containing the required data. Afterwards, a CAS (Column Access Strobe) specifies the column. RAS and CAS are used to identify a location in a DRAM chip. RAS access is the speed of the chip while CAS is half the speed. When you have slow DRAM, you should use 1 state delay.

◆ **DRAM Refresh Method:** Selects the timing pulse width of RAS from RAS Only or CAS before RAS (which one is better?).

◆ **RAS Precharge Time:** Technically, this is the duration of the time interval during which the Row Address Strobe signal to a DRAM is held low during

normal Read and Write Cycles. This is the minimum interval between completing one read or write and starting another from the same (non-page mode) DRAM. Techniques such as memory interleaving, or use of Page Mode DRAM are often used to avoid this delay. Some chipsets require this parameter in order to set up the memory configuration properly. The RAS Precharge value is typically about the same as the RAM Access (data read/write) time. The latter can be used as an estimate if the actual value is unavailable. At least one BIOS describes the precharge and access times as RAS LOW and RAS HIGH Times. For a 33 MHz CPU, 4 is a good choice, while lower values should be selected for slower speeds.

◆ **RAS Active Time:** The amount of time a RAS can be kept open for multiple accesses. High figures will improve performance.

◆ **RAS to CAS Delay Time:** Amount of time a CAS is performed after a RAS. The lower the better, but some DRAM will not support low figures.

◆ **CAS Before RAS:** Reduces refresh cycles and power consumption.

◆ **CAS Width in Read Cycle:** The number of wait states for the CPU to read DRAM. The lower the better.

◆ **Interleave Mode:** Controls how the CPU access different DRAM banks.

◆ **Fast Page Mode DRAM:** This speeds up memory access for DRAM capable of handling it (most do). When access occurs in the same memory area, RAS and CAS are not necessary.

Plug and Play/PCI

A system intended to make fitting of expansion cards easier (yes, really!). In this context, ISA cards are known as Legacy Cards, and are switched as normal to make them fit in. Have as few of these as possible, as accesses to them are slow. With Concurrent PCI, The T II (or 430HX/VX) chipset's Multi Transaction Timer allows multiple transfers in one PCI request, by reducing re-arbitration when several PCI processes can take place at once. Passive Release allows the PCI bus to continue working when it's receiving data from ISA devices, which would normally hog the bus. Delayed Transaction allows PCI bus masters to work by delaying transmissions to ISA cards. Write merging combines byte, word and Dword cycles into a single write to memory.

The idea is that plug and play cards get interrogated by the system they are plugged into, and their requirements checked against those of the cards already in there. The BIOS will feed the data as required to the Operating System, typically Windows '95. Here you will be able to assign IRQs, etc to PCI slots and map PCI INT#s to them. Although Windows '95 or a PnP BIOS can do a lot by themselves, you really need the lot, e.g. a Plug and Play BIOS, with compatible devices and an Operating System for the best performance. Be aware that not all PCI (2.0) cards are PnP. PC (PCMCIA) cards are also "Plug and Play", but are not considered here.

PnP itself was originally devised by Compaq, Intel and Phoenix. Your chipset settings may allow you to choose of two methods of operation:

- All PnP devices are configured and activated.
- All PnP ISA cards are isolated and checked, but only those needed to boot the machine are activated. The ISA system cannot produce specific information about a card, so the BIOS has to isolate each one and give it a temporary handle so its requirements can be read. Resources can be allocated once all cards have been dealt with (recommended for Windows '95, as it can use the Registry and its own procedures to use the same information every time you boot).

ESCD (Extended System Configuration Data), a system which is part of PnP (actually a superset of EISA), that can store data on PnP or non-PnP EISA, ISA

or PCI cards to perform the same function as the Windows '95 Registry above, that is, provide consistency between sessions. It occupies part of Upper Memory (E000-EDFF), which is not available to memory managers. The default length is 4K, and problems have been reported with EMS buffer addressing when this area has been used.

PCI Slot Configuration

Although an unlimited number of PCI slots is allowed, in practice 4 is the maximum, due to loading considerations. PCI cards and slots use an internal interrupt system, with each slot being able to activate up to 4, labelled either INT#A-INT#D, or INT#1-INT#4. These are nothing to do with IRQs, although they can be mapped to them if the card concerned needs it. Typically IRQs 9 and 10 are reserved for this, but any available ones can be used.

- **Latency Timer (PCI Clocks).** Controls the length of time an agent on the PCI bus can hold the bus when another has requested it, so everything gets its fair share. Since the PCI bus runs faster than the ISA bus, the PCI bus must be slowed during interactions with it. This setting allows you to define how long the PCI bus will delay for a transaction between the given PCI slot and the ISA bus. This number is dependent on the PCI master device in use and varies from 0 to 255. AMI defaults to 66, but 40 clocks is a good place to start at 33MHz (Phoenix). The shorter the value, the more rapid access to the bus a device gets, with better response times, but the lower becomes the effective bandwidth and hence data throughput. Normally, leave this alone, but you could set it to a lower value if you have latency sensitive cards (e.g. audio cards and/or network cards with small buffers). Increase slightly if I/O sensitive applications are being run.
- Using IRQ. Affected by the Trigger method. With PCI, you assign IRQs, etc to a slot, rather than adjusting the card, but only if the card needs an IRQ. There are two methods of IRQ usage, Level or Edge triggered (see Expansion Cards). Most PCI cards use the former, and ISA the latter.
- PCI Slot x INTx. Assigns PCI INT#s to slots 1/2/3 (or whatever). See Slot

X using INT#, overleaf.

- Edge/Level Select. Programs PCI IRQs to single-edge or logic level. Level or Edge sensitivity is programmed per controller. Select Edge for PCI IDE.
- PCI Device, Slot 1/2/3. Enables I/O and memory cycle decoding.
- Enable. As slave
- En Master. Enables PCI device as bus master.
- Use Default Latency Timer Value. If yes, you don't need Latency Timer (above).
- **Slot X Using INT#.** Selects an INT# channel for a PCI Slot, and there are four (A, B, C & D) for each one, that is, each PCI bus slot supports interrupts A, B, C and D. #A is allocated automatically, and you would only use #B, #C, etc if the card needs to use more than one (PCI) interrupt service. For example, select #D if your card needs four. Using Auto is simplest. Most graphics cards don't need this.
- **Xth Available IRQ.** Selects (or maps) an IRQ for one of the available INT#s above. There are ten selections (3, 4, 5, 6, 7, 9, 10, 11, 12, 14, 15). 1st available IRQ means the BIOS will assign this IRQ to the first PCI slots (order is 1, 2, 3, 4). NA means the IRQ has been assigned to the ISA bus and is therefore not available to a PCI slot.
- **1st-6th Available IRQ.** As above.
- **PCI IRQ Activated by.** The method by which the PCI bus recognises an IRQ request; Level or Edge (see Expansion Cards). Use the default unless advised otherwise by your manufacturer or if you have a PCI device which only recognizes one of them.

- **Configuration Mode.** Sets the method by which information about legacy cards is conveyed to the system.
- Use ICU--the BIOS depends on information provided by Plug and Play software (e.g. Configuration Manager or ISA Configuration Utility). Only set this if you have the utilities concerned.
- Use Setup Utility. The BIOS depends on information provided by you in the following settings. Don't use the above utilities.
- **ISA Shared Memory Size.** Sets a block of system memory which will not be shadowed. Should be disabled, unless you have an ISA card that uses the upper memory area. If you use this setting you will also get the following:
 - ISA Shared Memory Base Address. If you choose 64K, you can only choose D000 or below.
 - **IRQ 3-IRQ 15.** Used to indicate what IRQs are in use by ISA Legacy cards. If not used, set to Available. Otherwise, set Used by ISA Card, which means that nothing else can use it.
 - **PCI IDE Prefetch Buffers.** Disables a set of prefetch buffers in the PCI IDE controller. You may need to do this with an operating system (like NT) that doesn't use the BIOS to access the hard disk and doesn't disable interrupts when completing a programmed I/O operation. Disabling also prevents errors with faulty PCI-IDE interface chips that can corrupt data on the hard disk (with true 32-bit operating systems). Check if you've got a PC-Tech RZ1000 or a CMD PCIO 640, but disabling is done automatically with later boards.
 - **PCI IDE 2nd Channel.** Disable this if you're not using the 2nd channel on the PCI IDE card, or you will lose IRQ 15 on the ISA slots.
 - **PCI IDE IRQ Map to.** Allows you to configure your system to the type of IDE disk controller; an ISA device is assumed. The more apparent

difference is the type of slot being used. However, if you have a PCI IDE controller, this setting allows you to specify which slot has the controller and which PCI INT# (A, B,C or D) is associated with the connected hard drives. Note that this refers to the hard disk rather than individual partitions. Since each IDE controller supports two drives, you can select the INT# for each. Note also that the primary has a lower interrupt than the secondary, as described in Slot x Using INT#.

- **PCI-Auto.** If the IDE is detected by the BIOS on one of the PCI slots, then the appropriate INT# channel will be assigned to IRQ 14.
- **PCI-Slot X.** If the IDE is not detected, you can manually select the slot.
- **Primary IDE INT#, Secondary IDE INT#.** Assigns 2 INT channels for primary and secondary channels, if supported.
- **ISA.** Assigns no IRQs to PCI slots. Use for PCI IDE cards that connect IRQs 14 and 15 directly from an ISA slot using a table from a legacy paddleboard.
- **PCI Bus Parking.** Sort of bus mastering; a device parking on the PCI Bus has full control of the bus for a short time. Improves performance when that device is being used, but excludes others. Try with NICs and Hard Disk Controllers.
- **IDE Buffer for DOS & Win.** For IDE read ahead and posted write buffers, so you can increase throughput to and from IDE devices by buffering reads and writes. Slower IDE devices could end up slower, though.
- **IDE Master (Slave) PIO Mode.** Changes IDE data transfer speed; Mode 0-4, or Auto. PIO means Programmed Input/Output. Rather than have the BIOS issue commands to effect transfers to or from the disk drive, PIO allows the BIOS to tell the controller what it wants, and then lets the controller and the CPU perform the complete task by themselves. Modes 1-

4 are available.

- **HCLK PCICLK.** Host CLK vs PCI CLK divider; AUTO, 1-1, 1-1.5.
- **PCI-ISA BCLK Divider.** PCI Bus CLK vs ISA Bus CLK divider; AUTO, PCICLK1/3, PCICLK1/2, PCICLK1/4.
- **CPU to PCI Byte Merge.** See Byte Merging for explanation (below).
- **PCI Write-byte-Merge.** When enabled, this allows data sent from the CPU to the PCI bus to be held in a buffer. The chipset will then write the data in the buffer to the PCI bus when appropriate.
- **CPU-to-PCI Read Buffer.** When enabled, up to four Dwords can be read from the PCI bus without interrupting the CPU. When disabled, a write buffer is not used and the CPU read cycle will not be completed until the PCI bus signals that it is ready to receive the data. The former is best for performance.
- **PCI-to-CPU Write Buffer.** See above.
- **CPU-to-PCI Read-Line.** When On, more time will be allocated for data setup with faster CPUs. This may only be required if you add an Intel OverDrive processor to your system.
- **CPU-to-PCI Read-Burst.** When enabled, the PCI bus will interpret CPU read cycles as the PCI burst protocol, meaning that back-to-back sequential CPU memory read cycles addressed to the PCI will be translated into fast PCI burst memory cycles. Performance is improved, but some non-standard PCI adapters (e.g. VGA) may have problems.
- **PCI to DRAM Buffer.** Improves PCI to DRAM performance by allowing data to be stored if a destination is busy. Buffers are needed because the PCI bus is divorced from the CPU.
- **Latency for CPU to PCI write.** Delay time before CPU writes data to the PCI bus.
- **PCI Cycle Cache Hit WS.** Similar to above. With the latter, the CPU has less to do, so performance is better.

- Normal--Cache refresh during normal PCI cycles.
- Fast--Cache refresh without PCI cycle for CAS.
- **Use Default Latency Timer Value.** Whether or not the default value for the Latency Timer will be loaded, or the succeeding Latency Timer Value will be used. If Yes is selected (default), no further programming is needed in the Latency Timer Value option (below).
- **Latency Timer Value.** The maximum number of PCI bus clocks that the master may burst. A longer latency time gives the CPU more of a chance to control the bus. See also Latency Timer (PCI Clocks).
- **Latency from ADS# status.** This allows you to configure how long the CPU waits for the Address Data Status (ADS). It determines the CPU to PCI Post write speed. When set to 3T, this is 5T for each double word. With 2T (default), it is 4T per double word. For a Qword PCI memory write, the rate is 7T (2T) or 8T (3T). The default should be OK, but if you add a faster CPU to your system, you may find it necessary to increase it. The choices are:
 - 3T--Three CPU clocks
 - 2T--Two CPU clocks (Default)
- **PCI Master Latency.** If your PCI Master cards control the bus for too long, there is less time for the CPU to control it. A longer latency time gives the CPU more of a chance. Don't use zero.
- **Max burstable range.** The maximum bursting length for each FRAME# asserting. FRAME# is an electrical signal. Dunno what it does, yet.
- **CPU to PCI burst memory write.** If enabled, back-to-back sequential CPU memory write cycles to PCI are translated to PCI burst memory write

cycles. Otherwise, each single write to PCI will have an associated FRAME# sequence. Enabled is best for performance, but some non-standard PCI cards (e.g. VGA) may have problems.

- **Fast Back To Back.** Possibly as above, but working on it!
- **CPU to PCI post memory write.** Enabling allows up to 4 Dwords of data to be posted to PCI. Otherwise, not only is buffering disabled, completion of CPU writes is limited (e.g. CPU write does not complete until the PCI transaction completes). Enabled is best for performance.
- **CPU to PCI Write Buffer.** As above. Buffers are needed because the PCI bus is divorced from the CPU; they improve overall system performance by allowing the processor (or bus master) to do what it needs without writing data to its final destination; the data is temporarily stored in fast buffers.
- **PCI to ISA Write Buffer.** When enabled, the system will temporarily write data to a buffer so the CPU is not interrupted. When disabled, the memory write cycle for the PCI bus will be direct to the slower ISA bus. The former is best for performance.
- **DMA Line Buffer.** Allows DMA data to be stored in a buffer so PCI bus operations are not interrupted. Disabled means that the line buffer for DMA is in single transaction mode. Enabled allows it to operate in an 8-byte transaction mode for greater efficiency.
- **ISA Master Line Buffer.** ISA master buffers are designed to isolate the slower ISA I/O operations from the PCI bus for better performance. Disabled means the buffer for ISA master transaction is in single mode. Enabled means it is in 8-byte mode, increasing the ISA master's performance.
- **CPU/PCI Post Write Delay.** Delay time before the CPU writes data into the PCI bus.
- **Post Write CAS Active.** Pulse width of CAS# when the PCI master writes to DRAM.
- **PCI master accesses shadow RAM.** Enables the shadowing of a ROM on a PCI master for better performance.

- **Enable Master.** Enables the selected device as a PCI bus master and checks whether the card is so capable.
- **AT bus clock frequency.** AT bus speed in a PCI system. Choose whatever divisor gives you a speed of 6-8.33 MHz, depending on the speed of the PCI bus.
- **ISA Bus Clock Frequency.** As above.
- **Base I/O Address.** The base of the I/O address range from which the PCI device resource requests are satisfied.
- **Base Memory Address.** The base of the 32-bit memory address range from which the PCI device resource requests are satisfied.
- **Parity.** Allows parity checking of PCI devices.
- **ISA Linear Frame Buffer.** Set to the appropriate size if you use an ISA card that features a linear frame buffer (e.g. a second video card for ACAD). The address will be set automatically.
- **ISA VGA Frame Buffer Size.** This is to help you use a VGA frame buffer and 16 Mb of RAM at the same time; the system will allow access to the graphics card through a hole in its own memory map; in other words, accesses made to addresses within this hole will be directed to the ISA bus instead of main memory. Should be set to Disabled, unless you are using an ISA card with more than 64K of memory that needs to be accessed by the CPU, and you are not using the Plug and Play utilities. If you have less than 8 Mb memory, or use MS-DOS, this will be ignored.
- **Residence of VGA Card.** Whether on PCI or VL Bus.
- **ISA LFB Size.** LFB=Linear Frame Buffer. See above.
- **Memory Map Hole; Memory Map Hole Start/End Address.** See ISA VGA Frame Buffer Size. Where the hole starts depends on ISA LFB Size. Sometimes this is informative only. If you can change it, base address should be 16Mb, less buffer size.
- **Memory Hole Size.** Options include 1 Mb, 2 Mb, 4 Mb, 8 Mb, Disabled.

These are the amounts below 1 Mb assigned to the AT Bus, and reserved for ISA cards.

- **Memory Hole Start Address.** To improve performance, certain parts of memory are reserved for ISA cards, which must be mapped into the memory space below 16 MB for DMA reasons. The selections are from 1-15 with each number in Mb. This is irrelevant if the memory hole is disabled (see above).
- **Memory Hole at 15-16M.** See above.
- **Local Memory 15-16M.** To increase performance, you can map slower device memory (e.g. on the ISA bus) into much faster local bus memory. Local memory is set aside and the start point transferred from the device memory to local memory. The default is enabled.
- **15-16M Memory Location.** The area in the memory map allocated for ISA option ROMs. Choices are Local (default) or Non-local.
- **Byte Merging.** This exists where writes to sequential memory addresses are merged into one PCI-to-memory operation, which increases performance for older applications that write to video memory in bytes rather than words--not supported on all PCI video cards. Enable unless you get bad graphics. See also next for a variation.
- **Byte Merge Support.** 8- or 16-bit data en route from the CPU to the PCI bus is held in a buffer where it is accumulated, or merged, into 32-bit data, giving faster performance. In this case, enabling means that CPU-PCI writes are buffered (Award).
- **Multimedia Mode.** Enables or disables palette snooping for multimedia cards.
- **Video Palette Snoop.** Controls how a PCI graphics card can "snoop" write cycles to an ISA video card's colour palette registers. Snooping essentially means interfering with a device. Only set to Disabled if:
 - An ISA card connects to a PCI graphics card through a VESA connector

- The ISA card connects to a colour monitor, and
- The ISA card uses the RAMDAC on the PCI card, and
- Palette Snooping (RAMDAC shadowing) not operative on PCI card.
- **PCI/VGA Palette Snoop.** Alters the VGA palette setting while graphic signals pass through the feature connector of PCI VGA card and are processed by MPEG card. Enable if you have MPEG connections through the VGA feature connector; this means you can adjust PCI/VGA palettes. VGA snooping is used by multimedia video devices (e.g. video capture boards) to look ahead at the video controller (VGA device) to see what color palette is currently in use. It is only in exceptional circumstances that you might ever need to enable this, so disable for ordinary systems. (Award BIOS).
- **Snoop Filter.** Saves the need for multiple enquiries to the same line if it was inquired previously. When enabled, cache snoop filters ensure data integrity (cache coherency) while reducing the snoop frequency to a minimum.
- **E8000 32K Accessible.** The 64K E area of upper memory is used for BIOS purposes on PS/2s, 32 bit operating systems and Plug and Play. This setting allows the second 32K page to be used for other purposes when not needed, in the same way that the first 32K page of the F range is useable after boot up has finished.
- **P5 Piped Address.** Default is Disabled
- **PCI Arbiter Mode.** Devices gain access to the PCI bus through arbitration. There are two modes, 1 (the default) and 2. The idea is to minimize the time it takes to gain control of the bus and move data. Generally, Mode 1 should be sufficient, but try mode 2 if you get problems.
- **Stop CPU When Flush Assert.** See below.

- **Stop CPU when PCI Flush.** When enabled, the CPU will be stopped when the PCI bus is being flushed of data. Disabling (default) allows the CPU to continue processing, giving greater efficiency.
- **Stop CPU at PCI Master.** When enabled, the CPU will be stopped when the PCI bus master is operating on the bus. Disabling (default) allows the CPU to carry on, giving greater efficiency.
- **I/O Cycle Recovery.** When enabled, the PCI will be allowed a recovery period for back-to-back I/O, which slows back-to-back data transfers; it's like adding wait states, so disable (default) for best performance.
- **I/O Recovery Period.** Sets the length of time of the recovery cycle used above. The range is from 0-1.75 microseconds in 0.25 microsecond intervals.
- **Action When W_Buffer Full.** Sets the behaviour of the system when the write buffer is full. By default the system will immediately retry, rather than wait for it to be emptied.
- **Fast Back-to-Back.** When enabled, the PCI bus will interpret CPU read cycles as the PCI burst protocol, meaning that back-to-back sequential CPU memory read cycles addressed to the PCI will be translated into the fast PCI burst memory cycles. Default is enabled.
- **CPU Pipelined Function.** This allows the system controller to signal the CPU for a new memory address, even before all data transfers for the current cycle are complete, resulting in increased throughput. The default is Disabled, that is, pipelining off.
- **Primary Frame Buffer.** When enabled, this allows the system to use unreserved memory as a primary frame buffer. Unlike the VGA frame buffer, this would reduce overall available RAM for applications.
- **M1445RDYJ to CPURDYJ.** Whether the PCI Ready signal is to be synchronized by the CPU clock's ready signal or bypassed (default).
- **VESA Master Cycle ADSJ.** Allows you to increase the length of time the VESA Master has to decode bus commands. Choices are Normal (default) and Long.

- **LDEVJ Check Point Delay.** This allows you to select how much time is allocated for checking bus cycle commands. These commands must be decoded to determine whether a local bus device access signal (LDEVJ) is being sent, or an ISA device is being addressed. Increasing the delay increases stability, especially the VESA sub-system while very slightly degrading the performance of the ISA sub-system. Settings are in terms of the feedback clock rate (FBCLK2) used in the cache/memory control interface.

- 1 FBCLK2=One clock
- 2 FBCLK2=Two clocks (Default)
- 3 FBCLK2=Three clocks

- **CPU Dynamic-Fast-Cycle.** Gives you faster access to the ISA bus. When the CPU issues a bus cycle, the PCI bus examines the command to determine if a PCI agent claims it. If not, then an ISA bus cycle is initiated. The Dynamic-Fast-Access then allows for faster access to the ISA bus by decreasing the latency (or delay) between the original CPU command and the beginning of the ISA cycle.

- **CPU Memory sample point.** This allows you to select the cycle check point, which is where memory decoding and cache hit/miss checking takes place. Each selection indicates that the check takes place at the end of a CPU cycle, with one wait state indicating more time for checking to take place than zero wait states. A longer check time allows for greater stability at the expense of some speed.

- **LDEV# Check point.** The VESA local device (LDEV#) check point is where the VL-bus device decodes the bus commands and error checks, within the bus cycle itself.

- 0 Bus cycle point T1 (Default)
- 1 During the first T2
- 2 During second T2
- 3 During third T2

- **Local memory check point.** Allows you to select between two techniques for decoding and error checking local bus writes to DRAM during a memory cycle.

- Slow=Extra wait state; better checking (default).
- Fast=No extra wait state used.
- **FRAMEJ generation.** When the PCI-VL bus bridge is acting as a PCI Master and receiving data from the CPU, a fast CPU-to-PCI buffer will be enabled if this selection is also enabled. Using the buffer allows the CPU to complete a write even though the data has not been delivered to the PCI bus. This reduces the number of CPU cycles involved and speeds overall processing.
- Normal Buffering not employed (Default)
- Fast Buffer used for CPU-to-PCI writes.
- **PCI to CPU Write Pending.** Sets the behaviour of the system when the write buffer is full. By default, the system will immediately retry, but you can set it to wait for the buffer to be emptied before retrying.
- **Delay for SCSI/HDD (Secs).** The length of time in seconds the BIOS will wait for the SCSI hard disk to be ready for operation. If the hard drive is not ready, the PCI SCSI BIOS might not detect the hard drive correctly. The range is from 0-60 seconds.
- **Master IOCHRDY.** Enabled, allows the system to monitor for a VESA master request to generate an I/O channel ready (IOCHRDY) signal.
- **VGA Type.** This data is used when the video bios is being shadowed. The BIOS uses this information to determine which bus to use. Choices are Standard (default), PCI, ISA/VESA.
- **PCI Mstr Timing Mode.** This system supports two timing modes, 0 (default) and 1.

- **PCI Arbit. Rotate Priority.** Typically, the system manages or arbitrates access to the PCI bus on a first-come-first-served basis. When priority is rotated, once a device gains control of the bus it is assigned the lowest priority and every other device is moved up one in the priority queue.
- **I/O Cycle Post-Write.** When Enabled (default), data being written during an I/O cycle will be buffered for faster performance.
- **PCI Post-Write Fast.** As in the above I/O Cycle Post-Write, enabling this will allow the system to use a fast memory buffer for writes to the PCI bus.
- **CPU Mstr Post-WR Buffer.** When the CPU operates as a bus master for either memory access or I/O, this item controls its use of a high speed posted write buffer. Choices are NA, 1, 2 and 4 (default).
- **CPU Mstr Post-WR Burst Mode.** When the CPU operates as a bus master for either memory access or I/O, this item controls its ability to use a high speed burst mode for posted writes to a buffer.
- **CPU Mstr Fast Interface.** This enables/disables what is known as a fast back-to-back interface when the CPU operates as a bus master. When enabled, consecutive reads/writes are interpreted as the CPU high-performance burst mode.
- **PCI Mstr Post-WR Buffer.** When a PCI device operates as a bus master for either memory access or I/O, this item controls its use of a high speed posted write buffer. Choices are NA, 1, 2 and 4 (default).
- **PCI Mstr Burst Mode.** When a PCI device operates as a bus master for either memory access or I/O, this item controls its ability to use a high speed burst mode for posted writes to a buffer.
- **PCI Mstr Fast Interface.** This enables/disables what is known as a fast back-to-back interface when a PCI device operates as a bus master. When enabled, consecutive reads/writes are interpreted as the PCI high-performance burst mode.
- **CPU Mstr DEVSEL# Time-out.** When the CPU initiates a master cycle using an address (target) which has not been mapped to PCI/VESA or ISA space, the system will monitor the DEVSEL (device select) pin for a period

of time to see if any device claims the cycle. This item allows you to determine how long the system will wait before timing-out. Choices are 3 PCICLK, 4 PCICLK, 5 PCICLK and 6 PCICLK (default).

- **PCI Mstr DEVSEL# Time-out.** When a PCI device initiates a master cycle using an address (target) which has not been mapped to PCI/VESA or ISA space, the system will monitor the DEVSEL (device select) pin for a period of time to see if any device claims the cycle. This item allows you to determine how long the system will wait before timing-out. Choices are 3 PCICLK, 4 PCICLK (default), 5 PCICLK and 6 PCICLK.
- **IRQ Line.** If you have installed a device requiring an IRQ service into the given PCI slot, use this item to inform the PCI bus which IRQ it should initiate. Choices range from IRQ 3 through IRQ 15.
- **Fast Back-to-Back Cycle.** When enabled, the PCI bus will interpret CPU read or write cycles as PCI burst protocol, meaning that back-to-back sequential CPU memory read/write cycles addressed to the PCI will be translated into fast PCI burst memory cycles.
- **State Machines.** The chipset uses four state machines to manage specific CPU and/or PCI operations. Each can be thought of as a highly optimized process center designed to handle specific operations. Generally, each operation involves a master device and the bus it wishes to employ. The four state machines are:

CPU master to CPU bus (CC)

CPU master to PCI bus (CP)

PCI master to PCI bus (PP)

PCI master to CPU bus (PC)

Each have the following settings:

- **Address 0 WS.** This refers to the length of time the system will delay while the transaction address is decoded. Enabled=no delay.
- **Data Write 0 WS.** The length of time the system will delay while data is being written to the target address. When Enabled, there will be no delay.

- **Data Read 0 WS.** The length of time the system will delay while data is being read from the target address. When Enabled, there will be no delay.
- **On Board PCI/SCSI BIOS.** You would enable this if your system motherboard had a built-in SCSI controller attached to the PCI bus, and you wanted to boot from it.
- **PCI I/O Start Address.** I/O devices make themselves accessible by occupying an address space. This allows you to make additional room for older ISA devices by defining the I/O start address for the PCI devices.
- **Memory Start Address.** This is for devices with their own memory which use part of the CPU's memory address space, allowing you to determine the starting point in memory where PCI device memory will be mapped.
- **VGA 128k Range Attribute.** When enabled, this allows the chipset to apply features like CPU-TO-PCI Byte Merge, CPU-TO-PCI Prefetch to be applied to VGA memory range A0000H-BFFFFH.
 - Enabled=VGA receives CPU-TO-PCI functions.
 - Disabled=Retain standard VGA interface.
- **CPU-To-PCI Write Posting.** The Orion chipset maintains its own internal read and write buffers which are used to help compensate for the speed differences between the CPU and the PCI bus. When this is Enabled, writes from the CPU to the PCI bus will be buffered. When Disabled (default), the writes will not be buffered and the CPU will be forced to wait until the write is completed.
- **CPU Read Multiple Prefetch.** A prefetch occurs during a process (e.g. reading from the PCI or memory) when the chipset peeks at the next instruction and actually begins the next read. The Orion chipset has four read lines. A multiple prefetch means the chipset can initiate more than one

prefetch during a process. Default is Disabled.

- **CPU Line Read Multiple.** A line read means that the CPU is reading a full cache line. When a cache line is full it holds 32 bytes (eight DWORDS) of data. Because the line is full, the system knows exactly how much data it will be reading and doesn't need to wait for an end-of-data signal, freeing it to do other things. When this is enabled, the system is allowed to read more than one full cache line at a time. The default is disabled.
- **CPU Line Read Prefetch.** See above. When this is enabled, the system is allowed to prefetch the next read instruction and initiate the next process.
- **CPU Line Read.** This Enables or Disables (default) full CPU line reads.
- **CPU Burst Write Assembly.** The (Orion) chipset maintains four posted write buffers. When this is enabled, the chipset can assemble long PCI bursts from the data held in them. Default is Disabled.
- **VGA Performance Mode.** If enabled, the VGA memory range of A 0000-B 0000 will use a special set of performance features. This has little or no effect using video modes beyond the standard VGA most commonly used for Windows, OS/2, UNIX, etc, but this memory range is heavily used by games such as DOOM.
- **Snoop Ahead.** This is only applicable if the cache is enabled. When enabled, PCI bus masters can monitor the VGA palette registers for direct writes and translate them into PCI burst protocol for greater speed, to enhance the performance of multimedia video.
- **DMA Line Buffer Mode.** This allows DMA data to be stored in a buffer so as not to interrupt the PCI bus. When Standard is selected, the line buffer is in single transaction mode. Enhanced allows it to operate in 8-byte transaction mode.
- **Master Arbitration Protocol.** This is the method by which the PCI bus determines which bus master device gains access to it.
- **PCI Clock Frequency.** Allows you to set the clock rate for the PCI bus, which can operate between 0-33 Mhz. CPUCLK/3 means the PCI bus was operating at 11 Mhz ($33/3 = 11$).

CPUCLK/1.5 CPU speed / 1.5 (Default)

CPUCLK/3 CPU speed/3

14 Mhz 14 Mhz

CPUCLK/2 CPU speed/2

- **Max, Burstable Range.** Sets the size of the maximum range of contiguous memory which can be addressed by a burst from the PCI bus, a half or one K.
- **ISA Bus Clock Frequency.** Allows you to set the speed of the ISA bus in fractions of the PCI bus speed, so if the PCI bus is operating at its theoretical maximum, 33 Mhz, PCICLK/3 would yield an ISA speed of 11 Mhz.

7.159 Mhz (default)

PCICLK/4 A quarter speed of the PCI bus

PCICLK/3 One third speed of the PCI bus

- **8 Bit I/O Recovery Time.** The recovery time is the length of time, measured in CPU clocks, which the system will delay after the completion of an input/output request to the ISA bus, needed because the CPU is running faster than the bus, and needs to be slowed down. Clock cycles are added to a minimum delay (usually 5) between PCI-originated I/O cycles to the ISA bus. Choices are from 1 to 7 or 8 CPU clocks. 1 is the default.
- **16 Bit I/O Recovery Time.** As above, for 16 bit I/O. Choices are from 1 to 4 CPU clocks. 1 is the default.
- **I/O Recovery Time.** A programmed delay which allows the PCI bus to exchange data with the slower ISA bus without data errors. Settings are in fractions of the PCI BCL:

2 BCLK=Two BCLKS (default)

4 BCLK=Four BCLKS

8 BCLK=Eight BCLKS

12 BCLK=Twelve BCLKS

- **PCI Concurrency.** Enabled (default) means that more than one PCI device can be active at a time (Award). With Intel Chipsets, it allocates memory bus cycles to a PCI controller while an ISA operation, such as bus mastered

DMA, is taking place, which normally requires constant attention. This involves turning on additional read and write buffering in the chipset. The PCI bus can also obtain access cycles for small data transfers without the delays caused by renegotiating bus access for each part of the transfer, so is meant to improve performance and consistency.

- **PCI Streaming.** Data is typically moved to and from memory and between devices in discrete chunks of limited sizes, because the CPU is involved. On the PCI bus, data can be streamed, that is, much larger chunks can be moved without the CPU being bothered. Enable for best performance.
- **PCI Bursting.** Consecutive writes from CPU will be regarded as a PCI Burst cycle. Enable = best performance; some cards might not like it.
- **PCI (IDE) Bursting.** As above, but this one enables burst mode access to video memory over the PCI bus. The CPU provides the first address, and consecutive data is transferred at one word per clock. The device must support burst mode.
- **Burst Copy-Back Option.** When enabled, if a cache miss occurs, the chipset will initiate a second, burst cache line fill from main memory to the cache, the object being to maintain the status of the cache.
- **Preempt PCI Master Option.** When enabled, PCI bus operations can be preempted by certain system operations, such as DRAM refresh, etc. Otherwise, they can take place concurrently.
- **IBC DEVSEL# Decoding.** Allows you to set the type of decoding used by the ISA Bridge Controller (IBC) to determine which device to select. The longer the decoding cycle, the better chance the IBC has to correctly decode the commands. Choices are Fast, Medium and Slow (default).
- **Keyboard Controller Clock.** Sets the speed of the keyboard controller (PCICLK_I = PCI bus speed).

7.16 Mhz Default

PCICLK_I/2 1/2 PCICLK_I

PCICLK_I/3 1/3 PCICLK_I

PCICLK_I/4 1/4 PCICLK_I

- **CPU Pipeline Function.** This allows the system controller to signal the CPU for a new memory address even before all data transfers for the current cycle are complete, resulting in increased throughput. Enabled means that address pipelining is active.
- **PCI Dynamic Decoding.** When enabled, this setting allows the system to remember the PCI command which has just been requested. If subsequent commands fall within the same address space, the cycle will be automatically interpreted as a PCI command.
- **Master Retry Timer.** This sets how long the CPU master will attempt a PCI cycle before the cycle is unmasked (terminated). The choices are measured in PCICLKs which the PCI timer. Values are 10 (default), 18, 34 or 66 PCICLKs.
- **PCI Pre-Snoop.** Pre-snooping is a technique by which a PCI master can continue to burst to the local memory until a 4K page boundary is reached rather than just a line boundary.
- **CPU/PCI Write Phase.** Determines the turnaround between the address and data phases of the CPU master to PCI slave writes. Choices are 1 LCLK (default) or 0 LCLK.
- **PCI Preempt Timer.** This item sets the length of time before one PCI master preempts another when a service request has been pending.

Disabled No preemption (default).
 260 LCLKs Preempt after 260 LCLKs
 132 LCLKs Preempt after 132 LCLKs
 68 LCLKs Preempt after 68 LCLKs
 36 LCLKs Preempt after 36 LCLKs
 20 LCLKs Preempt after 20 LCLKs
 12 LCLKs Preempt after 12 LCLKs
 5 LCLKs Preempt after 5 LCLKs

- **CPU to PCI POST/BURST.** Data from the CPU to the PCI bus can be posted (buffered by the controller) and/or burst. This sets the methods.
- **POST/CON.BURST.** Posting and bursting supported (default).

- NONE/NONE. Neither supported.
- POST/NONE. Posting but not bursting supported.
- **PCI CLK.** Whether the PCI clock is tightly synchronized with the CPU clock, or is asynchronous. If your CPU, motherboard and PCI bus are running at multiple speeds of each other, e.g. Pentium 120, 60 MHz m/b and 30 MHz PCI bus, choose synchronise.
- **IRQ 15 Routing Selection.** MISA=Multiplexed ISA for asynchronously interrupting the CPU. IRQ 15 is usually used for Secondary IDE channels or CD-ROMs.
- **CPU cycle cache hit same point.** Working on this.
- **PCI cycle cache hit sam point.** As above.
- **Arbiter timer timeout (PCI CLK) 2 x 32.** Working on this.

Hard Disk Utility

Hard Disk Format

Will format your hard disk so it can receive new partitions.

IT WILL SMASH EVERYTHING ON YOUR HARD DISK!!! USE WITH CAUTION. A lot of inexperienced users have lost their sanity with this one. Several computer stores have made extra money with it! There's no need to do this unless you experience errors or if you want to change the interleave. **DON'T TOUCH THIS IF YOU HAVE AN IDE DRIVE.** It will perform a low level format and probably SCRAP your IDE hard drive. IDE is the standard drive type that nearly everyone has now. SCSI or ESDI drives shouldn't be low-level formatted. The new drives actually don't perform the low level format, but some old AT-Bus (IDE) drives you can scratch with this... This entry is only sensible for old MFM or RLL hard disks! Please refer to your hard disk manual to see how or if your hard disk can be low level formatted. Don't tell us we did not warn you.

Many manufacturers provide utilities to low level format their IDE drives (or any other types). Please refer to the comp.sys.ibm.pc.hardware.storage

FAQ for more technical information about this procedure. If normal (high level) hard disk formatting is required, you can use DOS FDISK to first erase and create partitions and then use FORMAT. It is also a good idea when your hard disk becomes inaccessible to see if it is just the system files that are corrupted. Most of the time, it is the case. SYS will do the job of replacing system files. Several packages (PC-Tools, Norton, etc.) provide utilities for repairing "damaged" HDD and FDD. Therefore, low level format is always of LAST RESORT when you encounter HDD problems.

Auto Detect Hard Disk

Handy when you "forgot" the specs of your hard drive. The BIOS will detect the number of cylinders, heads and sectors on your hard disk. In some BIOS versions, this option is in the main SETUP menu.

Auto Interleave

Determines the optimum interleave factor for older hard disks. Some controllers are faster than others, and you don't want the sectors laid out so reading consecutive sectors usually results in just missing the sector you wanted and having to wait a whole disk rotation for it to come around again. On modern ones, it's always 1:1 (and even if it wasn't, you cannot reformat anyway).

Interleaving is specified in a ratio, $n:1$, for small positive integers n . Basically, it means that the next sector on the track is located n positions after the current sector. The idea is that data on a hard drive might spin past the heads faster than the adapter can feed it to the host. If it takes you more than a certain amount of time to read a sector, by the time you're ready for the next sector, the heads will have passed it already. If this is the case, the interleave is said to be "too tight". The converse, where the CPU spends more time than necessary waiting for the next sector to spin under the heads, is too "loose" of an interleave. Clearly, it is better to have too loose an interleave than too tight, but the proper interleave is better still. Especially since any controller with read-ahead cacheing can pull the whole track into its buffer, no matter how slow the CPU is about fetching the data down.

The 1:1 interleave arranges the sectors on a track as follows:

0 1 2 3 4 5 6 7 8 9 a b c d e f g (17-sectors, using base 17 for convenience, this is clearly the in-order arrangement, one after another)

This is 2:1 interleaving:

0 9 1 a 2 b 3 c 4 d 5 e 6 f 7 g 8

The CPU has a whole sector's worth of time to get the a sector's data taken care of before the next sector arrives. It shows which logical sector goes in each physical sector.

Anyway, an $n:1$ interleave restricts the transfer rate to $1/n$ the speed of a 1:1 interleave (which is better than 1 revolution per sector if the interleave is too tight!). No modern PC should require interleaving. Only MFM and RLL (maybe also ESDI) and floppy drives which are capable of it (you could format a 1.44 meg floppy to 21 sectors/track, which would require a 2:1 interleave to not exceed the 500 mbps speed of the controller...but why?).

Media Analysis

Scan the hard disk for bad blocks. It is performing a LOW LEVEL FORMAT on the track where bad sector is encountered to mark that sector as a bad. It could cause damage on user data, even if scanning itself is non-destructive (also on MFM, RLL disks). Therefore, DON'T USE this option to on AT-Bus (IDE), SCSI or ESDI drives. These drives store the bad block data themselves, so you don't have to tell them or scan the media! Recommendation: use a media analysis program provided by an utility package or your hard drive manufacturer.

Power Management

This menu appears on computers having the "**Green PC**" specification, an initiative of the EPA (Environmental Protection Agency of the United States) with its Energy Star program. The main purpose is to minimize power usage when the system stays inactive for a while. The standard is still not yet achieved among manufacturers, so expect to see several variations. On most cases, the power management strategies are incremental, meaning that the longer a system stays inactive, the more parts will close down.

There exists three power management schemes: **APM** (Advanced Power Management) proposed by Intel and Microsoft. **ATA** (AT Attachment) for IDE drives. **DPMS** (Display Power Management Signaling) matches video monitors and video cards so they may simultaneously shut down.

◆ **Green Timer of Main Board:** Allows to setup the time after which the CPU of an idle system will shut down. Disabled or a time interval ranging from 1 to 15 minutes are the usual options. 5 to 10 minutes recommended.

◆ **Doze Timer:** Amount of time before the system will reduce 80% of its activity.

◆ **Standby Timer:** Amount of time before the system will reduce 92% of its activity.

◆ **Suspend Timer:** Amount of time after the system goes in the most inactive state possible, which is 99%. After this state, the system will require a warm up period so the CPU, hard disk and monitor may go online.

◆ **HDD Standby Timer:** Allows to setup the time after which the hard disk of an HDD idle system (no HDD access) will shut down. A terrific option if you have a somewhat noisy hard drive unit. The choice of a time interval depends on how hard disk intensive is your operating system. This may depends also on the amount of memory available. You should setup a longer time interval, like 10 minutes, if you only have 8MB of RAM and running OS/2 or Windows. For a plain/standard DOS environments, 2 to 5 minutes are recommended. If you have a comfortable 16 MB or more, the time lapse can be shorter. There are some reports that this option may cause problems with slave hard drives (AMI BIOS only?).

Frequently Asked Questions

How do I clear the BIOS memory?

Three alternatives are available depending of your type of motherboard:

- ◆ Enter BIOS Setup and change to settings to Power-On Defaults.
- ◆ Disconnect the battery.
- ◆ Insert appropriate jumper and wait until the BIOS memory is cleared (see mainboard documentation, the jumper is often located near the battery).


Sometimes this is possible with DIP switches on the motherboard. Sometimes (if not), you will have to remove the battery. And sometimes (if no DIP's and no removable battery, and not willing to desolder the battery), you can short the battery with a resistor to lower the current available for the CMOS.

This is only recommendable as a very last option. The NiCad cells often employed have a very low internal resistance, so the resistor will have to be of a very low value for the voltage to drop significantly. The corresponding current would be quite high, which is not very good for battery life. A better option would be to use a resistor to discharge the battery. Obviously, this only makes sense when you have a NiCad cell (which will be recharged every time you turn the computer on) as opposed to a lithium cell (which cannot be recharged). In the former case, a resistor of 39 Ohm will discharge the battery in under half an hour relatively safely.

Another good way to discharge the NiCad is to put a 6 volt lantern lamp across it, and let it discharge completely. Not only does it provide an effective load, it also gives a visual indication of the charge state. It's a good way to prevent "ghost memory" that's so common to NiCads. Metal Nickel Hydride batteries are now being seen in some systems. They don't have this problem and they are more \$\$.

Can I upgrade my BIOS?

Most BIOSes are specifically designed for a motherboard and its chipset.



Therefore, on rare occasions you can upgrade your BIOS for a newer version. It is often less troublesome to buy a new motherboard that comes with its own BIOS and transfer your CPU (memory, cache memory and adaptor cards...) than start hunting around for a new BIOS chip. I know very few computer stores who sell BIOS chips separately. However, it is possible to upgrade your BIOS so it may support new hardware. By browsing in computer magazines (like Computer Shopper, PC Magazine, etc.) you will find adds on companies that specialize on that sort of thing. The information they need is the Serial Number for the BIOS chip. It is the **long** number that prints out when you boot up. It includes the BIOS date, the chipset, etc. The price tag can vary greatly (from \$10 to \$80), so are the BIOS upgrades offered. (Does anyone have supplementary information on this, like good-bad experiences with BIOS upgrade? I already had 2 feedback on this, and they all agree it is a little tricky, but it works).

The Bus Speed Guide



Introduction

The bus speed of modern computer systems is still a fairly neglected subject. If you want to know about the performance of a system, you always look first at the CPU speed. The higher the CPU speed, the faster the system -- that's how most computer users judge system performance

Two variables determine the speed at which your CPU runs: *the frontside bus speed* and *the clock multiplier*. By manipulating these variables you determine the CPU's clock speed.

Understanding Bus Speed : The CPU uses the FSB (frontside bus, also called the system or external bus) to communicate with system memory and peripherals.

FSB Speed x Clock Multiplier = CPU Speed

(This is what publicised speed of CPU e.g. Pentium III 450 MHz here FSB speed is 100 MHz and internal speed of CPU is 450 MHz using the clock multiplier of 4.5 i.e. $100 \times 4.5 = 450$ MHz).

The motherboard chipset controls the clock multiplier, which, in conjunction with the FSB speed, determines the core speed of the CPU. Multiply the FSB speed by the clock multiplier to get the CPU speed. By manipulating the clock multiplier and/or the FSB speed, you can increase the core speed at which the CPU runs.

The Importance of the Bus Speed

First of all, I think I should explain to all those who are unsure what 'Bus Speed' means and what it does.

The 'Bus Speed' is the frequency a Pentium, Pentium II, or compatible is tacted with externally. Only in the first days of the Pentium did the CPU run with the same tact/speed externally as internally. These were the days of the P60 and P66. A Pentium 60 was clocked at 60 MHz from the motherboard and its core was running at the very same speed. Times were changing and faster CPUs were needed, so Intel came out with the P90 and P100. These two fellows were still running at 60/66 MHz externally, but the core was multiplied by a factor of 1.5, which made a 90/100 MHz CPU. A short while later, Intel released the P75 to replace the P60 and P66. The P75 also runs with a multiplied speed: 50 MHz x 1.5 = 75 MHz. You know what has happened from there. Now the current CPU's, which runs at 100 MHz externally and 600 MHz (100x6) internally.

So what does this '*runs externally*' mean?

To start with, the external clock is obviously supplied by the motherboard. At this speed/frequency/tact the CPU is communicating with all of the system components that it's directly connected to. These are

- the Chipset
- the RAM
- the Second Level Cache

We know about the important performance components of a PC system?

1. the CPU - well, that's covered, it's specified by it's clock speed
2. the Second Level Cache - aha, determined by type, size and ... the bus speed
3. the RAM - aha again, determined by the type, the size, the timing and ... the bus speed
4. the Video Card - ugh, determined by the type and the PCI speed ... = bus speed/2
5. the Harddisk(s) - alright, determined by the type, the bus type and of

course the PCI Speed ... hence the bus speed

Well, it seems there's nothing in that box, which isn't influenced by the bus speed. That's why a Pentium II @ 350 MHz is much faster than Pentium II @ 333 MHz. None other cause but only of the external frequency difference. So the faster the external speed the least CPU has to wait for instructions to fetched in. Now in coming days after the arrival of althon from AMD the comparisons will not only be by the internal speed of the cpu but also how much the external speed the cpu supports.

The Chipset Guide

Overview

Every computer purchase, every motherboard upgrade, every CPU buying decisions comes back to the same thing: the system chipset. If the CPU is the brain of your PC, the chipset is its heart. It controls the flow of bits that travel between the CPU, system memory, and the motherboard bus. Efficient data transfers, fast expansion bus support, and advanced power management features are just a few of the things the system chipset is responsible for.

Slot 1 AGP Chipsets

Intel's Slot 1 CPUs are currently the fastest CPUs available for PCs and Slot 1 is a CPU interface patented by Intel. This is why there is no other CPU manufacturer producing Slot 1 CPUs currently. Slot 1 offers Intel's GTL+ bus protocol, known from the Pentium Pro CPU, which hosts the so called dual independent bus (DIB), which means that the Slot 1 CPUs have one bus which they use to communicate with the memory and system devices like the PCI and AGP bus and another bus for communication with the second level cache, which is built into the Pentium II CPU cartridges. GTL+ is particularly helpful for multi CPU systems and is a protocol that is patented by Intel as well. This means to other chipset manufacturers that they have to license the GTL+ protocol from Intel if they want to design their own Slot 1 chipsets.

Current Slot 1 CPUs are the Pentium II in the two flavors of 66 and 100 MHz front side bus (FSB) and the Celeron, which is nothing else than a Pentium II core without the second level cache and only designed for 66 MHz FSB.

Intel 440LX AGP Chipset - 66 MHz Front Side Bus

The Intel® 440LX AGPset introduced Accelerated Graphics Port capability to

the Intel architecture PC platform. This popular chip set provides a fundamental building block for balanced platform performance in a wide variety of Pentium® II processor-based desktop systems.

This chip set is a great solution for designers who wish to improve performance in volume desktop PCs, especially with data or memory-intensive applications such as multimedia graphics and video that require concurrent data transfers within the system.

The Intel 440LX AGPset also supports a manageable client or server in networked business computing environments to reduce the total cost of ownership.

The Intel 440LX AGPset introduced Intel's Quad Port Acceleration optimization for the Pentium(r) II processor. Enhanced features include distributed arbitration and support for up to four outstanding transactions from the CPU in the buffer.

Today's Intel® 440LX AGPset Highlights:

- High-volume, high-quality platform
 - Accelerated Graphics Port (AGP) improves overall system bandwidth for superior 3D graphics and video performance
 - SDRAM support enables faster memory reads and writes and supports faster data streaming between the Pentium II processor and graphics controller
 - Dynamic Power Management Architecture (DPMA) enables enhanced power management and fast resume from powered-down states
 - Ultra DMA/33 improves access to all IDE devices including hard disk drives and supports smoother audio/video performance
 - Quality and reliability of Intel's pre-silicon simulations and post-silicon validation processes
 - World-class manufacturing and support capabilities
-

Intel 440EX AGP Chipset - 66 MHz Front Side Bus

Designed with a feature set tuned for the **Basic PC market segment**, the Intel® 440EX AGPset balances the **Intel® Celeron™** processor system performance. The result: a new generation of cost-effective PCs with the added performance feature of the **Accelerated Graphics Port (AGP)**.

Based on the core technology of the **Intel® 440LX AGPset**, the Intel 440EX AGPset allows designers to easily implement support for popular multimedia computing graphics and applications on volume-priced Basic PC platforms. The Intel 440EX AGPset helps deliver the popular multimedia computing applications such as 3D graphics, DVD-ROM applications and video on the Basic PC platform. The chip set complements AGP even further by supporting 66-MHz SDRAM for improved memory performance.

Good Performance and Low Cost of Ownership

The Intel 440EX AGPset is also **ACPI-compliant**, providing the cost-of-ownership advantages of advanced power management to small business and home office PCs.

Support for Ultra DMA/33 allows designers to implement the latest high-speed disk drives. The 440EX AGPset also brings **Universal Serial Bus (USB)** capability to the Basic PC market segment.

Intel® 440EX AGPset Highlights:

- Cost optimization with tuned feature set for Basic PC market segment
- Optimizes the performance of Celeron processor-based Basic PCs
- Enables design of Net PCs with Celeron processor performance
- Brings AGP functionality and multimedia computing capability for today's graphics and video applications
- **ACPI-compliance** enables enhanced power manageability
- Micro ATX motherboard form factor for cost-effective production
- **Ultra DMA/33** supports high-speed disk drives
- Support for 66-MHz SDRAM enhances graphics and video performance
- **USB** enables instant connectivity to peripherals

- Quality and reliability of Intel's pre-silicon simulations and post-silicon validation processes
- World-class manufacturing and support capabilities
-

Intel 440BX AGP Chipset - 66/100 MHz Front Side Bus - NOW ONLY THESE ARE IN RUN

Intel's BX chipset is the first PC chipset that is offering a system or front side bus that runs at 100 MHz and therefore enables a higher memory bandwidth as the common 66 MHz FSB used to.

Intel® 440BX AGPset

The Intel® 440BX AGPset is the first chip set to deliver 100-MHz system level bandwidth to unleash the full performance of the **Intel® Pentium® II** processors at 350-MHz or 400-MHz. In addition, the Mobile Intel® 440BX chip set is the first chip set to bring Pentium II processor performance to mobile systems.

The Intel 440BX AGPset improves system performance by increasing the bandwidth of the system bus from 66MHz to 100MHz.

Highest System Performance through Technology Innovation

The chip set features Intel's Quad Port Acceleration (QPA) to improve bandwidth between the Pentium II processor, the Accelerated Graphics Port, 100-MHz SDRAM and the PCI bus. QPA combines enhanced bus arbitration, deeper buffers, open-page memory architecture and ECC memory control to improve system performance.

Ultimate Design Flexibility

The Intel 440BX AGPset provides developers with maximum design flexibility

by delivering one chip set for both 66MHz and 100MHz system designs. Now, one basic motherboard design supports all Pentium II processor speeds from 233MHz to 400MHz. This flexibility can save engineering resources while maximizing use of investments in BIOS, software development and compatibility testing.

Closing the Gap Between the Desktop and Mobile

The Intel 440BX chip set architecture supports Mobile Intel Pentium II processor performance—bringing desktop-like performance to mobile PCs. It also enables mobile systems to benefit from advanced power management and Wired for Management capabilities, allowing mobile platforms to function much like desktop systems when connected to a network.

Intel® 440BX AGPset Highlights:

- Optimizes Pentium II processor 350/400-MHz performance by enabling 100-MHz SDRAM and memory.
- Supports both 100-MHz or 66-MHz system and memory bus designs
- Quad Port Acceleration improves system performance
- Enables "best-of-class" 3D and multimedia and gaming platforms with the Intel740™ graphics accelerator-based graphics card
- ACPI-compliant power manageability
- Supports Wired for Management capabilities in desktop and mobile platforms
- The first chip set to support the Mobile Intel Pentium II processor
- Quality and reliability of Intel's pre-silicon simulations and post-silicon validation processes
- World-class manufacturing and support capabilities

The CPU Guide

AMD K5

The AMD K5 has been coming slowly, but now it's there and getting better and better. By now, you can get it from a PR 75 up to a PR 166. It's FPU is still slower than the FPU of the Pentium, but it's better than the FPU of the 6x86. It also will soon be released as a split voltage version, which will reduce heat problems of this CPU as well.

Here a table to show the weird CPU speed/bus speed/multiplier relations of the K5 CPUs:

Processor Performance	Clock Speed (MHz)	Bus Speed (MHz)	PCI Bus Speed (MHz)	Multiplier
K5 PR75	75	50	25	1.5
K5 PR 90	90	60	30	1.5
K5 PR 100	100	66	33	1.5
K5 PR 120	90	60	30	1.5
K5 PR133	100	66	33	1.5
K5 PR 166	116.66	66	33	1.75

As you can see, the PR120 runs at the same speed as the PR90, the PR100 runs at the same speed as the PR133. These CPUs obviously have got a different die. The PR120 and PR133 are more advanced than the PR90/100. This however only touches the integer performance. The FPU performance of the PR133 is the same as of the PR100.

Due to these strange multiplier settings, the K5 has got a fixed multiplier, which cannot be changed from out side via multiplier jumpers. The PR75 to PR133 will run at 1.5 multiplier regardless what the jumpers say, the PR166 will also always run at x1.75.

The IDT WinChip C6 CPU

When IDT announced the C6 CPU in the first half of 1998 it came as quite a surprise to most of us. After this first introduction it became pretty quiet around IDT again and only now this new Socket 7 CPU starts shipping. Targeted against the Intel Pentium MMX, IDT claimed that the C6 has got about the same integer performance at the same clock speed as the Pentium MMX. IDT was also not making any secret of the fact that the floating point as well as the MMX unit of the C6 will not be as fast as in the Intel Pentium MMX. The C6 is targeted to the lower end market, similar to the lower clocked versions of the AMD K6 and the IBM/Cyrix 6x86MX CPUs. The special trick of IDT is supposed to be the fact that the chip is fairly small as well as requiring only little power and IDT is therefore pretty confident of raising the clock speed to up to 400 MHz soon as well as using the C6 in notebooks where low power consumption is crucial.

The architecture of the C6 is much less sophisticated than the architecture of its competitors AMD K6 and IBM/Cyrix 6x86MX. No 'register renaming' or 'out of order execution' is used here for reaching Pentium MMX performance but simply a large L1 cache in cooperation with a pretty classic but straight forward microprocessor design.

At the Computex in June IDT only showed a few systems with C6 CPUs running at 150 or 180 MHz. I also had one 150 MHz engineering sample a few months ago which couldn't really impress me a lot. Now IDT shipped a whole system to me with a 200 MHz C6 so that I could test if this new Socket 7 CPU is living up to its claims. I compared the C6 200 to the Intel Pentium MMX 200, the AMD K6 166 and 200 and the IBM/Cyrix 6x86MX PR2 166 (running at 133 MHz) and PR2 200 (166 MHz). The price of the C6 200 lies pretty much somewhere in between these CPUs, with the Intel Pentium MMX as the most expensive, the IBM/Cyrix 6x86 PR2 166 as the cheapest.

The IDT WinChip C6 CPU is certainly no particularly great performer, but that was never claimed by IDT either. It is a 'low trouble' CPU with a decent integer performance under Windows 95 that might have some impact at the low end sector of IBM's and Cyrix's 6x86MX market. I personally like this CPU and would prefer it to the 6x86MX PR2 166, because it's so much easier to handle, however as long as the 6x86MX is cheaper, I wouldn't see a sensible reason for spending more money on the C6. Although my expectations were extremely low,

the C6 shows that it can beat at least the slowest 6x86MX or offer pretty much the same performance. The C6 will be sold for \$135 (200 MHz), \$90 (180 MHz). The \$135 for the 200 MHz version are pretty much the same what the PR2 200 rated K6 and 6x86MX CPU will cost from beginning of November. Both CPUs are faster than the C6 200, so that IDT will have to lower prices pretty soon if they want to sell the chip.

What I would like to see in the near future is a *real* low power version of the C6, which is running at 2.8V or even less. Then I hope that IDT will be able to crank up the clock rate yields, which could result in quite a market share, probably taking over customers of IBM and Cyrix. I think that the biggest future of the C6 will be in the notebook market in case they can reduce the voltage to about 2V. This and a clock rate of 400 MHz will even give Intel and its upcoming notebook version of the Pentium II quite a little bit to think about.

The road map doesn't look too bad regarding the above said. IDT is planning a 225 (3x75) MHz and a 240 (4x60) MHz version for this November and next year there will be a WinChip C6+ with improved FPU performance (as fast as Pentium MMX) and much improved MMX performance (faster than Pentium MMX) as well as extended x86 instructions that shall improve 3D performance up to threefold. These new instructions will be distributed within Microsoft's next Direct3D. The WinChip C6+ will use a 1/4 micron die and run at only 2.5 V.

For late '98 we can expect a 300 MHz C6+ as well as a new version with on-chip 256kB L2 cache.

The Cyrix 6x86 CPU

The Cyrix 6x86 or M1, as it was called before, is a strong contender to the Intel Pentium and is able to outclass the Pentium in quite a few respects. This CPU might also be able to push technology forward more than the current Pentium CPUs, because since the 6x86 P200+ was released, a new **bus speed of 75 MHz** turned up at the horizon to replace the old maximum of 66 MHz of Intel CPUs. This will require new memory technology as well as new chipsets and therefore new motherboards.

This however is also the downside of the new 6x86 P200+, because you can't run this chip in a Triton FX or HX board for they don't support a bus speed of 75 MHz. Only to remind you - the 6x86 P200+ runs at a clock speed of 150 MHz, which is a multiplying of the 75 MHz bus speed with 2. Would you run this CPU with a bus speed of only 50 MHz and a multiplier of 3, you would lose all the speed advantages due to this pathetic slow bus speed.

Technical Specifications of the 6x86

Processor Performance	Clock Speed (MHz)	Bus Speed (MHz)	PCI Bus Speed (MHz)	Multiplier	Special
6x86-P200+	150	75	37.5 or 33	2	needs special motherboard and chipset to support 75 MHz bus speed !
6x86-P166+	133	66	33	2	no
6x86-P150+	120	60	30	2	no
					needs motherboard

6x86-P133+	110	55	27.5	2	that supports 55MHz bus speed
6x86-P120+	100	50	25	2	no

- Clocking 2x, 3x bus-to-core clock multiplier, **no 1.5 or 2.5 clock multiplier!**
- L1 Cache 16-KByte; write-back; 4-way associative; unified instruction and data; dual-port address
- Bus 64-bit external data bus; 32-bit address bus
- Pin/Socket P54C socket compatible (296-pin PGA)
- Compatibility Fully compatible with x86 operating systems and software including Windows 95, Windows, Windows NT, OS/2, DOS, Solaris and UNIX
- Floating Point Unit 80-bit with 64-bit interface; parallel execution; uses x87 instruction set; IEEE-754 compatible
- Voltage 3.3V core with 5V I/O tolerance
- Power Management System Management Mode (SMM); hardware suspend; FPU auto-idle
- Multiprocessing Supports SLiC/MP(TM) and OpenPIC(TM) interrupt architecture
- Burst Order 1-plus-4 or linear burst

Problems with the 6x86 and the solution

Problem	Solution(s)
<p>The 6x86 needs lots of power, more than a Intel Pentium. This leads to two different problems.</p> <p>1. The voltage regulator has to supply much more current, which some regulators are not</p>	<p>There are two solutions from Cyrix to this problem</p> <p>1. 6x86 of the revision 2.7 or higher need considerably less current or power than the previous versions. If you should want to buy a single</p>

up to. The regulators get very hot.

2. The 6x86 CPU gets much hotter than a Pentium, which means that you have to take special precautions for good cooling of the CPU. I only can recommend the original Cyrix P200+ heat sink with fan, which is very good.

There are so far three different voltage specifications of the single voltage 6x86.

- Type C028, which requires 3.52 V
- Type C016, which requires 3.3 V
- 'voltage switching' types, which can work with any of the above voltages

It's a matter of fact, that at a lower voltage the chip needs less current and produces less heat.

voltage 6x86, you should at least get a revision 2.7 or higher.

2. The 6x86L won't have this problem at all, since it runs at 2.8 V and needs under some conditions even less current than a Pentium of the same P-rating. This is the best 6x86 you can get. I wouldn't buy any other 6x86 model.

The 6x86 is very sensitive to reflections on the different buses. This was leading Microsoft to switch off the write back L1 cache of the 6x86 completely in Windows NT 4.0. This results in a serious performance decrease under this OS for 6x86 users.

There are also several different solutions to this problem:

1. Cyrix offers a replacement to all registered NT4 users who are affected by this 25% performance decrease. Read more about this on Brian Davis' [excellent 6x86 site](#).
2. You also can download a [software patch](#) to enable the L1 cache again. However this *can* lead to instability

	<p>under NT4.</p> <p>3. The revisions 2.7 and higher are not affected by this NT4 problem. You'll best get one of these or even better a 6x86L in the first place.</p>
<p>The FPU (Floating Point Unit) of the 6x86 is considerably slower than the FPU of the Pentium. A 6x86 P166+ has a FPU performance of a Pentium 90 only. This is something most people noticed in the game Quake.</p>	<p>Unfortunately there is no solution to this. Please consider however, that the FPU is only needed in very few applications like CAD or other calculating applications. It really depends on how important Quake is to you. It is expected, that more future games will use the FPU. Let's hope that the M2 will have a better FPU than the 6x86.</p>



The Cyrix 6x86L CPU

The 6x86 CPU was not a child that came to the world and was making every person, that had to deal with it, happy straight away. The 6x86 had its problems and Cyrix had to fight with a lot of trouble since the 6x86 release. However this doesn't say that Cyrix wasn't working extremely hard to survive by fixing one problem after the other. The latest improvement the 6x86 has received from its creator is the most pleasing one and eventually now I am happy with the 6x86 as well. The letter 'L' behind the CPU name is making all the difference to me and hopefully to all the possible 6x86 buyers as well.

The 'L' stands for 'low voltage' and doesn't say anything else as that the 6x86 now also has got the 'split voltage' feature known from the newly released Pentium MMX. Now for the Pentium this isn't making much of a difference, since the Pentium never had to fight any serious current or heat problems. For the 6x86 however, this could be the final liberation from his bad name as 'chip fryer'.

The 6x86L has just as the Pentium MMX a 2.8 V core and a 3.3 V I/O interface voltage. Cyrix expects a power reduction of more than 25% of the previous models, which will keep the CPU much cooler and won't ask for more current than a Pentium MMX.

This is probably the solution for all the people who have been scared by the recent problems with the 6x86. Unfortunately it comes a little bit late, since everyone is waiting for the M2 and the Pentium MMX 200 is much faster than a 6x86 P200+ (which is even on a lower level than the Pentium MMX 166 in most applications). However, who thinks that a P200+ is still by far fast enough and who doesn't want to spend as much money as for Intel CPUs, should get a 6x86L and nothing else. Getting an old 6x86 without the 'L' now is considered to be completely crazy, unless the old single voltage 6x86 shouldn't be considerably cheaper than the 6x86L.

The Intel Pentium Classic

I don't think there's too much to say about the Pentium, because there's that much information already being released on the net and in magazines. However there's one important thing, which doesn't seem to be known too well. It's the fact that Pentiums of the same clock speed are not all the same, but can be VERY different, and you should know that before you get one ! There are now three reasons for that to be considered:

- **Second class Pentiums, which require a higher voltage or are not capable of multi processing**

- Intel is actually releasing different kinds of Pentiums. Some are not capable of multi processing, others have to have a higher supply voltage to run correctly. You have to realize, that Intel is not having a different production line for each clock speed, but is in real now producing more or less only one dye. After the chips have been produced, they are tested, and regarding how they performed in the test, the clock rates are chosen. Some chips would work correctly on a special clock rate only with a higher voltage, which means also getting hotter. Others wouldn't work with multi processing. To find out which kind of CPU you have or are offered, you have to have a look on the bottom side of the chip, where's written something like e.g. SK 106 SSS. These last three letters can tell you what kind of CPU you have, and here's **the list**. It's obvious, that you should try hard getting a chip, that runs at 3.3 V(Standard) rather than 3.4 to 3.6 V (VRE), and maybe it makes you feel better to know that you've got a *real* Pentium, that works with multi processing as well, though you won't be able to use multi processing on a Triton board. If you should be into overclocking, it's almost inevitable to get a SSS-chip, because than you still can increase the voltage on the motherboard to get it running at a higher clock speed (as long as you cool it well!).

The Intel Pentium MMX

The Technical Advantages of the Pentium MMX over the Pentium Classic

- doubling of the L1 caches, 16kB data and 16 kB instruction cache
- data cache now 4 way set associative - only 2 way set associative in Pentium classic
- doubling of the Write Buffers, now 4 instead of 2
- new branch prediction unit, taken from the Pentium Pro
- implementation of a Return Stack, as known from the Cyrix/IBM 6x86
- increase of the U and V Pipelines by one step
- improvement of the Parallel Processing Ability of the two pipelines
- the MMX instruction set with the ability to pair up two MMX instructions in one go, hence 2 SIMD instructions can be processed with 16 byte data in one clock cycle
- new split voltage technology, 2.8 V for CPU core, 3.3 V for CPU I/O interface

The Intel Pentium-Pro CPU

The family consists of processors at 150 Mhz and higher and is easily scalable to up to four microprocessors in a multiprocessor system and is suitable specially for servers.

Features

- Available at 150MHz, 166MHz, 180MHz and 200MHz core speeds
- **Optimized for 32-bit applications running on advanced 32-bit operating systems. (16 bit programs runs slower).**
- Single package includes Pentium® Pro processor CPU, cache and system bus interface
- Scalable up to four processors and 4 GB memory (suitable for servers)
- Separate dedicated external system bus, and dedicated internal full-speed cache bus
- 8K/8K separate data and instruction, non-blocking, level one cache
- Available with integrated 256 KB, 512 KB or 1 MB non-blocking, level two cache on package

Core Frequency	L2 Cache Size	L1 Instruction/Data Cache Size
150 MHz	256-Kbytes	8-Kbytes/8-Kbytes
166 MHz	512-Kbytes	8-Kbytes/8-Kbytes
180 MHz	256-Kbytes	8-Kbytes/8-Kbytes
200 MHz	512-Kbytes	8-Kbytes/8-Kbytes

The Intel Pentium II (P-II)

It is simply Pentium Pro + MMX. The Intel Pentium II processors deliver good performance and are fully compatible with existing software. The Intel Pentium II processor family includes 233MHz, 266MHz, 300MHz, 333MHz, (FSB Speed 66 MHz) and 350MHz, 400MHz, 450MHz (FSB Speed 100 MHz) versions. The processor core is packaged in the Single Edge Contact (SEC) cartridge.

Features of P-II:

- Available in speeds from 233MHz up to 450MHz.
- Includes MMX technology.
- At 450MHz, delivers a 25% performance boost over the 333MHz Pentium II processor
- Dual Independent Bus (DIB) architecture increases bandwidth and performance over single-bus processors.
- 350MHz, 400 MHz and 450MHz versions improve system bandwidth and performance by increasing the system bus speed from 66MHz to 100MHz. The Intel 440BX AGP set enables the 100MHz system bus to increase peak processor data transfers to the rest of the system by 50%.
- 32K (16K/16K) non-blocking, level-one cache provides fast access to heavily used data.
- 512K unified, non-blocking, level-two cache.

The Intel Celeron Processor

The Intel Celeron processor, is for the Basic PC (sub-\$1000) market segment. The Intel Celeron processor meets the core needs and affordability requirements common to many new users. Intel Celeron processor benefits from the same P6 microarchitecture core as the Pentium II processor. It is same as P-II minus the L2 cache. Intel was forced to release this as a cost effective solution because of the competition it has been facing in CPU market. Currently Celeron processor is available with 128 KB L2 cache.

- Provides a base level of functionality to meet the core needs and affordability requirements common to new home and business PC users.
- Based on the Intel P6 microarchitecture, the same microarchitecture the Pentium II processor is based on.
- Utilizes the Intel P6 microarchitecture's multi-transaction system bus at 66MHz.
- Includes a 32K (16K/16K) non-blocking, level-one cache that provides fast access to heavily used data.
- Available in the single edge processor package (S.E.P.P.) form factor, which maintains compatibility with Slot 1.

AMD ALTHON (CURRENTLY FASTEST CPU ON EARTH)

Back in end 1998 at the Microprocessor Forum in San Jose, California, the PC-world watched and listened in amazement to Dirk Meyer's first presentation of K7's or now Athlon's architecture. It was quite obvious to experts as well to most other listeners, including Intel employees, that this new AMD processor would mark a new era in the processor world, if AMD could make its promises come true. Now finally, the waiting is over and we can look at a new processor that is indeed living up to all the positive expectations that arose at the end of last year.

Athlon was designed with the future in mind. The architecture is designed for very high clock rates and obviously superior to Intel's P6-architecture, including the Pentium III. If AMD can deliver Athlon and move over to .18 μ -process soon, Athlon will be a continuous threat for Intel, even beyond Coppermine. And currently Althon beats Intel flagship chip Pentium-III in almost all spheres.

The Video Card and 3D Accelerator Guide

What does a Video Card do, what affects the performance?

We have to realize that the data as soon as it leaves the CPU has to go through 4 steps until it finally reaches the monitor:

- From the bus into the Video Chipset, where it's processed (digital data)
- From the Video Chipset into the Video Memory, to store a mirror of the screen picture here (digital data)
- From the Video Memory into the Digital Analog Converter (= RAM DAC), to read out the screen mirror and convert it for the monitor (digital data)
- From the Digital Analog Converter to the Monitor (analog data)

As you can see, except the step from the RAM DAC to the monitor, each step is some kind of a bottleneck and crucial for the overall performance of the graphical subsystem. The slowest step is the one which determines the overall speed. Lets now discuss, what these single steps mean and what actually happens:

The transfer of data between CPU and the Video Chipset

This bottleneck is mainly depending on the bus type and speed, the mainboard and its chipset. The fastest bus system at present is the PCI bus, so you will have slower performance with VL bus, ISA, EISA and NuBus (only for Macs). The PCI bus however doesn't always run at highspeed of 33 MHz, so with a Pentium 75, P90, P120, P150 you'll have a PCI bus speed of only 25 (P75) or 30 Mhz, which already here decreases the performance of the graphical subsystem. Later chipsets also offer faster PCI performance, so the Intel 430HX chipset offers a faster PCI performance than the Intel Triton 430FX chipset. As the name already says, it's not a bus, it's a port. This means you can only run one device on it, the

graphic device. It can run at 66 MHz and can transfer data at the rising and falling edge of a clock cycle (x2 mode). This makes it at least double as fast as PCI, but this does not necessarily result in double performance of AGP graphic cards, because the data transfer bandwidth is not the limiting factor of current graphic cards.

The transfer of data between Video chipset to Video RAM and from Video RAM to the RAM DAC

I have been taking these two steps together because here lies the key for the performance of a video card as long as you exclude special chipset features. The big problem of a video card is that the video memory lies in between two very busy devices and has to serve both of them all the time. Each time the screen has to change the chipsets has to alter the video memory (and it changes continuously, e.g. mouse pointer, cursor blinking, etc.). Also the RAM DAC has to read out the video memory continuously, to maintain the screen. You can see, the video memory is caught in between them and here all these smart ideas like using VRAM, WRAM, MDRAM, SGRAM, EDO RAM, or increasing the video bus size like 32 bit, 64 bit and now 128 bit come in.

The higher the screen resolution and the higher the colour resolution, the more data has to be transferred from the video chipset to the video memory and the faster the data has to be read by the RAM DAC to be sent to the monitor. You can see that the video memory has to be accessed all the time by the chipset **and** the RAM DAC. Normal dynamic RAM can only be accessed at a max. frequency, so after the video chipset finished accessing (r/w) the video memory, the RAM DAC has to wait until it's allowed to read and vice versa.

The Video Card Manufacturers found 3 different ways to fight that problem:

1. Here comes the idea in, to make the video RAM **dual ported**. This means, that the video chipset reads or writes from/to the video memory via one port, but the RAM DAC reads out the video memory through an independent second port. The video chipset doesn't have to wait for the RAM DAC anymore and the RAM DAC doesn't have to wait for the video chipset anymore. This kind of video memory is called **VRAM**. It's obviously more complicated by having double the ports and therefore more expensive to produce. That's the simple reason why VRAM cards are more expensive and also faster. The **WRAM** used by Matrox and a few other cards is also dual ported, but organised somehow smarter so that

it's faster than VRAM but also 20% cheaper to produce. If you should wonder why typically cards which offer a high refresh rate and high colour depth have these two kinds of memory, you should consider the following. A higher refresh rate means that the RAM DAC feeds the monitor with a complete screen picture more often than at a lower refresh rate. Therefore the RAM DAC has to read out the video memory more often. This only can be achieved with either VRAM/WRAM, by accessing the video memory via the second port, or by a considerable decrease of video performance of DRAM/EDO cards. If you don't believe it, just run your favourite video benchmark at a low and then at a high refresh rate - you'll see a considerable difference if you've got a DRAM/EDO card. The same is valid for a higher colour depth. At a 8 bit colour resolution (=256 colours) a 1024x768 screen needs 786,432 bytes to be read by the RAM DAC to send a complete screen picture to the monitor. At 24 bit colour resolution (16,777,216 colours) the same screen needs 2,359,296 bytes to be read by the RAM DAC - and this takes more time. This btw is also the reason why you often can't have the same high refresh rate at true colour as you had at low colour in cheaper cards.

2. The other way to fight this problem is to increase the **video memory bus size**. Years back everybody was amazed by the new 32 bit video cards. These cards had a 32 bit data path between video chipset, video memory and RAM DAC. With 32 bit data path you can transfer 4 bytes in one go. Later there came the 64 bit video cards = 8 byte in one go, which are the standard at present and only recently some new chipsets were born, to have a 128 bit data path = 16 byte in one go. It's easy to see, that video cards with both (VRAM/WRAM & wide data path) will be the best performers, but with a really wide data path you could get around VRAM/WRAM. **Now by getting completely excited about these wide data paths we shouldn't forget one very important thing: a normal 8x1Mbit memory chip, as used on most video cards has a data bus of 32 bits !!! Therefore even a 128 bit chipset can access this memory chip only 32 bit wide !!! This is the reason why all 64 bit video cards are a lot slower if only fitted with 1 MB of video RAM ! Don't get a 64 bit video card with less than 2 MB !!!!** Chipsets with 128 bit data path usually need at least 4MB local memory, otherwise their performance is cut in half. The NVidia Riva chipset e.g. is able to talk to only 2 MB as well, via a 64 bit data path. Riva cards with only 2 MB are therefore castrated. However, due to the architecture of the card you won't use 128 bit data path even if you upgrade to 4 MB, because the data path just stays the same. This is probably the case in many video cards, so be careful not getting a 1 MB 64 bit card or a 2 MB 128 bit card!

3. The third and to us maybe most obvious way to get the video RAM accessed faster is to simply increase the **clock speed of the video chipset/video RAM/RAM DAC**. Years back the video chipsets ran at clock speeds high above the mainboard memory bus speeds already. SGRAM is nowadays running at 100 MHz clock and some graphics chip manufacturers are already talking of 125 or even 133 MHz video RAM clock using 7 ns SGRAM. SGRAM is nothing but a special graphics version of SDRAM (synchronous DRAM), so we know this is able to run at clock speeds up to 133 MHz.


Summarizing all these performance aspects, we learn that for optimal performance we should have an AGP or at least PCI system with the latest chipset and 33 MHz PCI bus speed, a video card with a high performance chip and either SGRAM or WRAM, a wide data path or a high clock frequency of the video chipset or best all these three things together!

Now which parts of this video card/monitor-combo plays which role?





- The **Monitor** plays a crucial role in terms of sharpness, brightness, stability and max. screen resolution of the picture. If you want to have a high quality picture you're asking for a high quality monitor with a big screen, at least 17". Your video card can be as good as it wants, as long as the monitor is crap the screen will still look horrible.

On the video card side, the **RAM DAC** is the part that is responsible to send the data for a decent picture to the monitor. Two factors are important, the quality of the RAM DAC, e.g. is it stand alone or integrated into the video chipset, and the max. pixel frequency, measured in MHz. A 220 MHz RAM DAC is not necessarily but most likely better than a 135 MHz one and it certainly offers higher refresh rates - will tell you why further down on this page. RAM DACs tend to be included into the graphic chips more and more now, since it can decrease costs of graphic cards considerably and the quality of modern internal RAM DACs is coming close to the external ones.

- The **Amount of Video RAM** is responsible for the **colour resolution in combination with the screen resolution in 2D**, in 3D, which is getting more and more important, the amount of local card memory is also determining the **maximal 3D resolution**. 3D needs much more local memory than 2D for the same resolution. This is due to the fact that 3D



needs a front, a back and a Z-buffer. The front buffer holds what you see, the back buffer holds the next picture while it's being processed and the Z-buffer holds the 3rd dimension value (z-value, as x and y make two dimensions, z holds the third). That is the reason why a card with 4 MB local memory can offer a resolution of 1600x1200 at high color (16 bit) in 2D, because it needs $1600 \times 1200 \times 2$ byte = 3.7 MB. However games that are using z-buffer information (and the good ones do, offering you real 3D) can only run at 800x600 x 16 bit color x 16 bit z-buffer, $800 \times 600 \times 6$ byte (2 byte color front buffer, 2 byte color back buffer, 2 byte 16 bit z-buffer) = 2.74 MB. 3D at 1024x768 would require 4.5 MB and can't be displayed by a 4 MB 3D card.

- The **Type of Video RAM** in combination with the **Video Chipset** is responsible for all performance aspects of the video card/monitor-combo. However we shouldn't forget that the bus system (PCI/VL/ISA/EISA/MCA/NuBus) and therefore also the mainboard and the mainboard chipset are responsible for how fast the data reaches the video card. **AGP**, the advanced graphic port can offer much higher transfer bandwidth than PCI.
- 
- 
- 
- 

Basics and 2D Considerations

The video card, which is inevitable in each computer system, is responsible to process the special video data received from the CPU into a format that a monitor can understand to make a rastered picture on the screen out of it. More or less the monitor screen is still the output device of a computer system, it's the most important port through which we humans get data transferred from the computer. Therefore the Video Card/Monitor-Combo is one of the most important parts of our computer and we should take very good care of it.

Now what do we have to ask of this Video Card/Monitor Combo ?

Picture Quality is very important, so that it's not gonna be a pain for our eyes to sit and look at the screen. Here are the factors for picture quality:

- Sharpness
- Brightness
- Stability
- Colour Resolution
- Screen resolution

Speed of the video device. Due to the work we do with the computer, the picture changes eversooften and these changes should take place as fast as possible. It depends on what in particular we do on the computer, so there are also some catagories:

- **2D performance**, also called **GUI** or **Windows performance**, due to Windows being the most popular GUI OS. This used to be the most inportant performance so far, it determines how fast your office applications perform, e.g. how fast you can scroll text/graphics or how fast you can open and close new windows. Since the days of the Matrox Millennium 2D performance of graphic cards got pretty close to the limit and nowadays the latest graphic cards don't differ much in 2D

performance anymore, most of them are pretty fast, faster than the old standard of the Matrox Millennium.

- **3D performance** is the most important topic to distinguish between different graphic cards today. Cards without 3D acceleration will soon disappear from the market and Matrox had to learn this the hard way when releasing the Millennium II with hardly any 3D features. They lost their market leader position in an instance. S3 used to be a big player in the graphic chip market, but the mediocre 3D performance of their chip got them almost completely out of business. Similar things seem to happen to Cirrus Logic and others.
- **DOS performance**, which nowadays is to be equalized with **game performance for DOS based games** for all professional applications today are running under a graphical user interface operating system. DOS based games are disappearing as well, so that this performance is getting less and less important.
- **Video Display performance** in my eyes is still not that important to most of us, but whoever likes to watch and process videos on his computer will have to look for a fast video processing card. DVD will probably bring a significant change here.

Considerations for Gamers

In case gaming is most that you do on your system and you couldn't care less for Windows NT, true color and OpenGL, you want to go for a pure 3D gaming card or get an add on card.

- **Direct3D or Proprietary 3D Engine?**

You'll now have to decide what kind of games are important to you. Currently the graphically best games are often designed for a special graphic chip, or at least they look best with this one chip. The number one supported 3D graphic chip is nowadays the 3Dfx Voodoo, found on add on cards like the Orchid Righteous3D, Diamond's Monster3D and several

others. It looks as if upcoming games will still support this particular chip and since the Voodoo 2 is already on the horizon, you can expect 3Dfx's 'Glide' engine staying supported by many games for a long time.

Alternatively to a special 3D chip support, many new games are using Direct3D's new features quite heavily, so that it depends on how well the 3D card's driver translates Direct3D to their proprietary engine.

PowerVR's PCX1 and PCX2 chips are quite powerful 3D chips, but the cards that use them are highly incompatible. I've seen only very few games that run on this chip properly. If the PCX engine is used directly, the games look awesome though. The only 3D chip to my knowledge, that doesn't have a dedicated 3D engine, but is using Direct3D as its API directly, is NVidia's RIVA 128 chip, currently the fastest Direct3D chip available on the market. The RIVA 128 is wonderful for Direct3D games, but games that are only supporting a bunch of proprietary 3D engines will not run on the RIVA 128. The future will bring almost any game in Direct3D, which will help NVidia's RIVA a lot.

■ 3D Performance

It is not easy to measure pure 3D performance, because there are so many different ways a 3D engine can be used. Most official benchmarks are using the Direct3D engine of DirectX, like e.g. 3D Winbench or VNU's Final Reality. These benchmarks can only show you the card's Direct3D performance, hence how well the driver translates Direct3D into the chip's own 3D engine. NVidia's RIVA 128 doesn't need this 'translator', it uses Direct3D as its own API. This is only one reason why the RIVA scores by far best in Direct3D benchmarks. However some games written for that specific 3D engine of a chip can run much faster than the 3D Winbench score would let you expect them to. VQuake for Rendition's Verite 1000 is one good old example. The Verite 1000 was never scoring well in 3D Winbench, but VQuake looked good and ran fast.

■ 3D Quality!

Now 3D performance is only one thing, 3D quality is another. There are a lot of 3D features used nowadays, most of them supported and used by DirectX 5, but there will be even more 3D features implemented in DirectX 6. A 3D chip can only support a special amount of 3D features,

others are either not supported at all, or special drivers are used that emulate these features. In my latest test I came across only one chip that supports virtually every current 3D feature properly and this is 3Dfx's Voodoo chip. The big let down of the Voodoo chip leads to the other aspect for quality, the 3D screen resolution. The Voodoo chip can only do 640x480 in case of 2 MB frame buffer memory (4 MB cards), as in most of the Voodoo cards, or maximal 800x600 in case the card comes with 6 MB RAM (e.g. **Quantum 3D Obsidian 100SB**), 4MB hereof as frame buffer. NVidia's RIVA 128 chip has got a similar problem, it can't support more than 4 MB onboard memory, only good for a 3D resolution of maximal 800x600. Now it doesn't have to be that bad, since we are quite pleased with our good old television as well, which has a lower resolution than 800x600. The 3D chip and the system CPU have to be powerful enough for running smoothly at this resolution as well. However, I've seen 'Forsaken' at 1024x768 on a PII 300 with an ATI XPERT card and it looks pretty awesome.

■ **How Powerful is Your CPU?**

Some 3D chips are taking a lot of workload from the CPU, others want decent CPU performance for its operation. PowerVR's PCX chips want at least a Pentium MMX 166 for decent quality, 3Dfx's Voodoo lets games run fast even on systems with weak CPUs and Rendition's new Verite 2100/2200 chip gives a huge improvement to slow CPUs, but fast CPUs are reaching its limit and don't really benefit of this chip anymore. NVidia's RIVA chip seems to scale linearly from 6x86 CPUs up to Pentium II CPUs. Under Direct3D its always the fastest chip.

■ **Price!**

Another thing you obviously want to take in consideration is the price you've got to pay for the card. Many cards that have good 2D performance as well are pretty expensive. This is often due to the more expensive memory they are using, but it could also be the additional features like e.g. TV out or video compression. Cards with more memory are also more expensive, but they offer higher resolutions in 3D, higher color depth and higher resolutions in 2D. Make sure you don't pay for something you won't need.

Considerations for Professionals

■ Picture Quality

If you are working on your computer professionally one of the most important things is the picture quality. This is achieved by a high quality and high clocked RAMDAC. Most of the new graphic chips have included the RAMDAC internally, thus saving cost, but the best picture quality is still produced by an external RAMDAC. The most popular cards with external RAMDACs are Matrox Millennium I and II and Number Nine's Revolution 3D. These cards are still offering you the sharpest and cleanest picture on the screen. If you have got an expensive monitor, you want to use the high refresh rates your monitor supports. As a simple rule you should at least have a refresh rate of 85 Hz available for all the resolutions you want to use. Refresh rates of 120 and more sound nice, but they won't give you much of an advantage anymore. Responsible for this is again the RAMDAC. The higher its clock rate, the higher are possible refresh rates.

■ 2D Performance

The 2D performance was what used to determine the quality of a graphic card in the past. Now 2D acceleration seems pretty close to the limit and almost all cards are offering a good 2D performance, at least at 16 bit color mode. Good 2D performance in true color is a virtue that's pretty rare still though. Matrox and Number Nine always used to fight about the 2D performance crown and it hasn't changed much still. If you are working really professionally at your computer, you can't possibly use Microsoft's mouse driver collection called Windows 95. Hence you are either using Windows NT or some really good OS that's not from the monopolist. NT drivers are very important for professional cards and the NT performance should be more important than the Windows 95 performance. There is often quite a bit of a difference between NT and 95 performance.

■ OpenGL, Heidi

For people that use a real graphic workstation with CAD and/or 3D rendering, SGI's OpenGL as well as Heidi are of major importance. Nowadays if you hear 'PC' and 'OpenGL' one company comes into your minds ... 3DLabs. I will not discuss the real high end chips of 3DLabs, since this is off topic on this website, but 3DLabs' new Permedia 2 chip is one of the most impressive graphic chips on the market today in my opinion. For 3DLabs the Permedia 2 is nothing but a low end chip for the mass market, but amongst its competitors it's quite a gem in terms of professional work. In mid to higher prized systems cards with the Permedia 2 will offer you the best OpenGL performance combined with a good 2D and a fairly impressive Direct3D performance.

3D Accelerators

The introduction of real 3D-accelerators started with the 3Dfx Voodoo chip, and that was the time when the first hard core 3D gamers started to evolve, producing high expectations as well as crazy hypes and some pathological fanaticism too. Voodoo (1) and 3Dfx remained at the top of the 3D graphics scene for more than a year, completely underestimated by the big players in the graphics scene at this time, S3 and Matrox. 3Dfx won the second round also, the Voodoo2 was a worthy successor of Voodoo (1). The idea of Voodoo2 was the same as for Voodoo (1), offer an add-on 3D-card to the normal 2D card, use a pass-through cable and offer a lot more power than Voodoo1. Voodoo2 was using two parallel texture units, a lot more graphics memory and the user had the chance to run two of those cards in parallel, called 'SLI'-mode, offering at that time mind blowing 3D performance. This was certainly no cheap solution, but the hard core gamers jumped on this bandwagon right away, making Voodoo2 another huge success. It took six more months until the first decent 2D/3D-cards became available, NVIDIA shipped RIVA TNT and 3Dfx tried their slightly castrated version called 'Voodoo Banshee'. People who believed in the marketing hype, saying that those two chips were a really new generation of 3D chips, were disappointed, since Voodoo2 SLI was still offering the best performance, Banshee and TNT had only the chance to come close in some areas.

Voodoo3, TNT2, Permedia3 and what not are the first 3D-chips that show a real further development of the technologies implemented into Voodoo2. What took no less than six chips in the days of Voodoo2 SLI is now condensed into a TNT2-chip, beefed up with a faster RAMDAC, good AGP-support, support of larger texture-sizes and several other goodies, running at a much higher clock speed than Voodoo2 as well. Really new concepts are still not to be found though and can also not be expected before the second half of 1999. We are still working with the second generation of 3D-chips, but even this generation has still got a lot to offer and you can be sure that it will be a hot summer in the 3D-arena this year.

We can expect the first announcements of third-generation 3D-chips within the next couple of months and there's supposed to be a new quantum leap in 3D-performance and quality. Fill rates of 500-600 Mpixel/s, triangle rates of 20 million/s and more as well as the support of new quality enhancements will ask a

lot more from game developers, CPU, memory, RAMDAC and even display technologies than we can imagine right now.

The Second Generation of 3D Chips

Let's focus on those chips, that provide today's top performance, TNT2, Voodoo3, Rage128, G400, PowerVR SG, Permedia3, etc. Voodoo3 was the first of those chips and the expectations were high when Voodoo3 was announced for the first time, since 3Dfx had the history of supplying two top performers in the past, the Voodoo in 1997 and the Voodoo2 in 1998. At Comdex 1998 the journalist and card-makers weren't too impressed with the Voodoo3-features though. Even the following marketing campaign wasn't able to convince more than a few fanatic 3Dfx followers, when 3Dfx threw in all of their arrogance and ignorance towards the criticism of 3D analysts and journalists about missing features of Voodoo3. What is sold as 'Voodoo3' today should rather carry the name 'Banshee2', since it's a straightforward development of Banshee with performance measures that are close to Voodoo2. We certainly don't want to forget improvements as digital flat panel support, higher resolutions and better RAMDAC and the price/performance ratio of Voodoo3 is also way better than Voodoo2 SLI. However, the absolute performance as well as the long list of missing features has disappointed a lot of hard core gamers.

3D Card Requirements for Power Gamers

OpenGL (ICD) and DirectX (6.x) Support

The support of those two is of the highest importance for optimal 3D-game support. The majority of 1st person 3D-action shooters is based on OpenGL, Quake2, Quake Arena, Unreal, Unreal Tournament, Half-life and many more are the typical examples. The majority of all other games is covered by DirectX, particularly flight simulators and racing games are mainly based on this 3D-interface from Microsoft. A 3D-card that doesn't support those two 3D-platforms is pretty much out of the question for any gamer.

Fill Rates in excess of 300 Mpixel/s or Mtexel/s

The fill rates quoted by the chip vendors should be taken with a decent grain of salt. Firstly those data are 'synthetic', comparable to MIPS or MFLOPS of CPUs. The fill rate can be calculated from the 3D-chip clock, the number of independent pipelines of the chip, the graphics memory bandwidth and of course the hardwired features of the chip. The color depth, Z-buffer and particularly rendering quality varies significantly in between the vendors. A good example for this is TNT2, which has a lower theoretical fill rate than Voodoo3, but it scores higher frame rates in complex 3D scenes. Still the fill rate can give you some kind of idea of what kind of 3D-performance you can expect from a 3D-chip.

Up to 32 MB of Onboard Memory at Clock Speeds of Way over 100 MHz

The onboard memory of each 3D-card has a very high impact on the overall 3D-performance. 16 MB is enough to run 3D applications at up to 1024x768 in 32 bit color depth (this is not valid for PVRSG, which needs less memory for that).

Realistic 3D scenes at 1280x1024x32 bit color require 32 MB and that's what is state-of-the-art today. The memory clock still varies a lot between the different chips and so does the width of the memory interface, where currently 128 bit is state-of-the-art, moving towards 'dual-256 bit' = 2 interfaces with 128 bit or real 256 bit soon. The graphics memory bandwidth is an important limiting factor of the 3D-performance and you can say that a vendor who can safely handle highest memory clocks is well prepared for future developments.

32-Bit Rendering

More and more games are offering a true-color mode due to quality reasons, meaning that the displayed frames are using a 24 bit deep color depth (RGB, 16.7 million colors). This means that a huge amount of data needs to be computed for each frame, but the performance of the latest 3D-chips is starting to be capable of supplying this kind of computing power. A 3D-architecture that is limited to only 16 bit, as in case of the Voodoo3 requires a reduction of color information and is thus decreasing the quality and reality level. 32-bit color rendering with 32-bit Z-buffering will stay the optimal implementation for quite some time to come.

Fast RAMDACs with at least 250 MHz or Digital Flat Panel Support

The RAMDAC is the link between the digital display information stored in the frame buffer and the analog CRT. The higher the bandwidth of the RAMDAC, the higher is the number of pixels that can be displayed on the CRT each second. This means that a faster RAMDAC is crucial for decent refresh rates, its quality is crucial for the quality of the displayed picture on the screen and only fast RAMDACs can supply really high resolutions of more than 2000x1500, which is most important for 2D applications. If the user prefers a flat panel screen, he should make sure that there is no digital-analog-converter (DAC) and then a analog-digital-converter (ADC) between the graphics memory (or frame buffer) and the flat panel, because it is not only a waste of silicon, but it also decreases the display quality by a serious amount. A flat panel requires digital data and

should thus be supplied with nothing else than that. That requires a digital output of the graphics card, where DFP is one of the standards. For optimal 2D-quality you should therefore either look for a fast RAMDAC for a CRT or a digital output for a flat screen.

State-Of-The-Art AGP-Support

AGP vs. PCI is a question that doesn't need to be asked anymore today. AGP has succeeded in the upper performance area quite a while ago already. When AGP 4x will be released later on this year, we may finally see the visible advantage that many of us are still missing. A today's 3D-card should support AGP 2x to be at least able to handle the data bandwidth that's required for 'several million polygons/s'. Games with large textures are coming up, FlightSim 99 is only one example, so that the 3D-chip should also be able to do AGP-texturing, which is one of the things that Voodoo3 is not able to do.

So far about the requirements of a well performing 3D-chip or 3D-card. Of course there are a lot of other effects that 3D-chips can do today as well, like fog, transparency or different kind of shadings, but those 'common' features don't really need to be mentioned again. However, there are still cases where features are implemented at a low quality level, when 3D-chip makers rate frame rates over image quality.

Realism is what we really want, the users and game developers have to tell the 3D-chip makers what they expect. The hardware developers have made another leap forward in terms of realism with the latest generation of 3D-chips and 3D-games should follow soon to take advantage of it. A look at the CPU-roadmaps shows that we can expect a lot more performance in very short time.

The Harddisk Guide

The hard disk can have a huge impact on the performance of your PC: The fact is that the rotating magnetic media of the hard disk is one of the severest performance bottlenecks, causing second-long delays while fat programs spin off the disk and into RAM. Whereas disk access times are measured in milliseconds, system RAM performance is counted in nanoseconds. Understanding hard disk operation - and optimizing - can eliminate teeth-grinding delays.

The factors that affect the speed of a harddisk:

- Rotation speed
- Number of sectors per track
- Seek time / head switch time / cylinder switch time
- Rotational latency
- Data access time
- Cache on the HD
- How data is organized on the disks
- Transfer rates
- Interface (EIDE / SCSI)

What are sectors, tracks, heads and cylinders?

On a harddisk, data is stored in the magnetic coating of the disk. The so called head, held by an actor arm, is used to write and read data. This disk rotates with a constant turn time, measured in revolutions per minute (rpm). Data is organized on a disk in cylinders, tracks and sectors. Cylinders are concentric tracks on the surface of the disk. A track is divided into sectors. A harddisk has a head on each side of a disk. Nowadays, the actuator arm is moved by a servo-motor (not a step-motor which needs more time while swinging in after moving over the desired track). All harddisks have reserved sectors, which are used automatically by the drive logic if there is a defect in the media.

Rotation speed

Typical harddisks have a rotation speed from 4,500 to 7,200 rpm, a 10,000 rpm drive just hit the market. The faster the rotation, the higher the transfer rate, but also the louder and hotter the HD. You may need to cool a 7200 rpm disk with an extra fan, or its life would be much shorter. Modern HD's read all sectors of a track in one turn (Interleave 1:1). The rotation speed is constant.

Number of sectors per track

Modern harddisks use different track sizes. The outer parts of a disk have more space for sectors than the inner parts. Usually, HD's begin to write from the outside to the inside of a disk. Hence, data written or read at the beginning of a HD is accessed and transferred faster rate.

Seek time / head switch time / cylinder switch time

The fastest seek time occurs when moving from one track directly to the next. The slowest seek time is the so called full-stroke between the outer and inner tracks. Some harddisks (especially SCSI drives) don't execute the seek command correctly. These drives position the head somewhere close to the desired track or leave the head where it was. The seek time everyone is interested in is the average seek time, defined as the time it takes to position the drive's heads for a randomly located request. Yes, you are correct: seek time should be smaller if the disk is smaller (5 1/4", 3 1/2" etc.).

All heads of a harddisk are carried on one actuator arm, so all heads are on the same cylinder. Head switch time measures the average time the drive takes to switch between two of the heads when reading or writing data.

Cylinder switch time is the average time it takes to move the heads to the next track when reading or writing data.

All these times are measured in milliseconds (ms).

Rotational latency

After the head is positioned over the desired track, it has to wait for the right sector. This time is called rotational latency and is measured in ms. The faster the drives spins, the shorter the rotational latency time. The average time is the time the disk needs to turn half way around, usually about 4ms (7200rpm) to 6ms (5400rpm).

Data access time

Data access time is the combination of seek time, head switch time and rotational latency and is measured in ms.

As you now know, the seek time only tells you about how fast the head is positioned over a wanted cylinder. Until data is read or written you will have to add the head switch time for finding the track and also the rotational latency time for finding the wanted sector.

Cache

I guess you already know about cache. All modern HD's have their own cache varying in size and organization. The cache is normally used for writing and reading. On SCSI HD's you may have to enable write caching, because often it is disabled by default. This varies from drive to drive. You will have to check the cache status with a program like **ASPIID** from Seagate.

You may be surprized that it is not the cache size that is important, but the organization of the cache itself (write / read cache or look ahead cache).

With most EIDE drives, the PC's system memory is also used for storing the HD's firmware (e.g. software or "BIOS"). When the drive powers up, it reads the firmware from special sectors. By doing this, manufacturers save money by eliminating the need for ROM chips, but also give you the ability to easily update your drives "BIOS" if it is necessary (Like for the WD drives which had problems with some motherboard BIOS' resulting in head crashes!).

Organization of the data on the disks

You now know, a harddisk has cylinders, heads and sectors. If you look in your BIOS you will find these 3 values listed for each harddisk in your computer. You learned that a harddisk don't have a fixed sector size as they had in earlier days.

Today, these values are only used for compatibility with DOS, as they have nothing to do with the physical geometry of the drive. The harddisk calculates these values into a logical block address (LBA) and then this LBA value is converted into the real cylinder, head and sector values. Modern BIOS' are able to use LBA, so limitations like the 504 MB barrier are now gone.

Cylinder, heads and sectors are still used in DOS environments. SCSI drives have always used LBA to access data on the harddisk. Modern operating systems access data via LBA directly without using the BIOS.

Transfer rates

In the pictures you can see the several ways how data can be stored physically on the harddisk. With a benchmark program that calculates the transfer rate or seek time of the whole harddisk you can see if your drive is using a 'vertical' or a 'horizontal' mapping. Depending on what kind of read/write heads and servomotors (for positioning the actuator arm) are used it is faster to switch heads or to change tracks.

The Interface (EIDE / SCSI)

Currently there are 2 different common interfaces: EIDE and SCSI. You will find an EIDE controllers integrated with the motherboard and that EIDE harddisks are much cheaper than SCSI drives. For SCSI you need an extra controller, because there aren't a lot of motherboards with integrated SCSI controllers. Together with the higher price of a SCSI disk a SCSI system is more expensive than EIDE.

The EIDE interface has a primary and a secondary channel that will connect to

two devices each, for a total of four. That could be a harddisk, CD-ROM or disk changers. Lately there have been tape backups with EIDE connectors, but you need special backup software.

Scanners for example aren't available with EIDE interface, only with SCSI. You can connect up to 7 devices to a SCSI bus or 15 devices to a Wide SCSI. In a standard environment, the performance of single harddisk won't improve much from the SCSI interface. Rather, the power of SCSI is that several devices can use the bus at the same time, not using the bus while they don't need it. So, we see the best benefit from SCSI when several devices are all used on the same bus.

On one EIDE channel, the 2 devices have to take turns controlling the bus. If there is a harddisk and a CD-ROM on the same channel, the harddisk has to wait until a request to the CD-ROM has finished. Because CD-ROM's are relatively slow, there is a degradation of performance. That's why everybody tells you to connect the CD-ROM to the secondary channel and your harddisk to the primary. The primary and secondary channels work more or less independently of one another (it's a matter of the EIDE controller chip).

The SCSI interface comes in several types. 8-bit (50 wire data cable) or 16-bit (68 wire data cable, Wide SCSI). The clock can be 5 MHz (SCSI 1), 10 MHz (Fast SCSI), 20 MHz (Fast-20 or Ultra SCSI) or 40 MHz (Ultra-2 SCSI).

Possible transfer rates of the SCSI bus		
SCSI bus clock	8-bit (50 wire data cable)	16-bit (68 wire data cable, Wide SCSI)
5 MHz (SCSI 1)	5 MBytes/s	NA
10 MHz (Fast SCSI, SCSI II)	10 MByte/s	20 MByte/s
20 MHz (Fast-20, Ultra SCSI)	20 MByte/s	40 MByte/s
40 MHz (Fast-40, Ultra-2 SCSI)	40 MByte/s	80 MByte/s

The theoretical transfer rate of EIDE is up to 16.6 MByte/s in PIO mode 4 or **multi DMA mode 2** (soon **33.3 MByte/s**) with all the problems you may have already faced. Here you will find a **table of several interfaces** and their speeds. However, today's CD-ROM's often use PIO mode 3, while older device use PIO mode 0 to 2. Sometimes devices lie about the PIO mode they support. There are harddisks that say they are able to use PIO mode 2 but they only work reliably in PIO mode 1! Whenever you get errors accessing your harddisk, try to lower the PIO mode first!

Possible theoretcial transferrates of the IDE bus (ATA)	
single word DMA 0	2.1 MByte/s
PIO mode 0	3.3 Mbyte/s
single word DMA 1, multi word DMA 0	4.2 MByte/s
PIO mode 1	5.2 MByte/s
PIO mode 2, single word DMA 2	8.3 MByte/s
Possible theoretcial transferrates of the EIDE bus (ATA-2)	
PIO mode 3	11.1 MByte/s
multi word DMA 1	13.3 MByte/s
PIO mode 4, multi word DMA 2	16.6 MByte/s
Possible transferrates of Ultra-ATA (Ultra DMA/33)	
multi word DMA 3	33.3 MByte/s

It is not only the interface transfer rate that determines how fast a harddisk is. How fast the data can be written or read from the media, e.g. data density and rotation speed is more important. The fastest interface can't do anything faster than the 'inner values' of a harddisk are capable of. Today, most harddisks are still under 10 MByte/s transfer rate physically. A faster interface is advantageous on when data is read from or written to the cache in a multitasking environment with several devices accessed simultaneously.

Multitasking environments especially benefit from SCSI, since simultanous access occurs frequently. If you have a server or are working with large files like

audio, video or disk-intensive applications, you will benefit more from SCSI than EIDE. There are three reasons for this:

- All modern operating systems now supports SCSI very well. Windows 3.x didn't!
- Busmastering really works better with a SCSI busmastering controller.
- The fastest harddisks with the best performance are SCSI.

If you need large capacities and the highest transfer rates available on the market you need SCSI. This is not because EIDE is incapable of this, it's because of the market. High-end disks with high capacities and high performance are intended to be used in servers and aren't build with EIDE interface. At the moment, EIDE disks are only built with up to a 5 Gigabyte capacity (there is a problem with a 4 GB barrier with some BIOS's again and for drives bigger than 8 GB you need a new BIOS that supports the INT 13 functions AH=41h bis 49h) and transfer rates of about 9 MBytes/s. If you need more, you'll have to use SCSI. Also, SCSI harddisks have larger cache RAM than EIDE harddisks.

Performance, some thoughts

You need to know how a slow or fast harddisk affects your overall system performance in a standard enviroment. If your operating system isn't constantly swapping (e.g. you have enough memory) the speed of a harddisk is only a small part of a well balanced system. Let's say you have a harddisk that has 30% better performance than another older one; the benefit for standard applications would be from 2% up to 18%. Sometimes, you want or need the fastest components available. Other times, more capacity and reliability is needed.

There are several programs available that test the performance of a harddisk. Some are crap, others are good. In any case, if you have one, you get numbers that tell you something. But do you have a point of comparison? Different benchmarks mean different numbers. Different environments mean different numbers. Modern benchmarks are independent from existing data on the harddisk (only read performance testing can be done). But a benchmark could be affected by several things:

- To which channel is the harddisk connected

- Is the harddisk alone or together with other devices connected to the controller
- Under which operating system is the harddisk tested and used
- Which drivers are loaded or **not** loaded.
- Testing at Monday or Friday etc.

The File System

It is very important and I recommend you to go through my [FAT32 and Hard Disk Maintenance Guide](#)

The Bus Master DMA Feature of the Triton

The Intel PCIsets 430 FX, HX VX and TX include the PIIX (PCI ISA IDE XCELERATOR), which is capable of transferring data between an EIDE device and the memory via DMA (Direct Memory Access) by using Bus Mastering. This means that data is transferred (more or less) directly from the HDD to main memory without using very many CPU resources. Also the new Multi Word DMA Modes introduced by the EIDE/Fast ATA specifications are able to transfer data just as fast as the new PIO (Programmed Input/Output) Modes. The advantage of using the Bus Master DMA is that it uses less CPU resources than PIO and therefore it's of particular benefit in multitasking environments, where the CPU can work on a different program while data is transferred to or from the HDD.

However, the way of the data from the harddisk into the memory is divided into two parts: From the drive to the chipset's EIDE interface and from the EIDE interface into the memory. The BIOS of the PC sets the EIDE interface of the chipset to use busmaster transfer between the drive and the interface. The busmaster driver is needed, to initiate the DMA transfers between the interface and the memory. The busmaster driver can switch on the DMA transfer between the drive and the interface, but if this isn't properly supported by the BIOS it will use PIO mode instead. However, the transfer from the EIDE interface into the memory can still use DMA.

If you want to use this feature, you'll need a driver for your specific OS. There is no option in any BIOS so far to enable or disable it, so don't bother about your BIOS, you'll just need the right driver. Since Windows 95 seems to be the most popular OS so far, here the...

Windows 95 IDE Bus Master DMA Drivers

You can find any busmaster driver available at [BM Drivers.com](http://BMDrivers.com)

In the so called Windows95 service release 2 (which is only sold as an OEM-product to PC manufacturers) you now have a DMA capable driver included (name is still ESDI_506.pdr). DMA mode can be switched on or off in the drive property screen in system manager.

Windows98 supports Busmastering for most Chipsets

Problems with the Bus Master DMA Drivers under Windows 95

The most common problem with these drivers I came across so far are with an ATAPI CD-Rom drive or a NON EIDE HDD, which also is connected to the EIDE interface and isn't recognized by the driver. In other words: *You install the driver and your CD-Rom and older harddrive have disappeared.*

We also have all come across the **delayed Windows 95 boot up** problem, which seems to annoy a lot of impatient people. If you have a look into the System Manager you'll find a not working second IDE port in case you haven't connected anything to it. The solution to this problem is the same.

One very good answer to this problem is to get the DMA driver working on one EIDE port and the default PIO driver on the other (to connect the CD-Rom or old hard drives or to get rid of this non functioning second IDE port).

After installing the Bus Master Driver you simply have to change the registry (always back up registry before changing it !!!!!):

- find **HKEY_LOCAL_MACHINE/System/CurrentControlSet/control/Services/Class/hdc**
- there should be four subdirectories **0000-0003**

- find the one where **DriverDesc** is something like "**Primary Bus Master IDE controller**" or "**Secondary Bus Master IDE controller**", according to the port you want to change (should be 0002 or 0003).
- in this subdirectory you change **PortDriver** from "**ideatapi.mpd**" to "**ESDI_506.pdr**"
- if you want to, you can change **DriverDesc** to something like "**Standard IDE/ESDI controller**", to make it look more correctly in your Systems Manager
- reboot

Now this EIDE port is using the default PIO driver and you easily can use CD-Roms or non EIDE hard drives on this port.

A different solution to this problem seems to be keeping the data wires to the HDDs/CD-Roms as short as possible, which does the trick Mount the drives properly, they should be grounded !

Another trick that may shorten the startup time, is to start Windows95 in safe mode and delete all drives in System-Manager. Don't worry, with the next start up of Windows95 they will be recognized again.



Keyboards

The keyboard is one of the simplest parts of the computer to understand. This chapter goes beyond the keyboard basics, however, to show you how to do tasks such as inserting special characters or typing in another language. This chapter also looks at how to take care of your keyboard and how to take care of yourself when using it.

Understanding the Keyboard Layout

The basic function of almost every keyboard is compatible with what is called a *standard 101- key keyboard*. The standard keyboard is divided into four main groups of keys:

- *The typewriter or alphanumeric keys.* These are all of the standard letters, numbers, Tab, Shift, and the spacebar arranged in the usual QWERTY order. The Return key is replaced by an Enter key, and some special computer keys (Alt and Ctrl) are added at the bottom.
- *The function keys.* These are usually labeled F1 through F12, and have different functions in each program you use (some programs may not have any features that use them).
- *The cursor keys.* These are the arrows and other keys that move the cursor or insertion point on the screen. In addition to the up, down, left, right keys, some keyboards add diagonal keys. This block of keys also includes Insert, Delete, Home, End, Page Up, and Page Down keys.
- *The numeric keypad.* This has two functions. When the Num Lock light is on, this functions as a 10-key calculator. When the Num Lock light is off, this functions as another set of cursor movement keys.

Most keyboards carry a variety of additional keys, including Esc (Escape), Print Screen, Scroll Lock, and Pause. The uses of these specialty keys varies from system to system. In most systems:

- Esc can be used to undo commands or actions and in keyboard combinations.
- Print Screen can be used to copy the current screen contents to the Clipboard, from which you can then Paste the contents into a document for printing (see the Print Screen entry in "Troubleshooting Common Keyboard Problems" later in this chapter).

- Scroll Lock changes the action of the directional arrow cursors.

Check your user manual to find the details of the specialty keys on your system's keyboard.



Using Special Characters and Symbols

Many documents include characters that don't appear on the keyboard. These special characters include letters from the Greek alphabet, scientific values, symbols for foreign currency, and copyright and registration symbols. Some applications have shortcuts for entering these special characters; Windows 95 provides its own set of special characters in its Character Map accessory. To use the character map, follow these steps:

1. Check to see what font you are using in the document to which you want to add the characters.
2. Click the Windows 95 Start button and choose Programs, Accessories, Character Map.
3. In the Character Map application font list, select the name of the font you're using in your document.
4. Find the character(s) you want to insert in the program document and double-click them. This enters the characters in the Characters to Copy box.

You can select several characters for insertion in your document--all at the same time! The characters that you double-click are shown in the Characters to Copy box.

5. When you have the characters you want to use, click the Copy button.
6. Return to the program that you want to use these in and click the mouse where you want to insert them.
7. Open the Edit menu in that program and choose Paste. This pastes a copy of those characters in the document. You can then cut and copy the individual characters into their appropriate locations in your document.

Typing in Another Language

When you installed Windows, it came with a default language for the keyboard. If you bought your PC or your copy of Windows 95 in the U.S., that language is English (United States). If you need to type in another language, Windows allows you to select another language for the keyboard. To do this, follow these steps:

1. Click the Start menu; choose Settings, Control Panel.
2. Double-click the Keyboard icon in the Control Panel to open the Keyboard Properties dialog box.

NOTE: Only the CD-ROM version of Windows 95 includes multi-language support. If you installed from floppy disks, you'll need to download and install the multiple language support from Microsoft's Web site at <http://www.microsoft.com/windows/software/cdextras.htm> to add another language.

-
3. Click the Language tab so you see the dialog box.
 4. Click the Add button to open the Add Language dialog box.
 5. Select the new language from the drop-down list and click OK.

This dialog box shows what languages are installed and which one is the default.

Select the language from this list. Some languages (like English) have several varieties.

6. The language is now added. By default, Windows 95 will still use the original language. The dialog box indicates that if you want to switch between languages, you press the Left Alt key and Shift at the same time. By default, Windows will also show you an indicator of which language is being used in the taskbar. When you are done adding languages, click OK.

Hold the mouse pointer over the indicator to see a ToolTip pop up with the full name of the language. Click the indicator to produce a menu with which you can switch between languages.

You can now type in Windows in the new language. But, even with multiple language support, you need to keep in mind a few limitations when selecting a new language for your system:

- Some languages employ characters not used in English. The font that you use also has to have all of the available characters for that language.
- Your keyboard still has the U.S. English layout and keys. The foreign language feature is designed with the assumption that users are familiar with the keyboard for the language they've chosen, and the characters are matched to that placement on your current keyboard. If you're unfamiliar with the keyboard for the language you've chosen, you'll just have to type each key and find out which keys type what characters.
- Finally, not every world language is available. If the language you're seeking doesn't appear in the list, check with Microsoft Technical Support to see if they know a Windows 95-compatible source for the language.

Special Keys for Windows 95

Some keyboards have special keys that perform special tasks in Windows 95. The idea of these keyboards is to give you access to some common mouse functions on the keyboard so that you don't have to take your hands off of the keyboard to perform them. These Windows 95 keyboards usually have three special function keys bringing the total number of keys to 104. These are:

- The Windows logo keys (there are usually two of these, one each just to the outside of the Alt keys), both marked by the Windows logo
- The Menu key (usually located between the Windows logo key and the Ctrl key on the right side of the spacebar), marked by a stylized drop-down menu and an arrow pointer

The Windows logo key (referred to as WINDOWS in the following table) performs these actions:

Key Combination	Action
WINDOWS	Opens the Start menu.
WINDOWS+R	Opens the Run dialog box.
WINDOWS+M	Minimizes all windows.
Shift+WINDOWS+M	Undoes Minimize All.
WINDOWS+F1	Opens Help.
WINDOWS+E	Opens Windows Explorer.
WINDOWS+F	Opens the Find: All Files dialog box.
Ctrl+WINDOWS+F	Opens the Find Computer dialog box.
WINDOWS+Tab	Cycles through current taskbar buttons.
WINDOWS+Break	Opens System Properties dialog box.

The menu key has one simple application. Select any object and press the menu key and the same pop-up menu that appears when you right-click the mouse on that object appears.

Special Software for Keyboards

Several computer and keyboard makers have special programmable keyboards that you can make perform special tasks. For example, you can record a long set of keystrokes, and then play it back with just a few simple keystrokes (like speed dial on your phone).

If you have a programmable keyboard, you should also have received some keyboard control software and a manual explaining how to use the software. To master the use of a programmable keyboard, you'll need to read the instructions carefully. Even if you don't want to use it, it's a good idea to at least know how to take the keyboard out of "program" mode and put it back into "normal typing" mode, for those times when you accidentally trigger the program mode.

Keeping Your Keyboard Clean and Working

Never spill liquids on your keyboard. Coffee, soda, and other beverage spills can ruin your keyboard. Liquid spills on the keyboard have even been known to cause electrical damage to the PC itself. With that in mind, though you may not stop drinking coffee around your computer, you should at least get a spill-proof mug or keep the coffee on the other side of the desk.

Another enemy of keyboards is static electricity. Static electricity can have the same damaging effect on your keyboard as does liquid. If your keyboard doesn't respond properly after a strong static charge, you may just need to turn off the PC and turn it back on to reset the keyboard. In some cases, however, the static discharge can zap the keyboard and even parts of the PC. If you shuffle your feet across carpet or your PC is in a room with dry air, avoid touching the PC or the keyboard until you have touched something metal to discharge any static. If you don't have a metal desk or bookcase in your work area, consider buying an anti-static mat and keeping it where you can touch it before touching the PC.

Dust, dirt, food crumbs, and hair are other enemies of keyboards. Try to avoid eating over the keyboard and if your computer is in a dirty, dusty area, keep the keyboard covered when not in use.

Some dirt and dust is unavoidable. To keep the keyboard working well, you should occasionally clean it. Any time you clean the keyboard, turn off the PC first; then try any or all of these three techniques for cleaning the keyboard:

- Turn it upside down and shake it. This should shake loose some of the crumbs and dirt that collect between the keys.
- Get a can of compressed air (available for a few dollars at most PC or electronics stores) and use it to blow the dirt and dust from between the keys. Be sure to read the directions on the can before using it. If you hold the can wrong or use it incorrectly, you may blow cold liquid into the keyboard instead of air. Generally, you need to hold the can upright and tip the keyboard up at an angle so the debris will fall out when it is blown loose.

- If the keys themselves are dirty or sticky, wet a cleaning rag with some rubbing alcohol or a little bit of any non-abrasive household cleaner (something like Windex or Formula 409 is fine) and use it to clean the keys. Don't spray or pour any cleaner directly on the keyboard, and don't get the rag dripping wet.



Avoiding Keyboard-Related Injuries

Carpal tunnel syndrome, tendonitis, and other repetitive stress injuries have become serious health problems for many computer users. If you will use your computer a lot, you should know how to avoid these injuries and what the symptoms are.

First, if you feel any pain after typing or using the PC for extended periods, consult a physician immediately. Symptoms can include:

- Pain in the wrists
- Numbness or tingling of fingers

To avoid this type of injury, you need to use the proper posture and position of the body while typing. Always sit with your back straight. Your shoulders should be relaxed, and your upper arms should hang almost straight down.

At this point, adjust your chair and desk so that your forearms are parallel to the floor and your hands are just above the keyboard. Don't rest your wrists on the desk or the keyboard.

While typing, your knuckles, wrists, and forearms should form a straight line. While a lot of people have started using wrist pads to rest their wrists on while typing, you should know that these generally don't keep your hands and arms aligned correctly.

If you type for hours on end, you should stop and take a break for just a few seconds every few minutes. This is even more important for touch typists whose actions are probably more repetitive than those who "hunt and peck."

Beyond these issues, you may want to consider investing in a special ergonomic keyboard. The Microsoft Natural keyboard and several others like it are specially designed to reduce the risk of this type of injury. These keyboards are generally split down the middle and angled to allow your arms and wrists to rest in the proper alignment. If you think you want to try one of these keyboards, find a store with one set up in a display where you can sit down and try the keyboard. Many computer professionals who work constantly at their PCs won't use

anything but one of these special keyboards.



Special Notes on Notebook Keyboards

Notebook keyboards have special configurations designed to fit all the keys into a smaller-than-standard keyboard. For those accustomed to a standard-sized keyboard, a notebook's keyboard can take some getting used to. The keys themselves are usually smaller and spaced closer together. Further, notebook keyboards almost always have fewer keys.

Notebooks rarely have a built-in numeric keypad, and they don't have duplicate Alt and Ctrl keys (as do standard keyboards). Notebook keyboards do have special function keys which you use in combination with other keys to replicate the function of a numeric keypad, or to control color, contrast, or brightness. These keys usually have another set of labels printed on them in a color other than the alphabet color on the keys. Your notebook's manual or online help feature will explain the use of these and any other special function keys.

Using a Standard Keyboard with a Notebook

Most notebooks have a plug that allows you to plug in an external keyboard. If you have a notebook and you would really prefer to use a standard keyboard, check to see that you have a keyboard with the right size connector to fit your notebook. With the notebook off, plug in the keyboard and turn on the notebook. Most notebooks will automatically detect the external keyboard and know to use it and not the built-in keyboard.

CAUTION: Don't plug in or unplug the external keyboard while the notebook power is on.

If you plan to use the notebook a lot with the external keyboard, you may want to invest in a port replicator or docking station. The notebook snaps into place in one of these and the keyboard remains attached to the replicator or docking station. This saves you the trouble of plugging the keyboard into the notebook each time you want to use it. This is even more handy if you also use an external mouse and monitor with the notebook.

Troubleshooting Common Keyboard Problems

The following tips may help you diagnose (and perhaps resolve) some of the most common keyboard problems.

Keys Type Odd Symbols or Letters

To solve this dilemma, do any of the following:

- See the section "Using Special Characters and Symbols" earlier in this chapter, and be sure that the language selected in Windows 95 is English (United States) or whatever language you want the keyboard to be typing in.
- Check the font selection in your application to be sure that you haven't selected a symbol font like WingDings or Symbol. Try selecting Times or Arial and see if the problem is fixed.
- If the problem happens only occasionally and only with a few keys, the keyboard itself may be bad. Borrow a keyboard from a friend or buy a cheap spare at the store. Plug it in and see if it works. If you get the same problem, your problem may be in the PC itself, not the keyboard, and the PC will require professional service.

Print Screen Key Doesn't Work

In Windows 95, the Print Screen key doesn't print to the printer. What it does do is copy the screen to the Clipboard. To print a screen, press the Print Screen key to copy what you want to print. Next, click the Start menu and choose Programs, Accessories, Paint. In Paint, click the Edit menu and choose Paste. You now have a Paint picture of your screen. Click the File menu and choose Print. You can save the Paint picture if you want to save the picture of your screen.

Arrow Keys Type Numbers Instead of Moving the Cursor/Number Keys Move the Cursor

Look for the light that indicates the Num Lock on your keyboard. When this is lit, the numeric keyboard is in number mode and is used like a 10-key calculator. When the light is off, the numeric keypad is used for the arrow key movements indicated on the keys.

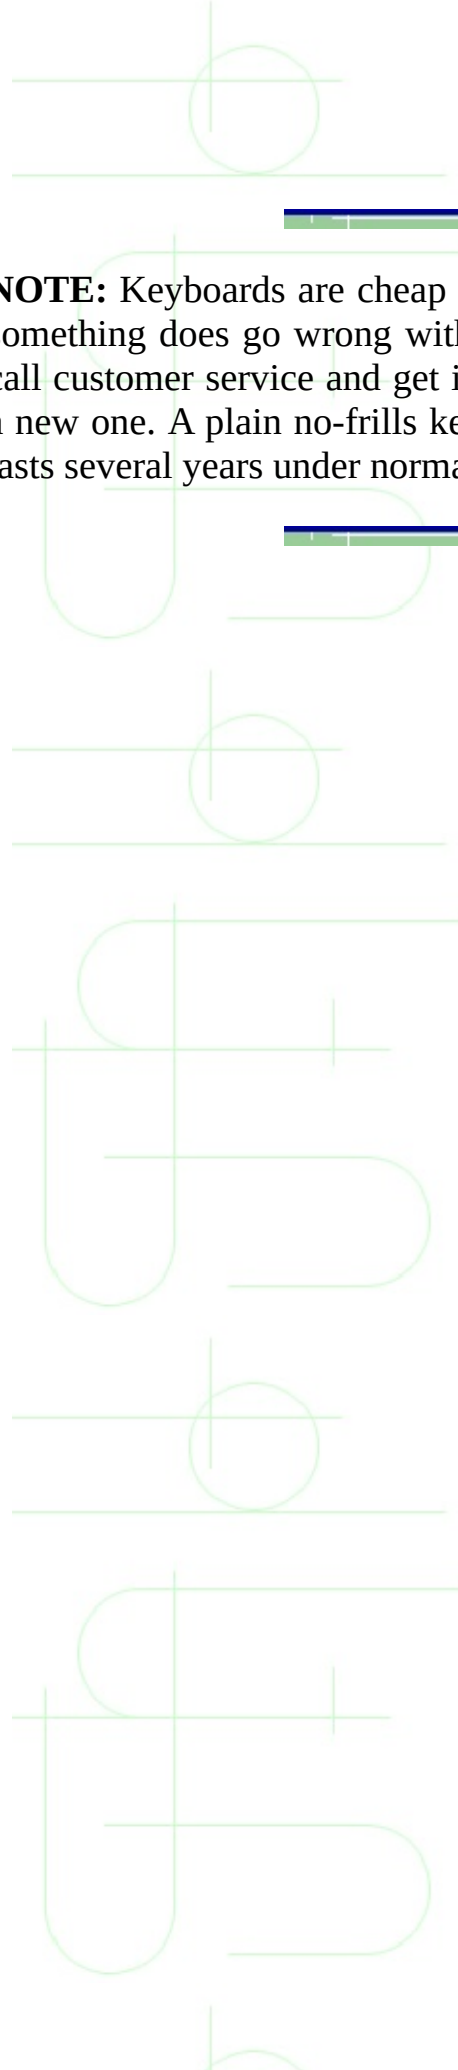
The Keyboard (or PC) Won't Stop Beeping

There's a key stuck somewhere. If there is an application running, look to see if the same letter or number is filling the screen. The key may be jammed, broken, or so dirty underneath that the contact is jammed shut. Try cleaning the keyboard and restarting the PC.

I Get a Message That Says *Keyboard error* or *Keyboard missing* When the PC Boots Up and Then Nothing Happens

Try any of the following:

- Check first to be sure that the keyboard is plugged in correctly.
- Your keyboard and mouse may have similar plugs and connectors. Check for an icon on the computer case or a description in the computer manual showing which plug is the correct one for the keyboard, and be sure it is in the right plug.
- Turn off all power to the PC, waiting a few minutes, and restarting it.
- Borrow a keyboard from a friend or buy a cheap spare at the store. Plug it in and see if it works. If you get the same error, your problem is in the PC itself-- not the keyboard-- and the PC will require professional service.



NOTE: Keyboards are cheap and generally not worth the trouble to fix. If something does go wrong with the keyboard and it is still under warranty, call customer service and get it replaced. If it isn't under warranty, just buy a new one. A plain no-frills keyboard costs anywhere from \$10 to \$30 and lasts several years under normal wear.

The Motherboard Guide

The motherboard is the foundation of any PC. All the critical subsystems, including the CPU, system chipset, memory, system I/O, expansion bus, and other critical components run directly off the motherboard. Likewise, the interconnections among these components are laid into the motherboard itself.

The mainboard is possibly the most important part of the computer. It manages all transactions of data between CPU and the peripherals. It houses the CPU and its second level cache, the chipset, the BIOS, main memory, I/O chips, ports for keyboard, serial I/O, parallel I/O ,disks and plug-in cards.

The first decision you have to make before buying a motherboard is nowadays which CPU and then which chipset you're gonna use and which motherboard to choose. There's no doubt about it - **you really should go for a brand motherboard**, preferably a brand that's present on the web, because that is by far the best way to get the latest **Flash BIOS update**, drivers and information about the board you might require.

Add-Ons

It is becoming pretty common to use a few more cards in your system than only a graphics card. A gaming system without a modem, ISDN or network card is certainly not worth being called a gaming system anymore, simply due to the fact that the only real gaming experience is generated by multiplayer games, my beloved Quake II is only one of many many others. Hence it's not out of the world if I expect that any network card should work flawlessly in any motherboard.

People who buy expensive Pentium III systems are certainly making a smart move when investing in SCSI rather than EIDE. SCSI still offers the highest disk performance, a great upgradeability for e.g. CDRoms, CD-recorders, scanners, streamers, ... and last but not least a very low trouble level. Thus I do appreciate if motherboards that are targeted towards expensive high end systems have got a SCSI adapter already onboard, a RAIDport is even better, and it's almost perfect

if it's even Adaptec's latest U2W SCSI adapter, as e.g. on DFI's new BX board. The least I would expect however, is that any SCSI adapter runs flawlessly in any board.

A sound system is nowadays a basic component of any PC. Thus I'd appreciate if there's either a decent sound system onboard or the board works fine with older ISA soundcards as well as the new PCI soundcards. In case of the latter it's useful having the new 'SBLink' onboard, which enables compatibility to the old ISA Soundblaster standard.

All in all do I think it's not really asked too much that a modern motherboard can host all these components together at the same time. If it doesn't, it may be as fast as it wants, it will still be pretty useless for any home or office user, system integrator or OEM.

Stability

Another requirement of a motherboard is certainly the stability. In the most cases boards become instable when they cannot work properly with the RAM that's plugged in. As we are fast moving towards the 100 MHz system bus as a standard, memory problems will become a lot more common. It can easily be that a board only works reliably with RAM of only a few memory vendors, other boards were designed and tested better, so that you can throw virtually any memory at it, as long as it applies to the basic specifications.

One way of testing this out is of course overclocking. If the board is running stable at a higher system bus than what it was designed for, it will most likely be rock stable at the specified system clock. However, testing a board to the limits is very difficult, because no board manufacturer and neither any CPU manufacturer would tell you which instructions are most sensitive to timing problems and overclocking.

So it's virtually impossible saying that a board or a CPU run absolutely stable at a particular clock speed, because it is very likely that the really touchy procedures haven't been ran at all. This means for the reader that you of course can be lucky as long as you are not using these procedures on your system, but it could as well be that you are using particularly the very software that will cause a crash in a board that was testified as stable.

Summary

Finally, the features of a board should be pointed out as well. I already mentioned onboard SCSI, network adapter and sound, but there are other things too. System monitoring can be an issue for people and it's certainly not wrong if a board is equipped with it. It can tell you if your fan stopped working, if your power supply fails or if your CPU gets too hot.

The new wake up features maybe worth a look at too, because it can save you from leaving your system running permanently, thus saving energy. Wake up on ring, on LAN and also on clock are features that I do appreciate. These features are used best in combination with the 'suspend to disk' feature, as well known from notebooks. AOpen is one of the few manufacturers who have this feature implemented into their boards for more than a year now. It starts your system exactly the same way you left it. The same programs are running, the same data is still there.

The above said leads to the following new evaluation scheme for motherboards in exactly this order:

- 1. Compatibility and Reliability (AGP, PCI, ISA cards, BIOS, RAM)**
- 2. Features (onboard features)**
- 3. Performance (office performance and gaming performance)**

Mice and Joysticks

A mouse is the primary input device for modern computers that feature operating systems with a graphical user interface, such as Windows 3.11 or Windows 95. While keyboards obviously excel at entering text, numbers, and symbols, your mouse is the tool you'll use to tell your computer what to do with all the data you've entered.

Joysticks are almost exclusively used with game software and help the user more effectively control the actions of computer-simulated airplanes or arcade-style games.

Defining Mice

All modern PC operating systems (Windows 3.11, Windows 95, and the Macintosh) rely on an on-screen pointer to select and execute commands. A mouse is simply an input device built to help the user control this on-screen pointer in as natural and efficient a manner as possible.

The pointer on the screen mimics the movements of your mouse. As you move your mouse, a ball encased in the bottom of your mouse rolls on the desk and in turn sends signals to the computer as to which direction to move the pointer on the screen. Move the mouse side to side, or up and down, and the on-screen pointer moves in a similar manner.

Once you have the mouse positioned to select the command or data you want to act on, you use the mouse buttons to execute the command. *The mouse controls the on-screen pointer and lets you select program icons, manipulate property sheets, and access data.*

Connecting a Mouse

Your mouse connects to your PC either through a dedicated mouse port or a standard 9-pin serial port. Once you're familiar with what the mouse connectors look like you can go ahead and connect your mouse using the following steps:

1. Save any open documents, and then close all active applications.
2. Identify which type of mouse connector you have and its corresponding connector on the back of your computer.
3. Carefully position the connector so that it matches the connector on the back of your PC, and gently press it into place.
4. In most cases, you'll have to restart your computer in order for the operating system to recognize your mouse.

NOTE: For computers running Windows 3.11, the process is a little more involved and requires installing MS-DOS mouse drivers (mouse.com and mouse.sys in most cases). If you have Windows 3.11, you'll need to refer to the manual that came with your mouse for complete installation instructions.

Using the Mouse Buttons

Most mice have two buttons. In Windows 95, the left button selects text and data, executes commands, and manipulates data, while the right mouse button accesses context menus.

Pressing a mouse button and then releasing it is known as *clicking your mouse*. You can click both left and right mouse buttons. Pressing the button and releasing it twice in quick succession is called *double-clicking*.

The Left Mouse Button

It seems simple, but there are a lot of things you can do by combining various types of mouse clicks with mouse movements. Table 13.1 gives you some examples.

Table 13.1 Left Mouse Button Operations

Task	Mouse Button
Select items	Press and hold down the left mouse button. Move the mouse to select desired text, numbers, or objects. Release mouse button. The selected text is highlighted.
Move selected items (also called <i>click and drag</i>)	Position mouse over highlighted text. Press and hold left mouse button down. While holding down left mouse button, move mouse (and the selected items) to their new location and release mouse button.
Access a menu or command	Position pointer over menu or property box button; press and release left mouse button quickly.
Start a program	Quickly press and release the left mouse button twice (double-click).

NOTE: Windows 95 lets you customize what each of the buttons does

through the Mouse Properties dialog box. This can be a help if you have a three-button mouse. See the section "Adjusting Mouse Properties" later in the chapter for examples.

The Right Mouse Button

The right mouse button is generally reserved for special uses. In Windows 95, the right mouse button accesses a context menu that lists the available options for the item you've just clicked.

The right mouse button does different things, depending on which type of item you click. See below table for some examples.

Right Mouse Button Operations

Action	Menu Options
Right-clicking a file	This pulls up a menu that asks you if you want to open, print, delete, or send the file somewhere.
Right-clicking a program	Presents you with a menu that lets you open, create a shortcut, or access that program's property sheet.
Click and drag a file (press and hold the mouse key while moving it)	Presents you with a menu that lets you choose to move or copy the file to its new location.

Using the IntelliMouse Wheel

Microsoft's newest mouse features a small wheel in between the two mouse buttons. The wheel directly controls an application's scroll bar (see Chapter 6, "Working with Applications") letting you move up and down in a document without having to move the pointer off to the right of the screen. You can also use the IntelliMouse to pan in documents by clicking the wheel and moving the mouse in the direction you'd like to pan. When you're ready to stop panning, click the wheel again.

Three-Button Mice

Some mice have three buttons. Usually the mouse comes with software that lets you customize what that third button does. Windows 95 also supports many of the more popular three-button mice and may have built-in support for that third button. If you have a three-button mouse, see the following section for tips on how to use the third button.

Adjusting Mouse Properties

Windows 95 allows you to customize your mouse to best suit your style. You can adjust the speed at which the mouse moves the pointer across the screen, the amount of time allowed between the clicks of a double-click, and how the pointer appears on the screen. Left-handed computer users can configure the mouse to work best with the left hand.

These adjustments are made through the Mouse Properties dialog box. The Mouse Properties program is in Windows 95's Control Panel. You can get to it by clicking Start, Settings, Control Panel. You should then see the Mouse icon. Double-click the Mouse icon to start the Mouse Properties program.

Configuring Mouse Double-Click Speed

To configure mouse speed when double-clicking:

1. Click Start, Settings, Control Panel and then double-click the Mouse icon.
2. Click the Buttons tab.
3. To adjust the amount of time needed between double-clicks, slide the bar in the bottom half of the Buttons tab of the Mouse Properties dialog box either left or right. (Hold the left mouse button down while moving the mouse, then release after you've moved the slider.)
4. You can test your setting by double-clicking the jack-in-the-box--if he pops up, you've clicked fast enough (see Figure 13.3).

Configuring Right- or Left-Hand Mouse Preferences

To change your mouse for right or left-handers:

1. Click Start, Settings, Control Panel and then double-click the Mouse icon.

2. To switch between right- and left-handed mouse configurations, click either the Right-handed or Left-handed option buttons in the top of the Buttons tab of the Mouse Properties dialog box (see Figure 13.3).



Adjusting Pointer Speed

You can control how fast the pointer moves on the screen in relation to your mouse movements on your desk. You can have the pointer move completely from one side of the screen to the other with only the slightest mouse movement, or you can slow the pointer down for greater control.

The pointer speed is set in the same Mouse Properties dialog box as mentioned earlier:

1. Click Start, Settings, Control Panel, then double-click the Mouse icon.
2. Once the dialog box is open, click the Motion tab

Cleaning a Mouse

If your pointer starts moving erratically, or your mouse isn't moving smoothly, it's probably time to clean your mouse. This is a simple process:

1. Carefully remove the mouse ball. On the bottom of your mouse there will be a cover plate that usually just twists off. If you have trouble getting the cover off, refer to your mouse manual.
2. Once the ball is out, roll it between your fingers to remove any dirt.
3. Inspect the inside of the mouse and carefully remove any dirt you find in there, especially any dirt that has collected on the little rollers that come in contact with the mouse ball.
4. Reassemble your mouse and give it try.

A thorough cleaning usually does the trick. If no amount of cleaning helps, you may be in need of a new mouse.

Defining Joysticks

Joysticks are basically sticks attached to a base unit that measures the distance the stick is moved left, right, up, down, or diagonally. Electronic sensors in the base unit translate those motions into motions that are understood by the computer and software.

Joysticks are almost exclusively used with game software; they are not designed to replace a mouse.

Types of Joysticks

Joysticks come in various shapes and sizes. Some are built to mimic the flight controls of an airplane, complete with buttons to fire guns and missiles. Others aren't sticks at all, but may be steering wheels designed to help the user more effectively play driving simulation games. In general, computer input devices used to help users play computer games are called joysticks.

Connecting a Joystick

Joysticks all connect with a 15-pin D-shaped connector. See Figure 13.5 for an example of a joystick connector.

Connecting a joystick is simple. To connect yours, follow these steps:

1. Save any open documents, and then close all active applications.
2. Identify the joystick connector on the back of your computer.
3. Carefully position the connector so that it matches the connector on the back of your PC, and gently press it into place.
4. In most cases, you'll have to restart your computer in order for the operating system to recognize your joystick.

If the computer doesn't initially see your joystick, you'll need to run Windows 95's Add Hardware Wizard. You can access the Add Hardware Wizard by clicking Start, Settings, Control Panel. The Add Hardware Wizard will be one of the icons in the Control Panel folder.

Calibrating Your Joystick

Calibrating your joystick lets you "center" your joystick so that when released it returns to a neutral position. It also makes certain that movements from side to side, or front to back result in smooth and predictable movement on the screen.

There are two ways to calibrate your joystick. One is through the Joystick Properties dialog box, and the other is through calibration sliders on the joystick itself.

Calibrating Your Joystick with Software

To calibrate your joystick through the Joystick Properties dialog click Start, Settings, Control Panel and double-click the Joystick icon. There, you'll see sliders that let you set the default X and Y axis positions of the joystick. You calibrate your mouse by adjusting these sliders.

Calibrating Your Joystick with Hardware

You can also calibrate most joysticks via slider bars, or wheels, that are on the joystick itself. Inspect your joystick to see if it has built-in calibration adjusters.

Modems and Faxes

Modems let you link your computer to the outside world. Connected properly, you can use a modem to send data and files to anyone else with a modem and a computer. You can also access information and download it for use on your PC.

This chapter explains what a modem is, and how to set one up for use with your PC. In addition, you learn about using a modem as a fax machine and how to troubleshoot modem problems.

Understanding Modems

A *modem* is a hardware device that lets two computers exchange information over telephone lines. Using a modem, you can connect directly to another computer that is also equipped with a modem, to an online service provider such as America Online, or to the Internet and World Wide Web. When you're connected, you can send and receive files or communicate in *real time*--which means you communicate with little or no delay, as if you were face-to-face with the person at the other computer.

Modems let you communicate over standard telephone lines with other computer users. You can transfer data, exchange electronic files, and even carry on a typed conversation in real time.

These days, a modem is a vital part of a PC; without one, you'll miss out on a wide range of opportunities. Here's a small sampling of things you can do with a modem:

- Exchange data directly with other computer users.
- Access an online service or the Internet to gather information or to distribute information about anything, anywhere.
- Exchange electronic mail.
- Play interactive online games.
- Make hotel or airline reservations.
- Check stock prices.

The word *modem* is actually an abbreviation of the terms *modulator-demodulator*, which describe how a modem works.

Here are the basic steps involved in transferring data using modems:

1. Software in your computer activates the modem, which listens on the

phone line for a dial tone, then dials a specified number.

2. The receiving modem on the other end answers the call and sends out a signal which the sending modem recognizes, establishing a connection.

3. The modem sending information converts the computer's digital signals into analog signals that can be sent over the phone lines (this is called *modulating*).

4. The receiving modem translates the incoming data back into digital signals that its computer can understand (this is called *demodulating*).

5. When the online session is over, the computers issue a command telling the modems to hang up.

In order to make a connection using a modem, you need the following:

- *Access to a telephone line.* Modems have two jacks where you plug in standard telephone cables. You use one jack to connect the modem to the telephone line, and the other jack to connect the modem to a phone.
- *Data communications software.* You need software to control the modem. Data communications software issues commands such as telling your computer when to dial and when to hang up. It also lets you set communications parameters such as the data transfer rate. The software must be compatible with both the modem and your PC. (Windows 95 comes with a basic data communications software called HyperTerminal.)
- *Software drivers.* You must have the necessary drivers for the modem to work with your operating system. Windows 95 comes with drivers for most modems.

Choosing a Modem

Choosing a modem is not as difficult as choosing some other computer components because there is less variety. There are two basic types of modems:

- *Internal.* These are hardware boards you plug into an expansion slot in your PC's system unit. Internal modems are convenient because they don't take up desk space, and they use the computer's power supply so they are on whenever the computer is on.
- *An internal modem plugs into an expansion slot inside your PC. The phone jacks are accessed through a port on the back of the system unit.*
- *External.* These are connected to the PC by plugging a cable into a port on the system unit. External modems also have their own power cords. You must turn on the modem when you want to make a connection. One benefit of an external modem is that you can move it to another computer if necessary.

An external modem has controls and status lights on the front that you can use to monitor transmissions.

NOTE: You can attach a telephone to either kind of modem so when you are not using the modem, you can use the phone to make regular calls.

Aside from choosing an internal or external modem, the other main factors to consider are:

- *Communications speed.* A modem's speed, called the *baud rate*, is measured in bits per second (bps). Currently, speeds range from 300 bps to 56Kbps, but to satisfy the need for speed, manufacturers are developing faster modems all the time.
- *Cost.* The faster the modem, the higher the cost--although faster transmission times may save you money on your phone bills. Also,

external modems generally cost about \$30 to \$50 more than internal modems. For example, a 56Kbps internal modem will probably cost between \$100 and \$200. A 56Kbps external modem would range from \$150 to \$225.

NOTE: If you only use your modem to send and receive e-mail messages, you may be able to get by with a modem as slow as 14.4Kbps. If you use your modem to transmit and download files, or if you like to surf the World Wide Web--which uses a lot of graphics images--you should purchase the fastest modem available.

-
- *Communications software.* Most communications software programs have features for storing telephone numbers as well as for automatic dialing and answering. Even more important, you use the communications program to set *communications parameters*, which are the rules the modem uses to make the call and transmit the data. You want to choose a modem that supports the communications parameters you plan to use.

NOTE: The communications parameters set for the sending modem must match the communications parameters set for the receiving modem.

Modems also operate under communications parameters. Modem communications parameters include:

- *Transmission speed.* As mentioned earlier, speed can range from 300 bps and up.
- *Transfer protocol.* This helps to make sure the data the modem transfers is error-free. Common transfer protocols include XMODEM, YMODEM, and ZMODEM. The Internet uses FTP (File Transport Protocol), and some online services have their own transfer protocols. For example, CompuServe uses CompuServe B.

- *Data format.* This is usually text for unformatted text files, and binary for graphics or formatted text files.
- *Compression standard.* This allows modems to compress data before transmitting it, effectively increasing transmission speeds. Common compression standards include MNP 5 and CCITT V.42bis.

Selecting a Fax Modem

One final consideration when purchasing a modem is whether or not you want fax capabilities. Because most modems currently available include fax capabilities, it's not much of a decision. Fax modems generally don't cost any more than standard modems, so in addition to using the modem to communicate with other computers, you can use it to communicate with facsimile machines.

With a fax modem, you can:

- Transmit standard computer files to a receiving fax machine.
- Receive fax transmissions and store them as computer files.

NOTE: If you already have a stand-alone fax machine, you do not have to use the fax options on the modem. The device will work fine as a plain modem.

To use the fax capabilities, you need to have specific fax software or data communications software that includes fax support.

Connecting a Modem

Connecting a modem involves attaching the modem device to your PC and to the telephone line. The modem should come with the necessary equipment, including a standard RJ-11 telephone cord. The steps you take depend on whether you have an internal modem or an external modem.

If you have an internal modem, you must install it by removing the cover of the system unit and plugging the modem card into a vacant expansion slot. Unless you are extremely confident about working inside your computer, you should have the modem installed by a professional computer technician.

If you have an external modem, however, you can easily connect it yourself. To connect an external modem to your PC, follow these steps:

- 1.** Locate the serial cable that came with the modem. It should have a 25-pin male connector at one end and a 25-hole female connector at the other end.
- 2.** Plug the female connector into an open 25-pin serial port on the PC's system unit, and plug the male connector into the modem device.
- 3.** Secure the connectors using the screws on either side.
- 4.** Locate the power cord for the modem device and plug it into an electrical wall socket, power strip, or uninterruptible power supply.

After the modem is connected to your PC, you must connect it to the phone line. This procedure is the same for both external and internal devices.

To connect a modem to the phone line, follow these steps:

- 1.** Locate the two phone jacks on the modem, labeled **Phone** and **Line**. They will be on the outside of an external modem, or on the back of the system unit if you have an internal modem. They look just like the phone jacks on your telephone.
- 2.** Plug one end of a standard RJ-11 telephone cord (it should come with the

modem) into the Line jack, and the other end into the wall jack for your telephone.

3. Plug one end of another RJ-11 telephone cord into the Phone jack, and the other end into your telephone.



TIP: If you have a phone line dedicated for use with your computer, you do not need to connect the modem to the telephone. You only need to connect to a phone if you want to be able to use the phone for regular calls when the modem is not in use.



Understanding Your Modem Controls

An internal modem is integrated into your PC; you will interact with it using your data communications software. If you have an external modem, however, you can control and monitor certain functions on the device itself. For example, status lights on the front of the modem let you know whether the modem is on, and whether it is in *send mode*, ready to dial and transmit data, or *receive mode*, ready to answer and download data.

If you look at the front of your external modem, these are some of the controls you are likely to find:

- **Power On light.** An external modem must be turned on in order to work. When you press the on/off switch, the power light will come on.
- **Speed indicator.** Most modems can operate at different speeds; this light tells you which speed is currently in use.
- **Mode indicator.** If you have your modem set to answer an incoming call, the Answer mode indicator will be lit. If it is set to dial, the Dial mode indicator will be lit.
- **Error indicator.** If there is a problem with the transmission, the Error light will be on.

Be sure to read the documentation that comes with your modem for complete instructions on understanding the modem controls and status lights.

Setting Up Your Modem to Work in Windows 95

Windows 95 uses the Install New Modem Wizard to help you identify and set up newly installed modems. Before starting the wizard, it helps to know the manufacturer and model of the modem. Also, if you have an external modem, make sure it is turned on.

To use the Install New Modem Wizard, follow these steps:

1. Double-click My Computer and then double-click the Control Panel folder icon. Alternatively, click the Start button, then open the Settings menu and click Control Panel.
2. Double-click the Modems icon in the Control Panel folder to open the Modems Properties dialog box. Click the Add button to start the Install New Modem Wizard, then choose Next.
3. Windows 95 immediately tries to detect and identify the modem. When it does, it displays the Verify Modem screen of the wizard, displaying the name or type of modem that it finds. It also identifies the communications port where the modem is connected. If the information is correct, choose Next, then choose Finish.
4. If Windows 95 cannot detect the modem, the wizard displays a screen where you can select the modem information yourself. (You can display the same screen by clicking the Change button if the information on the Verify Modem screen of the wizard is incorrect.) Select the modem's manufacturer in the Manufacturers list box. Select the specific modem in the Models list box, then choose OK.

NOTE: If your modem is not on the list, choose the Standard Modem Types manufacturer and select the model according to speed. Or, click the Have Disk button and follow the instructions to install a vendor-supplied driver.

5. Choose Next, then choose Finish. Windows 95 installs the necessary drivers for the modem, then closes the wizard. Choose Close to close the Modems Properties dialog box.

On the Verify Modem page of Windows 95's Install New Modem Wizard, you see the port and modem type Windows has detected. Click Change to select a different modem.



NOTE: If Windows 95 detects the new device when you first turn on your PC, it will prompt you to run the Add New Hardware Wizard to detect it and set it up. Go ahead and use the Add New Hardware Wizard; it performs the same function as the Install New Modem Wizard, but will detect and set up other types of hardware devices as well.



Troubleshooting Your Modem

When a modem doesn't work, there could be problems with the hardware, the communications software, or the software settings. Here are a few tips for locating the problem and fixing it.

The External Modem Isn't Working

If your external modem is not working, here are a few things to check:

- Make sure the modem is turned on.
- Make sure the cables between the modem and the computer attached correctly and securely.
- Make sure the modem power cord is plugged in.
- Make sure the modem is set to the correct mode. Use Answer mode if you are trying to receive a transmission, or Dial mode if you are trying to initiate a transmission.

The Modem Dials, but Doesn't Connect

If your modem dials, but doesn't connect, here are a few things to check:

- Make sure both the sending and receiving modems are set to the same communications parameters.
- Make sure the modem is dialing the correct number. If another modem doesn't pick up on the other end, the call can't go through.
- Make sure there isn't interference on the telephone line. Modems are more sensitive to static and other interference on a line. Someone picking up an extension phone, crossed wires, even thunderstorms in the area can cause enough disturbance to cancel a call. Make sure the line is clear, or try a different phone number.

The Modem Doesn't Get a Dial Tone

If your modem doesn't get a dial tone, here are some things to check:

- Make sure the RJ-11 telephone cords are connected correctly. You may have the Phone jack connected to the telephone line, and the Line jack connected to the telephone.
- Make sure the telephone is working.
- Make sure your communications software is correctly installed.
- And in some cases with pulse dial de-select in modem properties wait for dial tone before dialing

Windows or the Communications Software Displays a Message Saying It Can't Find the Modem

If Windows or your communications software displays a message telling you that it can't find your modem, here are some things to check:

- If you have an external modem, make sure it's plugged in, that the cable is attached securely to the port on the system unit, and that the device is turned on.
- Make sure the correct driver is installed. If necessary, run the Install New Modem Wizard again.
- Make sure the COM port setting is correct. If the modem is set up with COM1, but Windows 95 thinks it's in COM4, Windows won't be able to find it. Check the COM port setting on the Properties page of the Modems Properties dialog box, and make sure it is the same as the COM port setting in your communications software.
- The interrupt request (IRQ) setting may be conflicting with another device. This is a relatively technical problem, and you should have a qualified technician look into it.
- Check for an IRQ conflict. To do so, from the Start menu, choose Settings, Control Panel. Click the System icon, then click the Performance tab of

the System Properties dialog box. The status area of this tab will report an IRQ conflict if one exists. You're not advised to try to fix this conflict, but the information will help you accurately report the problem to a technician.

- The modem may be damaged. Have a technician check it out.



Monitors

Your monitor is plugged into a video card (also called a *video adapter* or *graphics adapter*) in your computer. In most PCs, the video card fits into an expansion slot on the motherboard. (A few PCs have the video card built into the motherboard, but the end result is the same.)

Your monitor is connected by a cable to the video card. (We'll look more closely at that connection in the later section "Hooking Up Your New Monitor.") The video card receives instructions from the processor and sends them to the monitor. Any information you see on-screen makes use of the video card as a messenger to the monitor. The monitor doesn't do any processing itself. The monitor only displays the information that the video card tells it to.

Monitor Sizes

The two most common monitor sizes are 15-inch and 17-inch. If you have an older, hand-me-down PC or a very inexpensive starter PC, you may have a smaller 14-inch monitor. 21-inch monitors are also available but mostly used by graphics professionals.

TIP: In the U.S., monitors are measured diagonally from one corner of the tube to the other, but part of the tube is hidden by the case so the actual viewing area is reduced. Lately, most stores have begun to also list the "viewable" size.

VGA and SVGA Monitors

The two most common acronyms you will see on current monitors are VGA or SVGA. Both of these terms generally refer to how many dots (or pixels) in each direction the monitor can display. VGA is 640x480 (width by height) and SVGA is 800x600. This measurement is called the monitor's *resolution*, and more is better! Most new monitors are capable of displaying at least SVGA quality. In fact, 1,024x768 is somewhat of a minimum to look for.

Almost any VGA or SVGA monitor made in the last few years is capable of displaying any of these resolutions. However, it's actually the video card that determines what resolution your monitor displays at any time. The monitor is capable of switching from one resolution to another on command from the video card.

You'll want to consider the size of your monitor when deciding which resolution to use. While these monitors can display many possible resolutions, higher resolutions may be too small to read on small monitors, and low resolutions will look huge and clunky on large monitors. Below table recommends some good resolutions for common monitor sizes.

Recommended Resolutions for Common Monitor Sizes

Monitor Size	Recommended Resolutions
14-inch	640x480 800x600
15-inch	800x600 1,024x768
17-inch	1,024x768 1,280x1,024

NOTE: EGA, CGA, and MDA are older, greatly outdated monitor types that do not work with Windows 95.

Interlaced Monitors

The picture that you see on the monitor at any time is constantly being *refreshed* (redisplayed) on the screen. Most monitors are capable of refreshing every pixel on the screen about 60 times per second. However, some monitors take a slight shortcut and refresh only every other line of pixels each time. So, the monitor still may refresh 60 times a second, but each line of pixels is refreshed only 30 times a second--a process called *interlacing*.

Interlaced monitors are always less expensive than non-interlaced monitors with similar features, but interlaced monitors flicker more than non-interlaced monitors. If you work under fluorescent lights or spend long periods of time in front of the screen, you should consider going with a non-interlaced monitor.

VGA and SVGA Video Cards and Resolutions

Like monitors, video card capabilities are measured in terms of resolutions. The two most common sets of resolutions are given special acronyms: VGA and SVGA. VGA resolution is 640x480. SVGA is 800x600 and is usually also used to refer to any resolution higher than this. (However, you may sometimes see XGA used when referring to super-high resolutions like 1,024x768 and higher.)

Any new PC will come with an SVGA card. Some older PCs will still have a VGA card. While it's not the most recent, a VGA card will do just fine as long as you don't want to play intense video games or have several windows visible on-screen at the same time.

Your video card and monitor work together as a pair. The highest resolution you will be able to use is limited to the highest that either can use. So, a video card that can display up to 1,600x1,200 resolution with a monitor that can only display 1,024x768 will be limited to 1,024x768 as a pair.

Video Cards and Colors

Another measure of video card performance is how many colors it can display at once. While you might think that a video card should be able to display "all" of the colors, it's not quite that simple.

Video displays use three different basic colors--red, green, and blue (called *RGB color*)--to display all of the possible colors. Each pixel on-screen actually has a red, blue, and green component. By varying the brightness of each of these colors for each pixel, the video card mixes the three basic colors to come up with all of the hues it displays.

But because a video card is digital, each of these colors has to be broken down into digital increments. This is where video color gets a little confusing, so follow a little math.

Color capabilities are coded with two sets of terms. In some settings, you'll see these described by the number of colors displayed, in others by the number of "bits" used to display the colors. Below Table puts these two measures together.

Terms for Color Capabilities

Bit Depth	Math	Colors Displayed	Other Name
4	$2^4=16$	16 color	
8	$2^8=256$	256 color	
16	$2^{16}=65,536$	65K	Hi-color
24	$2^{24}=16,777,216$	16 million	True-color

So, how does this relate to the digital increments? 256 colors means that the card can have 256 sets of increments across the three colors. A video card can actually display any of the total possible 16 million colors. What is really limited at any time is the number of colors it can remember and display at once. The more colors your programs need to display at once, the more memory your card will need. This is the topic of the next section.

Video Card RAM

RAM in your computer stores the information and programs that you are working with. RAM on the video card stores information about each screen pixel--what color it is, and how bright it is.

The more pixels you are displaying on-screen at once (higher resolution), the more RAM it takes on the video card to track the colors for all of the pixels. And the more colors you are displaying, the more RAM it takes to track the color for each pixel.

An 800x600 display has about 50 percent more pixels than a 640x480, and a 1,024x768 has about 160 percent more. So, to display the same number of colors in these higher resolutions will take about 50 percent and 160 percent more video RAM, respectively.

Below table shows maximum numbers of colors that can display at various resolutions with different amounts of RAM. The numbers shown are "general-rule-of-thumb" numbers, and will vary from card to card.

Maximum Number of Colors Displayed at Different Resolutions

RAM	Resolution	Maximum Colors
1M	640x480	16 million
	800x600	65,000
	1,024x768	256
2M	640x480	16 million
	800x600	16 million
	1,024x768	65,000
4M	640x480	16 million
	800x600	16 million
	1,024x768	16 million

Video Cards, Video Games, and Full-Motion Video

For general everyday computing needs like using a word processor, spreadsheet, or presentation program, video card speed is not a big issue. But for playing action-oriented video games or full-motion videos, the speed of the video card is important. The faster the video card, the less "jerky" the action will seem as you fly through space in your favorite flight simulator game, or the fewer frames of a movie the card will miss.

For video cards, there are a couple of key benchmarks. ZiffDavis Benchmarks has a bench-marking group that measures video cards in terms of *graphics winmarks*. This is essentially a measure of how fast the card is in Windows and is important for Windows games. Similarly, they have a DOS graphics rating for DOS games.

Another measure of graphics speed is *polygons per second*. This measures how many shapes per second the card can draw.

With any video benchmark, faster is better and higher scores are always faster.

Special Considerations for Notebook Displays

Notebook displays have a few unique factors and features that don't apply to desktop monitors. In general, these features relate to the type of display technology used in notebook construction.

Passive Matrix, Dual-Scan, and Active Matrix Displays

The first term you may encounter in regard to notebook displays is *passive matrix*. Most bargain-priced notebooks use this type of display technology. The display can be dim and hard to see if you aren't looking at it in the right light from the right angle. An improved version of this is called dual-scan; dual-scan screen notebooks are sometimes sold under the descriptions DSTN or CSTN.

Another type of notebook computer screen technology is *active matrix*. Active matrix screens--sometimes listed as TFT in system descriptions--are used in top-of-the-line notebooks. Though more expensive, active matrix displays are brighter and crisper.

Notebook Display Size

The other key difference in notebook displays is size. As explained earlier, viewable area on a desktop monitor is usually smaller than the monitor's stated size. But, a 12.1-inch notebook monitor measures 12.1 real inches diagonally. So, you can get away with a smaller display on a notebook because you don't lose any unviewable area. 12.1 inches is becoming a good entry point. 13-inch displays are starting to come on a lot of laptops, but right now they are probably not worth the big price premium.

Notebook Display Resolution

In general, notebook displays are limited to 640x480 resolution for a 10-11-inch display. Most 12-inch displays can run at 800x600, and some of the new 13-inch displays run at up to 1,024x768.



Hooking Up Your New Monitor

When you are ready to hook up your monitor, carefully unpack it and save all of the packing and boxes. If you ever need to move or have it serviced, you'll want the box and packing to transport it. Look in the box to see if there are any disks or CDs. If there are, read the labels; they may be software that you need to use the monitor. Or, they may just be freebies like games or giveaway bonus software.

The monitor should have two cables. One should be a familiar power cord with a standard three-prong plug. The other is the cable you will use to connect the monitor to the video card.

- 1.** Shut down the PC and turn off the power to the PC and monitor, then disconnect the monitor and PC from the electrical outlet.
- 2.** Look at the back of the PC and find where the monitor cable is attached. Carefully loosen the thumbscrews of the old monitor cable and detach it from the card. Be sure not to disturb other wires and make sure you see where you disconnected it from.
- 3.** Move the old monitor out of the way and move the new monitor into its place.
- 4.** Connect the new monitor cable to the video card connector that you detached the old monitor from. Don't force the connection. The connector is longer on one side than the other so you may need to turn it over to line it up.
- 5.** Once you feel the connector slide into place, tighten the thumbscrews. Some people like to leave these loose. I prefer not to have my monitor cable fall off while I'm working.
- 6.** Plug the monitor into the power outlet and then turn on the PC and monitor.

If the picture looks odd, you may need to adjust it. We'll look at adjusting the monitor settings in the next couple of sections.

Setting Up a New Monitor in Windows 95

To get the most out of your monitor in Windows 95, you'll want to tell Windows what type of monitor you have. To do this:

1. Click the right mouse button on any blank area of the desktop and choose Properties from the pop-up menu.

Or, click the Start menu, choose Settings, then choose Control Panel and double-click the Display icon in the Control Panel folder.

2. In the Display Properties dialog box, click the Settings tab, then click the Change Display Type button to open that dialog box

3. Click the Change button to open the Select Device dialog box, In the middle of the dialog box is the area where we choose the monitor type.

4. Scroll through the Manufacturer list in the dialog box and select the brand of monitor you have. Then, in the Models list (to the right) select the model. Click OK.

NOTE: If your monitor came with a disk with driver software for Windows 95, insert the disk and click the Have Disk button. Click the Browse button in the Install from Disk dialog box to select the drive with the installation disk, then click OK. This will then give you a list of possible monitors to choose from in the Select Device dialog box in step 4.

5. Windows now enters your selection in the Change Display Type dialog box. Click Close in that dialog box, then click OK in the Display Properties dialog box and you are ready to go.

CAUTION: Choosing the wrong monitor display type during the preceding procedure can damage your monitor. If your brand and model aren't listed, call the maker's technical support number to find out what setting you

should use.



NOTE: Some cool graphics cards come with their own utilities for changing the monitor settings. In fact, they may even add extra features in Windows 95's own Display Properties dialog box


Each of these cards is different, and if you want to make the most of the extras they provide, read the manual or help that came with it. If your display properties dialog box has some extras in it like this, some of the tasks in this chapter may work a little differently with your computer.




Adjusting Your Monitor Settings

On the front of your monitor, there are usually several controls that affect the appearance of your monitor display. These controls may be small dials or buttons. Here are some typical controls that you will find and what they do:

- **Brightness.** This controls the overall brightness of the picture on-screen. Use this when the picture is either too bright or too dark.
- **Contrast.** This changes the contrast between the light and dark areas displayed on-screen. Adjust this if your picture looks washed out, gray all over, or if the whites hurt your eyes.
- **Horizontal and vertical sizing.** If the picture on-screen doesn't reach close enough to the top and bottom or the sides, use these controls to resize it.



NOTE: On most monitors, there will be a small black border area around the picture when it is correctly sized. Adjusting the picture too large to try to get rid of this border will distort and curve the edges of what's on-screen.

- 
- **Centering.** If your picture isn't centered on-screen, use the centering controls to adjust it. You may need to work with the centering and sizing alternately to get it centered and sized.
 - **Deguass.** This is almost always a button. Monitors generate large magnetic fields; over time these fields can build up and distort the picture. Other magnetic items (like un-shielded speakers) can also generate fields that distort the picture. If this happens, press the Deguass button. You will usually hear a pop, and the picture will disappear and resize. When it comes back, it should look better. If not, you may want to see the troubleshooting section at the end of the chapter.

Most monitors' controls have "centers" that put these settings at their middle grounds or factory presets. On dial controls, you may feel a small click as you

turn the control to the center point. On button controls, there will usually be an indicator on-screen that shows the center.



TIP: Notebook computer users will usually find the monitor brightness and contrast controls as function keys on the keyboard. Notebooks don't have the ability to adjust the display size or recenter it, and there is no need to degauss them. See your notebook manual if you need details on using those.



Changing the Numbers of Colors Windows Can Display

How many colors do you want to be able to see on-screen? That answer depends a lot on what you use your PC for and what kind of video card you have. If all you use your PC for is word processing or spreadsheets, the standard display options of showing 16 or 256 colors at once should work for you.

But, if you like to look at pictures or video or play photo-realistic games, you may find that the colors look all wrong on-screen with the default settings. To change the number of colors Windows 95 can display at once, follow these steps:

1. Click the right mouse button on any empty space on the desktop and choose Properties from the pop-up menu.
2. Click the Settings tab in the Display Properties dialog box.
3. Click the Color Palette drop-down list and choose from one of the options. The options available here will depend on your video card and desktop display size. Windows won't display options that are over the limit for your card.
4. Click OK once you have made your selection.
5. In order for this change to take effect, Windows must restart. Click Yes to allow Windows to restart.

Changing the Size (Resolution) of the Display

You can increase the size of the desktop display area to be able to fit more visible open windows on-screen at once, or see a bigger area of the windows you have open. However, because your monitor doesn't actually change size when you change the resolution of the display area (that would take a small miracle), making the display area "bigger" makes each item on-screen look smaller.

To change the display area size, follow these steps:

- 1.** Click the right mouse button on any empty space on the desktop and choose Properties from the pop-up menu.
- 2.** Drag the slider for the Desktop Area to a higher or lower setting. If the slider is all the way to the right edge of the control area, then you can't make the display area larger. Likewise, when it is all the way to the left, it is as small as it will get.
- 3.** Click OK to resize. Windows will inform you that this setting may cause problems. Click OK to continue.
- 4.** If all goes well, the picture will "hiccup" a little and return to normal. If it looks normal, click OK to keep the new setting.

If things didn't go well and the display doesn't look right, click Cancel and Windows will return you to the previous settings. If the display is so messed up that you can't see it to cancel it, Windows will return to the old settings after 15 seconds if you do nothing.

Changing Wallpaper and Patterns in Windows

Wallpaper and patterns are just decoration on top of Windows. While they don't really relate to the monitor, you do see them on the monitor. So if you want to change them, follow these steps:

1. Click the right mouse button on any empty space on the desktop and choose Properties from the pop-up menu.
2. Click the Background tab of the Display Properties dialog box.
3. In the Pattern list, select a pattern. This pattern will be applied to your desktop. A preview of any changes you make here is shown in the drawing of the monitor in the top of this dialog box.
4. To apply one of the existing wallpapers, select it in the Wallpaper list. You will probably want to select the Tile option to fill the desktop with the wallpaper as most of these are very small.
5. If you have a picture that you would like to use as wallpaper, click the Browse button to find the right file and click OK after you have selected it. Usually, pictures are larger and look best with the Center option selected.

TIP: Windows can only use pictures in the Windows Bitmap format for wallpaper. If you are downloading picture files from the Internet, most of these are in GIF or JPEG format. You'll need to use a graphics program like Microsoft Photo Editor (that comes with Office 97) or a shareware program like PaintShop Pro to save these as bitmaps that can be used as wallpaper.

6. Click OK when you have made your changes to close the dialog box.

Cleaning and Protecting Your Monitor

Monitors, like most computer parts, are susceptible to damage and decay from dirt, dust, and misuse. A few simple guidelines will help extend the life of your monitor.

First, try to keep the exterior clean and free of dust. You can use any lint-free rag to dust the case and screen. For a small investment of a couple dollars, you can get monitor cleaning "wipes" that will do a better job keeping dust off of the screen. You can use a common, non-abrasive household cleaner if you want but it's not necessary. If you do, use it sparingly, be sure the monitor is off, and spray it on the rag, not the monitor. If you use the monitor in a really dusty area, you will want to get a cover for it and keep it covered when not in use.

Second, if you have speakers attached to your PC, be sure they are shielded to prevent electrical static from passing between your PC and speakers. If they aren't, move them away from the monitor. If you are shopping for computer speakers, be sure to buy shielded ones. Almost all speakers designed for PCs are shielded, but you will find some cheaper models that aren't. Likewise, speakers that you bought outside the computer department at the local electronics store to use with your Walkman may not be shielded.

Finally, your monitor throws off a lot of heat. It has a lot of openings on the top and maybe the sides and back to let it cool. Don't stack papers, books, or other junk on top of the monitor or around the sides. In addition to causing the monitor to overheat, these could be a fire hazard, or your cheap, plastic toy you used as a monitor hood ornament could melt and ruin the guts of the monitor.

Troubleshooting Display Problems

The following sections present some common monitor and display problems and offer some potential solutions.

The Picture Is All Wavy, Fuzzy, and Won't Work in Windows

This usually means that the video card or monitor settings are wrong. Shut down the computer and restart it in Safe mode; then check to be sure your display settings match the brand and model of monitor and card that you have. If this doesn't work, try setting the display settings to the generic VGA or SVGA display card and monitor types. Contact the video card vendor's technical support to see if they have a more recent set of drivers.

The Picture Changes Size Going from DOS to Windows

This is fairly normal and is only a problem if the change is so dramatic that either DOS or Windows screens are too small or large to use. If this is the case, check your display settings to be sure the correct monitor is selected.

The Picture Keeps Getting Smaller

Bad news: The picture tube in the monitor is going bad. If it's under warranty, send it back for repair. Otherwise, it's probably a better investment (and no more expensive) to shop for a new monitor.

I Can't Resize the Screen on My Notebook

Notebook screens are a fixed size. They can't be resized like a monitor. Some will allow you to resize the desktop resolution, and if you do, you may have to pan to see the whole desktop on-screen at once.

Windows or a Program Tells Me I Can't Display the Right Number of Colors

Your screen resolution may be too high to display more colors with your video

card. Try selecting a lower screen resolution in the Display Properties dialog box.

The Monitor Doesn't Come on When I Turn on the PC

Look for a separate on/off switch on the monitor itself and turn it on there.

The Monitor Doesn't Turn off When Windows 95 Shuts Down the Computer

Look for a separate on/off switch on the monitor itself and turn it off there.

Can I Fix or Upgrade the Inside of My Monitor Like My PC?

No! The inside of your monitor is a dangerous electrical place even when unplugged. Leave servicing the monitor to a professional.



Printers

Printing files and documents can be as simple as clicking an application's Print button. The hard part is deciding which printer is right for you. This chapter defines printers, and explains the differences between the most common printers available. In addition, you learn how to connect a printer to your computer, and how to make sure it will work with Windows 95.

Printers Defined

A *printer* is the computer component that lets you create copies of the information stored in your computer on paper. The printed material is often called *hard copy*, to differentiate it from the data stored on a disk, or held in the computer's memory.

There are three basic types of printers available for use with personal computers:

- *Laser printers.* These combine a magnetic roller with powdered ink called toner to transfer high-quality characters or images onto a page.
- *Inkjet printers.* These have small nozzles that actually spray fast-drying ink onto the page to form characters or images.
- *Dot-matrix printers.* These use a print head to strike an inked ribbon against paper, like a typewriter, creating characters out of a series of dots.

The type of printer you choose depends on your budget and the type of output you need. For example, if all you need to print are grocery lists, you may be happy with a dot-matrix printer. In general, dot-matrix printers are noisier, slower, and produce a poorer-quality image than do laser or inkjet printers, but they are also less expensive. If you need to print newsletters, brochures, or illustrated reports, you will probably want a high-quality laser printer.

In general, there are three main factors to consider when purchasing a printer:

- *Cost.* Printers are available ranging in price from a few hundred dollars to a few thousand dollars, but the purchase price is just the beginning. Paper, ink, toner, printer ribbons, and the other *consumables* are required to keep the printer running, and they all cost money.
- *Resolution.* Resolution is measured by how many dots per inch (dpi) the printer can print. The more dots per inch, the higher the resolution of the printed image. The higher the printer's resolution, the better is its quality.
- *Speed.* Printer speed is rated by how many characters printers output in a

second (cps) or how many pages printers output in a minute (ppm). The higher the number, the faster the printer.

Although the primary considerations involved in selecting a printer are usually cost, speed, and resolution, there are other issues that come into play as well. Available fonts, the ability to print in color, and the amount of maintenance involved may weight your decision:

- *Noise.* Impact printers such as dot-matrix printers are noisier than non-impact printers such as laser or inkjet printers. If you have a small or crowded office, noise can be an important factor.
- *Software support.* PCs communicate with printers using software programs called printer drivers, and some printers come with built-in printer languages, such as PostScript. Be sure that the software programs you use support the printer driver and printer language for the printer you purchase.
- *Fonts.* A font is the style of letters that your printer puts on the paper. They are important if you want control over the appearance of text on the printed page. Most printers come with built-in fonts called internal or resident fonts. Some dot-matrix printers may only have a few resident fonts, while some laser printers may have dozens. In addition, laser and inkjet printers can use soft fonts that come with your software applications.
- *Color.* Color spices up most printed documents, but it can significantly increase the cost of printing. Not only do color printers cost more, but the ink, toner, or ribbon cartridges cost more as well. Although the cost of laser color printers is coming down, color inkjet printers still offer the best combination of cost effectiveness and print quality.
- *Maintenance.* The amount of time you spend cleaning a printer, refilling a paper tray, or changing a toner cartridge may vary greatly from printer to printer.

Laser Printers

Laser printers are similar to copy machines--they use laser beams to burn special toner onto the page to create a permanent impression. They create high-quality output at a relatively fast speed, without making too much noise. The downside is the price; most black-and-white laser printers cost \$500 or more, and each toner cartridge (which lasts about 3,000 pages) can cost an additional \$50.

When considering a laser printer, here are some things to keep in mind:

- Laser printers offer resolutions ranging upwards from 300 dpi-600 dpi is common, and very high-quality printers can output at resolutions as high as 3,386 dpi!
- Laser printers are available in speeds of 4, 8, or 12 ppm. However, keep in mind that these are the top speeds. The actual speed will depend on whether you are printing plain text, or complex characters and images.
- The amount of desk space a printer uses is called the *footprint*. If you don't have much room, you'll want a small footprint printer.
- Some printers can hold only 50 sheets of paper, while others can hold 250 sheets. Some have more than one paper holder, which increases capacity, and makes it possible to switch between different sizes of paper. That means you don't have to reload when you want to print legal size pages or envelopes.

Inkjet Printers

Inkjet printers spray a fine, quick-drying ink through small nozzles to produce characters and images on paper. Although the results are not quite as sharp as those of laser printers, inkjet printers provide very good quality output at a lower cost. *Laser printers cost more than other printers, but they may be worth the price because they are fast, quiet, and produce high-quality text and graphics.*

When considering an inkjet printer, here are some things to keep in mind:

- Inkjet printers are quiet.
- Inkjet printers, which start at about \$200, are less expensive than laser printers, yet produce output almost as well.
- Inkjet printers are slower than laser printers, but faster than dot-matrix printers.
- Replacement ink cartridges are expensive, costing about \$45 each.
- Many inkjet printers can use color ink cartridges.

Dot-Matrix Printers

Dot-matrix printers are the cheapest printers available. They create text and images on the page by hammering several small pins against an inked ribbon. The more pins used, the better the image--9-pin and 24-pin are common options. The 24-pin printers produce a better quality output, but are somewhat slower than the 9-pin printers.

Print quality for dot-matrix printers is often described in terms of mode: draft mode (low resolution), near-letter-quality mode (medium resolution), or letter-quality mode (high resolution). The speed depends on the mode, with draft mode being the fastest.

When considering a dot-matrix printer, here are some things to keep in mind:


- Dot-matrix printers are the cheapest. Prices start at about \$100; the ribbon cartridges last a long time and cost about \$15 apiece.
- The sound made by the pins banging away can be quite loud; 24-pin printers are louder than 9-pin printers.
- Most dot-matrix printers can use different types of paper, including continuous feed as well as cut-sheet paper. In addition, they can be used to print on multipart forms. Some let you switch back and forth between the different types without reloading.
- Narrow carriage dot-matrix printers can accommodate standard letter size paper, while wide carriage dot-matrix printers can handle 11x17-inch pages. Wide carriage printers cost generally about \$100 more than narrow carriage printers.

Connecting Your Printer to Your PC


Most printers have two cables: one is a power cord, and the other is the printer cable. Connecting a printer involves attaching the printer cable to the correct port on your computer's system unit, then plugging the power cord into an electrical outlet.

To attach a printer cable to the system unit, follow these steps:

1. Locate the printer cable. The easiest way is to find the one that isn't a power cord. It is probably heavy, gray, about six feet long, and it ends in a D-shaped connector with 25 pins.



NOTE: For laser and inkjet printers, the printer cable is probably permanently attached to the printer. For dot-matrix printers, you will have to manually attach one end of the cable to the printer. Locate the end that *is not* the D-shaped 25-pin connector and attach it to the port on the printer. Secure it using two clips that snap into place.

- 
2. Locate the parallel port on the system unit. The parallel port also has 25 sockets, arranged in the same configuration as the 25 pins on the cable connector.
 3. Plug the connector on the cable into the port on the system unit.
 4. Secure the connector by tightening the screws on either side.

Understanding Your Printer Controls

Although most commands that affect your printer come from your PC and your PC applications, most printers also have a set of controls. For example, you may be able to select the mode of operation as well as which paper tray to use for a particular print job. Status lights on the front of the printer let you know which controls are active.

If you look at the front of your printer, these are some of the controls you are likely to find:

- *Online or ready light.* A printer must be online in order to work, which means it must be active and ready to accept data from the computer.

NOTE: Although some printers automatically come online when you print a file, most have an on/off switch. If none of the status lights on your computer are lit, locate the on/off switch and turn the printer on.

-
- *Mode indicator.* Dot-matrix printers are likely to have different modes such as draft mode, near-letter-quality mode, and letter-quality mode.
 - *Error indicator.* If there is a problem with printing such as no paper, or a paper jam, the error light will be lit.

Be sure to read the documentation that comes with your printer for complete instructions on using the printer controls.

Installing Your Printer to Work in Windows 95

Windows 95 uses the Add Printer Wizard to assist you in installing a printer. If your printer is not already installed, run the wizard to add the printer.

NOTE: These instructions are for installing a local printer, one that is attached directly to your computer. If you have a networked printer, ask your network administrator for information about installing it for use with your computer.

1. Double-click My Computer and then double-click the Printers folder icon. Double-click the Add Printer icon to start the Add Printer Wizard. Choose Next.
2. On the second screen of the wizard, select the Local Printer option button to install a printer directly attached to your computer. Then click Next again.
3. Select the printer's manufacturer in the Manufacturers list box. Select the specific printer in the Printers list box and then choose Next.

NOTE: If your printer is not on the list, either choose the Generic manufacturer and Generic/Text Only printer, or click the Have Disk button and follow instructions to install a vendor-supplied driver.

4. Choose the printer port, typically LPT1, and then click the Configure Port button. Check the Spool MS-DOS Print Jobs and the Check Port State Before Printing check boxes. Choose OK and then choose Next.
5. In the Printer Name text box, type a name for the printer or keep the name that is displayed. Select the Yes option button if you want Windows-based programs to use this printer as the default printer. (If a different

printer is the default printer, choose No.) Choose Next.

6. Select Yes if you want to print a test page, and then choose Finish. The test page prints (if you selected Yes). The wizard copies the printer drivers to your system, prompting for the Windows 95 disks or CD if needed.

After you install a printer in Windows, you can make changes to the configuration to customize it for different printing requirements. You make these changes in the printer's Properties dialog box:

1. Click the Start menu, select Settings, Printers, and then select the printer icon for the printer you want to configure.

2. Right-click the printer icon and then select Properties to open the Properties dialog box

Windows 95 makes it easy to install a new printer by using the Add Printer Wizard. Just pick the right make and model from the list.

3. The Properties dialog box contains several tabs. Select each tab to view the settings. Click the Help button (the question mark in the upper-right corner) and then click a feature to read details about that property. Properties vary according to each printer's capabilities.

4. Change settings as desired, and then choose OK to save the new settings. Choose Cancel to get out of the dialog box without making any changes.

If you no longer need a printer that is installed in Windows 95, you can delete it from the Printers folder window:

1. Double-click My Computer and then double-click the Printers folder icon.

2. Select the printer you want to delete. Open the File menu and select Delete.

3. Windows asks if you are sure that you want to delete the printer. Choose Yes. The printer icon is deleted.

4. Windows then asks if it can remove files that were used only for this printer. Choose Yes. (If you plan to reattach this printer in the future, choose No.)

Other Printing Considerations

Printing a document is similar in most Windows applications. You can open the application's **F**ile menu and select **P**rint to open the Print dialog box. In the dialog box, you can select options such as which pages you want to print and how many copies to print. To quickly print the current document without setting options, click the Print button on the toolbar.

A printer icon on the taskbar near the clock indicates that printing is in progress. A quick way to check the status of print jobs is to point to the printer icon on the taskbar.

After you issue the print command, the Print Manager takes control of the *print job*. You can use the commands in Print Manager to check the status of a print job, to cancel a print job, or to pause a print job.

Print Manager keeps track of documents waiting to be printed. There are several ways to open Print Manager:

- If a print job is in progress, double-click the printer icon in the taskbar.
- Click the Start button and select **S**ettings, **P**rinters. Then double-click the printer icon for the printer you want to manage.
- Double-click the My Computer icon and then double-click the Printers folder icon.
Double-click the printer icon for the printer you want to manage.

If your printer is not working, here are a few things to check:

- Is the printer turned on?
- Is the printer online?
- Are the cables between the printer and the computer attached correctly and securely?

- Is the printer power cord plugged in?
- Is there paper in the printer?
- Has the printer been paused in the Print Manager?

TIP: To check if a printer has been paused, open the Print Manager, then open the Printer menu. If there is a check mark next to the Pause Printing command, the printer has been paused. To restart it, select the Pause Printing command.

RAM Guide

Due to cost considerations, all but the very high-end (and very expensive) computers have utilized DRAM for main memory. Originally, these were asynchronous, single-bank designs because the processors were relatively slow. Most recently, synchronous interfaces have been produced with many advanced features. Though these high-performance DRAMs have been available for only a few years, it is apparent that they will soon be replaced by at least one of the protocol-based designs, such as SyncLink or the DRDRAM design from Rambus, Inc. and Intel.

Basic DRAM operation

A DRAM memory array can be thought of as a table of cells. These cells are comprised of capacitors, and contain one or more 'bits' of data, depending upon the chip configuration. This table is addressed via row and column decoders, which in turn receive their signals from the RAS and CAS clock generators. In order to minimize the package size, the row and column addresses are multiplexed into row and column address buffers. For example, if there are 11 address lines, there will be 11 row and 11 column address buffers. Access transistors called 'sense amps' are connected to the each column and provide the read and restore operations of the chip. Since the cells are capacitors that discharge for each read operation, the sense amp must restore the data before the end of the access cycle.

The capacitors used for data cells tend to bleed off their charge, and therefore require a periodic refresh cycle or data will be lost. A refresh controller determines the time between refresh cycles, and a refresh counter ensures that the entire array (all rows) are refreshed. Of course, this means that some cycles are used for refresh operations, and has some impact on performance.

A typical memory access would occur as follows. First, the row address bits are placed onto the address pins. After a period of time the RAS\ signal falls, which activates the sense amps and causes the row address to be latched into the row address buffer. When the RAS\ signal stabilizes, the selected row is transferred onto the sense amps. Next, the column address bits are set up, and then latched

into the column address buffer when CAS\ falls, at which time the output buffer is also turned on. When CAS\ stabilizes, the selected sense amp feeds its data onto the output buffer.

Page Mode Access

By implementing special access modes, designers were able to eliminate some of the internal operations for certain types of access. The first significant implementation was called Page Mode access.

Using this method, the RAS\ signal is held active so that an entire 'page' of data is held on the sense amps. New column addresses can then be repeatedly clocked in only by cycling CAS\. This provides much faster random access reads, since the row address setup and hold times are eliminated.

While some applications benefit greatly from this type of access, there are others that do not benefit at all. The original Page Mode was improved upon and replaced very quickly so you will likely never see any memory of this type. Even if you do, it wouldn't be worth even getting it for free, considering the advantages of later access modes.

Fast Page Mode

Fast Page mode improved upon the original page mode by eliminating the column address setup time during the page cycle. This was accomplished by activating the column address buffers on the falling edge of RAS\ (rather than CAS\). Since RAS\ remains low for the entire page cycle, this acts as a transparent latch when CAS\ is high, and allows address setup to occur as soon as the column address is valid, rather than waiting for CAS\ to fall.

Fast Page mode became the most widely used access method for DRAMs, and is still used on many systems. The benefit of FPM memory is reduced power consumption, mainly because sense and restore current is not necessary during page mode access. Though FPM was a major innovation, there are still some drawbacks. The most significant is that the output buffers turn off when CAS\ goes high. The minimum cycle time is 5ns before the output buffers turn off,

which essentially adds at least 5ns to the cycle time.

Today, FPM memory is the least desirable of all available DRAM memory. You should only consider using this if it is either free, or your system does not support any of the later memory types (such as a 486 based system). Typical timings are 6-3-3-3 (initial latency of 3 clocks, with a 3-clock page access). Due to the limited demand, FPM is actually more expensive now than most of the faster memories now available.

HyperPage Mode (EDO)

The last major improvement to asynchronous DRAMs came with the Hyperpage mode, or Extended DataOut. This innovation was simply to no longer turn off the output buffers upon the rising edge of /CAS. In essence, this eliminates the column precharge time while latching the data out. This allows the minimum time for /CAS to be low to be reduced, and the rising edge can come earlier.

In addition to a 40% or greater improvement in access times, EDO uses the same amount of silicon and the same package size. EDO has been shown to work well with memory bus speeds up to 83MHz with little or no performance penalty. If the chips are sufficiently fast (55ns or faster), EDO can be used even with a 100MHz memory bus. One of the best reasons to use EDO is that all of the current motherboard chipsets support it with no compatibility problems, unlike much of the synchronous memory now being used.

Even with all the stated advantages, EDO is no longer considered mainstream. Most manufacturers no longer produce it, or have limited production. It is only a matter of time before the prices begin to rise, and the equivalent size SDRAM module will be less expensive.

If you already own EDO memory, there is no real reason to jump to SDRAM unless you require bus speeds above 83MHz. With a typical EDO timing of 5-2-2-2 at 66MHz, there is almost no noticeable improvement with SDRAM over EDO, and at 83MHz it is still negligible. If you require 100MHz bus operation, EDO will lag far behind current SDRAM in performance even if it does operate at that speed due to the need for 6-3-3-3 timings. On the other hand, with EDO being phased out, you will likely find SDRAM to be equal to or even lower in price.

Burst EDO (BEDO)

Burst EDO, while a good idea, was dead before it ever was born. The addition of a burst mode, along with a dual bank architecture would have provided the 4-1-1-1 access times at 66MHz that many expected with SDRAM. Burst mode is an advancement over page mode, in that after the first address input, the next 3 addresses are generated internally, thereby eliminating the time necessary to input a new column address. Unfortunately, Intel decided that EDO was no longer viable, and SDRAM was their preferred memory architecture so they did not implement support of BEDO into their chipsets. In fact, several large memory manufacturers had put considerable time and money into the development of SDRAM over the past decade, and were not very happy with the BEDO design.

Except for support of bus speeds of 100MHz and faster, BEDO would probably have been a much faster and more stable memory than SDRAM. Essentially, BEDO lost support as much for political and economic reasons as for technical ones, it seems.

Synchronous Operation

Once it became apparent that bus speeds would need to run faster than 66MHz, DRAM designers needed to find a way to overcome the significant latency issues that still existed. By implementing a synchronous interface, they were able to do this and gain some additional advantages as well.

With an asynchronous interface, the processor must wait idly for the DRAM to complete its internal operations, which typically takes about 60ns. With synchronous control, the DRAM latches information from the processor under control of the system clock. These latches store the addresses, data and control signals, which allows the processor to handle other tasks. After a specific number of clock cycles the data becomes available and the processor can read it from the output lines.

Another advantage of a synchronous interface is that the system clock is the only timing edge that needs to be provided to the DRAM. This eliminates the need for

multiple timing strobes to be propagated. The inputs are simplified as well, since the control signals, addresses and data can all be latched in without the processor monitoring setup and hold timings. Similar benefits are realized for output operations as well.

JEDEC SDRAM

All DRAMs that have a synchronous interface are known generically as SDRAM. This includes CDRAM (Cache DRAM), RDRAM (Rambus DRAM), ESDRAM (Enhanced SDRAM) and others, however the type that most often is called SDRAM is the JEDEC standard synchronous DRAM.

JEDEC SDRAM not only has a synchronous interface controlled by the system clock, it also includes a dual-bank architecture and burst mode (1-bit, 2-bit, 4-bit, 8-bit and full page). A 'mode register' that can be set at power-on and changed during operation controls the burst mode, burst type (sequential or interleave), burst length and CAS latency (1, 2 or 3).

CAS Latency is one of several performance related timings for SDRAM. This measurement is the time it takes to strobe in the Row Address, and to activate the bank. When a burst read cycle is initiated, the addresses are set up and RAS\ and CS\ (chip select) are held low on the next clock cycle (rising edge of CLK), thereby activating the sense amplifiers on the bank. A period of time equal to tRCD (RAS\ to CAS\ delay) must pass after which CAS\ and CS\ are held low (again, at the next clock cycle). After the time period for tCAC (column access time) has passed the first bit of data is on the output line and can be retrieved (at the next clock cycle). The basic rule is that CAS latency times the clock speed (tCLK) must be equal or greater than tCAC (or $CL \times tCLK \geq tCAC$). This means that the column access time is the limiting factor for CAS Latency.

SDRAM was initially introduced as the answer to all performance problems, however it quickly became apparent that there was little performance benefit and a lot of compatibility problems. The first SDRAM modules contained only two clock lines, but it was soon determined that this was insufficient. This created two different module designs (2-clock and 4-clock), and you needed to know which your motherboard required. Though the timings were theoretically supposed to be 5-1-1-1 @ 66MHz, many of the original SDRAM would only run at 6-2-2-2 when run in pairs, mostly because the chipsets (i430VX, SiS5571) had

trouble with the speed and coordinating the accesses between modules. The i430TX chipset and later non-Intel chipsets improved upon this, and the SPD chip (serial presence detect) was added to the standard so chipsets could read the timings from the module. Unfortunately, for quite some time the SPD EEPROM was either not included on many modules, or not read by the motherboards.

SDRAM chips are officially rated in MHz, rather than nanoseconds (ns) so that there is a common denominator between the bus speed and the chip speed. This speed is determined by dividing 1 second (1 billion ns) by the output speed of the chip. For example a 67MHz SDRAM chip is rated as 15ns. Note that this nanosecond rating is **not** measuring the same timing as an asynchronous DRAM chip. Remember, internally all DRAM operates in a very similar manner, and most performance gains are achieved by 'hiding' the internal operations in various ways.

The original SDRAM modules either used 83MHz chips (12ns) or 100MHz chips (10ns), however these were only rated for 66MHz bus operation. Due to some of the delays introduced when having to deal with the various synchronization of signals, the 100MHz chips will produce a module that operates reliably at about 83MHz, in many cases. These SDRAM modules are now called PC66, to differentiate them from those conforming to Intel's PC100 specification.

PC100 SDRAM

When Intel decided to officially implement a 100MHz system bus speed, they understood that most of the SDRAM modules available at that time would not operate properly above 83MHz. In order to bring some semblance of order to the marketplace, Intel introduced the PC100 specification as a guideline to manufacturers for building modules that would function properly on their upcoming i440BX. With the PC100 specification, Intel laid out a number of guidelines for trace lengths, trace widths and spacing, number of PCB layers, EEPROM programming specs, etc.

There is still quite a bit of confusion regarding what a 'true' PC100 module actually consists of. Unfortunately, there are quite a few modules being sold today as PC100, yet do not operate reliably at 100MHz. While the chip speed rating is used most often to determine the overall performance of the chip, a

number of other timings are very important. tRCD (RAS to CAS Delay), tRP (RAS precharge time) and CAS Latency all play a role in determining the fastest bus speed the module will operate on to still achieve a 4-1-1-1 timing.

PC100 SDRAM on a 100MHz (or faster) system bus will provide a performance boost for Socket 7 systems of between 10% and 15%, since the L2 cache is running at system bus speed. Pentium II systems will not see as big a boost, because the L2 cache is running at $\frac{1}{2}$ processor speed anyway, with the exception of the cacheless Celeron chips of course.

DDR SDRAM

One limitation of JEDEC SDRAM is that the theoretical limitation of the design is 125MHz, though technology advances may allow up to 133MHz operation. It is obvious that bus speeds will need to increase well beyond that in order for memory bandwidth to keep up with future processors. There are several competing new standards on the horizon that are very promising, however most of them require special pinouts, smaller bus widths, or other design considerations. In the short term, Double Data Rate SDRAM looks very appealing. Essentially, this design allows the activation of output operations on the chip to occur on both the rising and falling edge of the clock. Currently, only the rising edge signals an event to occur, so the DDR SDRAM design can effectively double the speed of operation up to at least 200MHz.

There is already one Socket 7 chipset that has support for DDR SDRAM, and more will certainly follow if manufacturers decide to make this memory available. In this industry, many times it is the first to market that gains the support, rather than the best technology.

Enhanced SDRAM (ESDRAM)

In order to overcome some of the inherent latency problems with standard DRAM memory modules, several manufacturers have included a small amount of SRAM directly into the chip, effectively creating an on-chip cache. One such design that is gaining some attention is ESDRAM from Ramtron International Corporation.

ESDRAM is essentially SDRAM, plus a small amount of SRAM cache which allows for lower latency times and burst operations up to 200MHz. Just as with external cache memory, the goal of a cache DRAM is to hold the most frequently used data in the SRAM cache to minimize accesses to the slower DRAM. One advantage to the on-chip SRAM is that a wider bus can be used between the SRAM and DRAM, effectively increasing the bandwidth and increasing the speed of the DRAM even when there is a cache miss.

As with DDR SDRAM, there is currently at least one Socket 7 chipset with support for ESDRAM. The deciding factor in determining which of these solutions will succeed will likely be the initial cost of the modules. Current estimates show the cost of ESDRAM at about 4 times that of existing DRAM solutions, which will likely not go over well with most users.

Protocol Based DRAM

All of the previously discussed DRAM have separate address, data and control lines which limits the speed at which the device can operate with current technology. In order to overcome this limitation, several designs implement all of these signals on the same bus. The two protocol based designs currently getting the most attention are SyncLink DRAM (now called SLDRAM due to trademark issues) and Direct Rambus DRAM (DRDRAM) licensed by Rambus, Inc.

DRDRAM

Intel has placed their money on the proprietary memory design developed by Rambus, Inc. On the surface, this looks to be a very fast solution for system memory due to its fast operation (up to 800MHz). The reality is, however, that the design is only up to twice as fast as current SDRAM operation due to the smaller bus width (16 bits vs. 64 bits).

Despite the claims from Intel and Rambus, Inc., there are some potentially serious issues which need to be addressed with this technology. The higher speeds require short wire lengths and additional shielding to prevent problems with EMI. In addition, latency times are actually worse than currently available

fast SDRAM. Since most of today's applications do not actually utilize the full bandwidth of the memory bus even today, simply increasing the bandwidth while ignoring latency issues will likely not provide any real performance improvements. In addition, processors operating with 800MHz bus speeds will certainly require more than double the current memory bandwidth.

While these issues are serious enough, the biggest drawback to the technology is that it is proprietary technology. Manufacturers wishing to implement a solution with DRDRAM will be required to pay a royalty to Intel and Rambus, Inc., and will also have no real control over the technology. This is not an attractive outlook for most memory manufacturers who have no desire to essentially become chip foundries.

SLDRAM

Many memory manufacturers are putting their support behind SLDRAM as the long-term solution for system performance. While SLDRAM is a protocol-based design, just as RDRAM is, it is an open-industry-standard, which requires no royalty payments. This alone should allow for lower cost. Another cost advantage for the SLDRAM design is that it does not require a redesign of the RAM chips.

Due to the use of packets for address, data and control signals, SLDRAM can operate on a faster bus than standard SDRAM – up to at least 200MHz. Just as DDR SDRAM operates the output signal at twice the clock rate, so can SLDRAM. This puts the output operation as high as 400MHz, with some engineers claiming it can reach 800MHz in the near future.

Compared to DRDRAM, it seems that SLDRAM is a much better solution due to the lower actual clock speed (reducing signal problems), lower latency timings and lower cost due to the royalty-free design and operation on current bus designs. It appears that even the bandwidth of SLDRAM is much higher than DRDRAM at 3.2GB/s vs. 1.6GB/s

Though Intel initially intended to support only DRDRAM in future chipsets, competing chipset manufacturers, memory manufacturers and pressure from end users may force them to include support for SLDRAM as well. If the marketplace can successfully influence Intel to provide this support, we may

actually see a situation where the best technology wins over marketing hype.

Scanners

Defining Scanner Types

Scanners are peripheral devices used to *digitize* (convert to electronic format) artwork, photographs, text, or other items from hard copy. In a sense, a scanner works as a pair of eyes for your PC. Your eyes see an image and translate the image into electrical impulses that travel to and are interpreted by your brain. Similarly, a scanner captures images and converts them to digital data that travel to and are interpreted by the computer.

A scanner works by dividing an image into microscopic rows and columns and measuring, like the film in a camera, how much light (or lack thereof) reflects from each individual intersection of the rows and columns. Each reflection is recorded as a dot, or picture element (pixel). After the scanner collects information from each dot, it compiles the result into a digital file on the computer.

There are a wide variety of scanners that work in a number of different ways, but the technology behind them is essentially the same. The following sections discuss the more popular types of scanners available today.

ON THE WEB: Two excellent Internet resources for more information on scanners and scanning techniques can be found at:

<http://www.scanshop.com> <http://www.cis.ohio-state.edu/~sking/Scanners.html>

Flatbed Scanners

Flatbed scanners look and behave a lot like a photocopier. You lay the item to be scanned on a glass plate and the scanning head passes below the glass.

Flatbed scanners are very versatile: you can scan objects in a variety of sizes and shapes, including pages from a book, without damaging the original. While

flatbed scanners are the best choice for a wide variety of uses, if you plan to do a lot of text scanning (called *OCR* for *Optical Character Recognition*) work, keep in mind that flatbeds only accommodate one page at a time. Scanning multi-page documents can be a slow, tedious process, because you have to manually remove one page and insert the next.

Sheetfed Scanners

Sheetfed scanners look and act more like fax machines. The page or item is fed into the machine, scanned, then spit out on the other end. A sheetfed scanner is a good choice for large volumes of text, but not for handling delicate original photographs. Scanning directly from a book or other three-dimensional object is impossible.

Hand Scanners

Hand scanners are a low-cost alternative to their larger, more sophisticated cousins. As their name implies, hand scanners are manual devices you move over a flat surface, just as you do your PC's mouse

The hand scanner's advantages are many, but so are its disadvantages. Generally, hand scanners work best for small, uncomplicated images such as company logos or small black-and-white photographs. You might want a hand scanner if you don't plan to use it on a regular basis, because it usually doesn't require adding internal cards to your CPU, and it's easily disconnected and stored away. Most hand scanners can only scan a four-inch wide image at one time and require a steady hand. You're usually provided with software that helps you "sew up" a series of these 4-inch, side-by-side scans into one image, but this is obviously not as convenient as getting the full image at once.

NOTE: Other types of scanners also exist, such as barcode readers (like the ones at the grocery checkout) and slide scanners; however, because these are specialty items primarily used for commercial purposes, they are not covered here.

Color versus Grayscale Scanners

Scanners that can scan images in full color have become much more popular as their prices have dropped. Just a few years ago, color scanners cost several thousands of dollars, but can now be bought for a few hundred, depending on resolution and type.

Even so, grayscale (meaning shades of black and white only, no color) scanners are still available and are significantly cheaper. In many cases, they are perfectly adequate for the average user. Unless you have a color printer, or use your scanner to create artwork that will only be viewed on-screen (such as for a Web page), there's no point in having a color scanner. Consider this carefully before buying; however, what you think you'll never do now could change as you grow more experienced and interested in computer technology.

Understanding Resolution and Dots Per Inch

Like printers, the technical capability (*optical resolution*) of a scanner is measured in *dots per inch (dpi)*. The higher the dpi, the sharper your on-screen printable image will be. The reason for this is that the more dots that can be placed in the same area (in this case, one square inch), the smoother and more detailed the overall image will look.

On the other hand, the higher its resolution, the longer the scanner takes to scan the image, and the larger the resulting image's data file will be (more dots equals more information; more information equals a larger file). Typical flatbed scanners range from 300 dpi to more than 3,200 dpi; The most affordable resolution for the average consumer is 1,200-1,600 dpi.

NOTE: Don't be fooled by ads touting astronomical resolutions like 6,400 dpi interpolated or enhanced resolution. *Interpolated resolution* (as opposed to optical or true resolution) is just another term for the scanner's ability to guess at how the image would look at a higher resolution by adding new dots in between the old ones. For most uses, this process adds no value--and can often actually reduce the scan quality. The exception is that interpolated resolution can improve the quality of scanned line art images that will be output at very high resolution, such as magazines or commercially

produced documents.



System Requirements

Before you buy or try to install a scanner, make sure that your computer can actually run it. Scanners can really stretch the limits of a mid-range PC--and may not work at all on low-end ones. Carefully review the manufacturer's stated system requirements and verify that your system has at least the minimum amount of RAM, CPU speed, and hard drive space listed. For best results, *exceed* the minimum requirements by as much as possible. Most manufacturers require nothing less than a 486/33 processor (if you have a Pentium or better, you should be safe). Most experts recommend a minimum of 16-24M of RAM even if the manufacturer claims less.

Connecting a Scanner to Your PC

Like all peripheral devices, your scanner must be connected to your computer in some way. Most likely, your scanner came with all the equipment required to hook it up.

SCSI Cards and Cables

The vast majority of flatbed and sheetfed scanners use SCSI (Small Computer System Interface, pronounced "skuzzy") connections. If you're lucky, your computer already has a SCSI card installed and all you'll have to do is connect the provided cable to the SCSI port at the back of your computer. But most computers don't come that way, so you might need to open the system case and install a new card in one of your computer's expansion slots. If so, follow the instructions that came with the scanner, and check your computer system's manual to make sure you're installing it in the right place. See Chapter 27, "Upgrading Your Hardware," for details on how to install add-in cards.

NOTE: Hand scanners typically do not require SCSI connections; They usually connect to parallel ports, like the one used by your printer.

The great thing about SCSI devices is that one SCSI card can support up to seven different devices connected together in a "SCSI chain." You might, for example, have a scanner plugged into the SCSI port, a Jaz drive connected to the scanner, an external hard drive connected to the Jaz drive, and a CD-ROM drive connected to the external hard drive. Not only does this give you more flexibility in configuring your system, it also saves the few expansion ports in your computer for other devices that aren't able to share.

CAUTION: Be careful when comparing scanners and the SCSI cards included with them. Some manufacturers cut corners (and offer lower prices) by including a *proprietary* SCSI card, or one that only works with that particular scanner and does not support other devices. The card will

work, of course, but it's not a very efficient use of the technology. If you think the scanner is right for you but it comes with a proprietary card, you can always buy a general-purpose SCSI card separately.

NOTE: Most scanners and SCSI devices require a cable length of 12 feet or less--keep this in mind when you're setting up your system.

TWAIN Drivers

Every hardware device, whether internal or external, requires a software program called a *driver* to help it communicate with your computer's operating system. For scanners, the standard driver type is a TWAIN driver (supposedly, *TWAIN* stands for *Technology Without An Interesting Name*). By using TWAIN drivers, scanners can communicate with any TWAIN-compatible applications. Most of today's popular drawing and graphics applications such as CorelDRAW!, Paint Shop Pro, Adobe Photoshop, and even Microsoft Office 97 all support TWAIN, and thus use very similar steps, buttons, and the like. You don't have to learn a different method for every application you scan from. (If you're not sure, check the application's File menu for an item labeled Acquire, or look in the manual.)

Getting Windows 95 to Recognize Your Scanner

After you've made the physical connections, install the TWAIN driver according to the instructions that came with the scanner. If instructions for using the scanner under Windows 95 are not included, follow these steps:

1. Shut down the computer and wait 30 seconds. Power up the scanner, wait 10-15 seconds, then reboot your PC.
2. From the Start menu, choose Settings, Control Panel. When the Control Panel window opens, double-click Add New Hardware.
3. Click Next. The wizard will ask if you want Windows 95 to search for your new hardware; choose No and click Next again

- 4.** At the next screen, click Other Devices from the list of items click Next again.
- 5.** Look through the Manufacturers list and the Models list. If you find your scanner (you probably won't), click Next. If not, click Have Disk.
- 6.** In the Install from Disk dialog box, show the wizard where your install program is (usually on a floppy disk that came with the scanner), either by typing in the path name or clicking the Browse button and navigating to the proper location. Click Next twice.
- 7.** When you get to the final wizard screen, click Finish. Windows will continue the setup process, and may ask you for your Windows 95 CD-ROM.
- 8.** When the setup process is complete, restart your computer by choosing Shut Down from the Start menu, then choosing Restart the Computer?.
- 9.** As Windows 95 restarts, you may notice a message on the screen indicating it has found a new device and is installing drivers for it, or other similar messages. Respond and follow any directions if necessary.
- 10.** To confirm that Windows has found and set up your scanner, double-click the System icon in the Control Panel folder.
- 11.** Click the Device Manager tab, and look for the Other devices entry on the list of devices. Click the + next to Other devices; your scanner should be listed. Click OK to clear the dialog box. Your scanner is now ready to use.

Scanning a Document

Most scanners come with some kind of image-editing software such as Adobe Photoshop or Photoshop LE. Even if not, there are some excellent programs available for a relatively inexpensive price--some you can even download from the Internet, such as Paint Shop Pro or Graphics Workshop.

If the program is TWAIN-compliant, its scanning features will work about the same no matter which program you choose. The main differences will come from the TWAIN driver supplied by your scanner's manufacturer. You'll want to review the program's documentation or check its online Help system to learn exactly how to use that particular program.

Regardless of the program, the general tips and techniques that follow will help you produce better scans and get the most out of your image-editing software:

- Always keep the glass on your flatbed scanner clean and free of smudges and fingerprints. Be sure to use special lens-cleaning cloths, not household tissues or paper towels--these can permanently scratch your scanner's glass. You can get lens-cleaning cloths at most computer or office supply stores.
- If you have to move your flatbed scanner, make sure you lock the scanning head in place first. If you don't, it can become dislodged from its track and be costly to repair. If the locking mechanism isn't obvious, check the scanner's documentation for help.
- Know the desired result before you scan. How will the image be used--for a letterhead? On a Web page? In a newsletter? Different jobs require different approaches, so plan ahead.
- Use the lowest resolution possible that will still provide the quality you want. "Overscanning" only serves to make larger files and slower scans.
- Scan photographs at about twice the resolution of the expected output. If you're not sure what that will be, ask your print shop. Magazines are traditionally printed at 133 lines per inch (lpi), so scan the photo at 266

dpi.

- For line art images (hand-drawn artwork, mechanical drawings, and so on), always use the highest resolution possible.
- Don't overscan an image if it's only going to a laser or inkjet printer. Laser printers generally output at about 50 or 100 lpi, so a simple 200 dpi scan will do just fine.
- Images intended only for on-screen viewing need even less: 80 to 100 dpi is adequate.
- For OCR (see the next section), use 300-400 dpi.
- Get to know your image-editing software. Experiment with the various filters and effects you can apply to your scanned images--they can make all the difference in the world. If you're really interested in mastering image editing or you need to learn a lot quickly, consider taking a class or buying a book especially on that software program. *Inside Adobe Photoshop* from New Riders Publishing would be one recommendation. An edition of the book is available for each of the two current releases (3 and 4).

OCR Software

Using OCR, scanners can also be used to convert hardcopy text to a text file you can then edit in your word processor, saving hours of retyping. In addition to stand-alone programs like IBM's TextBridge, many of the better fax/modem software programs such as Delrina WinFax PRO also include OCR capability.

Keep in mind that while most OCR programs have a 95 percent or higher accuracy rate, they are not perfect. For best results, you need to start with a clean hard copy (preferably on white paper), in a simple typeface, such as that produced by a typewriter. Handwritten words, unless very neatly printed, usually don't translate very well. If your original is messy, or the print is too small or fancy for the scanner to read easily, you might end up spending as much time cleaning up the scanned text as you would retyping it after all.

Speakers and Sound

Sound is one of the favorite options on today's PCs. In fact, sound is a standard feature of most new PCs.

Everyone has their own uses for sound. You may just want to play a few music CDs with your CD-ROM drive while you are working, or you may use multimedia applications for presentations or educational programs. You may just like the sound of your jet engines roaring as you punch the throttle in a flight simulator. This chapter explains the basics of sound in Windows and shows you how to get the most sound out of your PC.

Components for Sound on Your PC

To produce sound on your PC, you need a sound card and speakers. The *sound card* is an expansion card that plugs into one of the slots on your motherboard. This card processes all of the instructions that have to do with sound, and then sends the sounds to the speakers to be played. The speakers plug into the sound card and usually have a power plug as well.

TIP: Many motherboard manufacturers include sound chip(s) on the motherboard itself. If you have a new PC, check the specifications of its make and model to see if it uses this technology.

Sound Cards

Sound cards plug into an expansion slot in your PC. The card has a set of connectors that are exposed on the back of the PC; you can plug a microphone and speakers into the connectors. Some cards include an optional connector for a *line input*, which is used to connect another input source such as the output from a tape player or portable CD player. Other optional connectors can include a joystick connector and a MIDI connector (for connecting a MIDI keyboard). The card may include a volume control, which controls the volume of the speakers and/or headphones.

Speakers

The built-in speakers in most PC cases are used just for making system sounds, such as warning beeps and action indicators. To play more sophisticated sounds on your PC, you need a set of external speakers. Usually speakers come in pairs, and there is a plug that connects them to your sound card. Arrange the speakers with one on the left and one on the right of your desk or work area to get a stereo effect.

Optionally, some speakers come with a *subwoofer*. This improves the bass (low

notes) sound. If you have a subwoofer with your speakers, it should go on the floor under your desk.



CAUTION: Be sure your speakers are shielded. Otherwise, if they get too close to the monitor, they will distort the display. Even with shielded speakers, take care not to get floppy disks too close to them as the magnets in the speakers can damage data on the disks.



Playing Sounds for Windows Events

By default, Windows 95 plays a sound when it starts and when it closes. You can change the sounds that are played for these events, and you can assign sounds to be played for other events. To do this, follow these steps:

1. Click the Start menu, choose Settings, and then choose Control Panel.
2. Double-click the Sounds icon in the Control Panel. This opens the Sounds Properties dialog box
3. To assign a sound to an event, scroll the list at the top of the dialog box to find the event.

NOTE: Some programs add additional events that you can assign sounds to. These will appear at the bottom of the list.

4. Click the Name box and choose a sound from the list. If you don't see a sound you want to use, click the Browse button and find any sound file in the WAV format.
5. To test the sound (to be sure it is the right one), click the Play button next to the Preview box.
6. Assign as many sounds as you want to different events, then click OK when you are finished.

Playing WAV Sound Files on Your PC

You're not limited to the built-in event sounds that came with your operating system; you can play other types of sound files as well. WAV files are the most common type of PC sound file. Windows has a built-in player, called Sound Recorder, for playing WAV files. To play a WAV file, use Explorer or My Computer to find the file you want to play, and double-click it to start it. Sound Recorder starts, and the WAV file begins playing immediately. When the file is finished playing, Sound Recorder closes automatically.

If you want to stop the playback, rewind it, and play it again, click the Stop button before it finishes playing. Once you click stop, you can play and rewind it as much as you want and the player won't close when the sound is finished playing.

You'll see how to use this same application to record your own sounds later in the section "Recording Sounds in Windows 95."

Getting New Sounds

If you are really into sounds and want more on your PC, there are several ways to get more sound files. The Windows CD-ROM installation includes many additional sound schemes, which you can add by using the Add/Remove Programs option in the Control Panel.

You can also buy any of a number of CD-ROMs that contain whole collections of additional sounds in the WAV format. And, you can find thousands of sound clips on the Internet.

ON THE WEB: Use your Web browser to go to the following site for a good starting list of sites with WAV sounds:

http://www.yahoo.com/Computers_and_Internet/Multimedia/Sound/Arch

Playing Music on Your PC

With a CD-ROM, sound card, and speakers, you can use your \$2,000 PC like a \$100 CD player. The sound probably won't be as good as your home stereo, but it is more convenient than setting up a stereo with your PC.

To use your computer to play audio CDs, put a CD into the CD-ROM drive. If your computer has AutoPlay features (as do all Windows 95 systems), it will automatically detect the CD and start the Windows CD player.

NOTE: If you installed Windows from floppy disks, you won't have the Windows CD player installed. If you installed Windows from a CD but didn't install the CD player, you need to use the Add/Remove Programs option in the Control Panel to add the CD Player.

Adjusting the Speaker Volume and Tone in Windows 95

There are three ways to adjust the volume to your speakers. First, you can just adjust the volume control on the speakers themselves. See the instruction manual that accompanied your speakers to find the specific location of and adjustment instructions for the volume controls on your speakers.

Within Windows, you can also adjust the speaker volume by clicking the speaker icon in the taskbar to bring up the volume control. Drag the slider up or down to adjust the volume, or click the Mute box to turn the speakers off all together. This Mute option is what you want to use if you are listening to a CD on headphones and don't want it to play on the speakers.

If you want more control over the speaker output, click the speaker icon in the taskbar with the right mouse button and select Volume Controls from the menu. This opens the Volume Control dialog box

The control at the far left is the master volume control. It controls the overall volume going to the speakers. Above it, you can adjust the balance to the left and right speakers.

Because the sound in your PC can come from several different sources, the Volume Control dialog box also provides controls for adjusting the relative volume of each of those sources. These sources are:

- *Wave*. Used to control the volume of files in the WAV format.
- *MIDI*. Used to control the volume of synthetic MIDI sounds, often used in games.
- *CD*. This is a separate control for the audio CD's volume or volume of audio CD sounds in games.
- *Line In*. If you have another external audio device plugged into the line in jack on your sound card, this controls its volume.

From the Volume Control dialog box, you can also control the bass and treble settings for your speakers. This adjustment is made for all sound, though, and doesn't differentiate between different sound sources. To adjust bass and treble settings, follow these steps:

1. Select the Options menu in the Volume Control dialog box and select the Advanced Controls item in the menu.
2. This adds an Advanced button to the Volume Control dialog box. Click it.
3. This opens the dialog box Adjust the Bass and Treble settings using the sliders, then click OK.

Listening to Your CDs on Headphones

Most CD-ROM drives have a headphone built-in jack and a volume control on the drive itself. To listen to a CD on the headphones instead of playing it over the speakers, follow these steps:

1. Put a CD into the CD-ROM drive.
2. Plug in the headphones.
3. Click the speaker icon in the taskbar and select the **M**ute check box to mute the speakers (otherwise, sound will issue from both the PC's speakers and the headphones).
4. Click the Play button in Windows' CD Player application to start the CD.
5. Adjust the volume using the control on the CD-ROM drive, and put on the headphones.

If your CD player doesn't have a volume control dial on it, you can use Windows to control the volume to the headphones. Follow these steps:

1. Click the speaker icon in the Windows taskbar; a Volume control box with slider pops up.
2. Click with the left mouse button on the slider, and drag the slider up to raise the sound level, or down to lower the sound level.
3. Click anywhere outside the control box to close it.

TIP: You can also plug a set of headphones into the Speaker jack on the sound card and listen to all of your sounds from CDs, games, and Windows on headphones instead of speakers. If you do that, all of the Windows speaker volume controls discussed in this chapter will control the volume for the headphones. In that case, be sure not to mute the speaker volume, or you won't hear anything on the headphones.



Microphones

Your computer may have come with a microphone. The microphone plugs into the sound card and is used to record your voice or other sounds. Another new use of a microphone is to make telephone-like conference calls over the Internet. Like any conference call, when you speak into the mic, the person you call hears you on his speakers.

Recording Sounds in Windows 95

The most common way to record a sound in Windows is by using a microphone. To do this, be sure the microphone is plugged into the microphone input jack on your sound card. If your mic has an on/off switch, be sure it is on. Then follow these steps:

1. Click the Start menu, and choose Programs, Accessories, Multimedia, Sound Recorder.
2. This opens the Sound Recorder. Click the Record button and begin speaking, singing, or making noise into the mic. When you are finished, click the Stop button. You can click Record again to record more at the end of the file.
3. When you are finished recording, select the File menu and choose Save. Find a directory to save the file in, then give it a name and click Save.
4. To hear how the recorded sound sounds, click the Play button.

You can also easily record a sound clip from a CD. To do this, follow these steps:

1. Click the Start menu, and choose Programs, Accessories, Multimedia, Sound Recorder to start the Sound Recorder.
2. Insert an audio CD into the CD to start the CD Player as discussed in the section "Playing Music on Your PC" earlier in this chapter.
3. Find the right song and the right point in the CD that you want to record.

Then click Pause in the CD Player.

4. Click the Record button in the Audio Recorder.

5. Click Pause again in the CD Player to start the playback. You should see the sound meter showing the recording in the Audio Recorder

6. When you're finished, click Stop in both the recorder and the player and save the file in the recorder.




NOTE: Recording sounds from CDs on your PC is a legal gray area. If you are just using a few clips for your own enjoyment on your PC, you won't get arrested or fined. But sharing these clips with others by trying to sell or distribute them, or by putting them on a Web site is a definite copyright infringement.




Adjusting Microphone and Input Levels in Windows 95

To control the level of incoming sounds from the microphone and other devices for recording, follow these steps:

1. Click the speaker icon in the taskbar with the right mouse button and select Volume Controls from the menu.
2. Select the Options menu in the Volume Control dialog box and choose Properties from the menu.
3. This opens a Properties dialog box . Be sure the Microphone option is selected at the bottom of the dialog box.
4. Select the Recording option in the Adjust Volume For section and click OK.
5. Now the controls that you see are for devices you can record from. Adjust the volumes for any of these. Close the Recording Control dialog box when you are finished.



TIP: If you leave the Recording Control open while you are recording, you can continue to adjust the controls as needed while you record.



Troubleshooting Sound Problems in Windows 95

Sound problems are common in Windows 95 and here are some common problems and solutions.

No Sound from the Speakers

Check to be sure that the speakers are plugged into the correct speaker output jack on the sound card and that the speakers are plugged into the power source or have live batteries.

If there are volume controls and an on/off switch on the speakers, be sure the controls are turned up and the speakers are turned on.

Does the sound card have a volume control dial where the speakers plug into it? If so, are these volume controls turned up?

Check the volume control to be sure that the Windows volume to the speakers isn't turned down too low and that the speakers aren't muted.

Check the volume controls as shown to be sure the volume isn't turned down or off there.

No Sound from Speakers When an Audio CD Is Playing

Check the volume controls as shown to be sure the CD volume isn't turned down or off.

There is a cable inside your PC that connects the CD-ROM audio to the sound card. If your PC wasn't assembled professionally or you added the CD-ROM or sound card yourself, check to be sure that this cable is there and connected properly.

Can't Record with the Microphone

Check to be sure that the microphone is connected to the mic input jack on the sound card. It should not be connected to line in.

If there is an on/off switch on the mic, check that it is in the On position.

Check the recording level controls in to be sure your microphone recording level is on.

Background Music in Games and Multimedia Doesn't Play

If all of your other sounds play okay, the problem may be with your MIDI volume. MIDI is a type of synthesized (rather than recorded) music used in a lot of games and multimedia. Check the settings to be sure the MIDI volume isn't down or off. Raise the volume for that if needed.

Many games and multimedia programs also use CD audio mixed in with the other sounds. If these seem to be missing, run through the checklist of solutions for the previous problem, "No Sound from Speakers When an Audio CD Is Playing."

Background Music in Games and Multimedia Sounds More Like a Kazoo than a Symphony

There are two different kinds of sound technologies for playing synthetic MIDI sounds. One is called *FM synthesis* and is used in older and less expensive sound cards. It tends to make all of the instruments sound the same. The better technology is called *wavetable synthesis*. If you are really particular about the quality of your sound in games and multimedia, you may want to upgrade to a wavetable sound card or see if a wavetable add-on is available for your current card.

Is your \$200 sound card connected to \$5 speakers? A better set of speakers may be the answer.

Getting Help with Hardware

Consider the components found in a PC, among them the power supply and fan, motherboard, memory, hard drives and floppy drives, video card, modem, sound card, CD-ROM, parallel and serial ports, and drive controller. Then there are the peripherals such as the keyboard, mouse, monitor, printer, scanner, and joystick, to name a few. With so many components, there are ample opportunities for problems to occur.

Simple Solutions First

At one time or another, most of us have been stumped by an apparently catastrophic computer glitch, only to have another user end the crisis with a few clicks of the keys or the throw of a switch. Even experienced users sometimes overlook the simple solutions and immediately begin taking more drastic measures. Any user--novice or expert--is better off starting with basic diagnostic procedures that take only a minute or two before going on to more exotic efforts. Remember two simple solutions before you begin dissecting your computer:

- Many hardware problems are the result of incorrect setup or settings.
- Sometimes a single-click or a reboot will solve a seemingly catastrophic problem.

The following sections discuss each of these approaches in more detail.

CAUTION: Do not open up your computer's power supply or monitor. Both can store enough energy to hurt or even kill you. Besides, there's little the average user can do to repair either of these components. Power supplies are inexpensive and normally are replaced when they fail. If you have a problem with your monitor, and reading the manual doesn't suggest a solution and the cables are firmly seated, take it to an authorized repair facility.

Getting the Settings Right

The best way to prevent annoying hardware problems from developing is to be sure the hardware is set up properly when you install it. While serious hardware problems--those where the hardware is damaged and in need of repair or replacement--can crop up occasionally, most hardware problems are related to improper settings that don't allow components to work together smoothly. These problems often occur when you make changes to the system's hardware. If you incorrectly install a new component or change a hardware setting, your well-behaved computer can become a monster.

A couple of quick tips for avoiding newly installed hardware problems:

- Read up on the system or component before you buy it. In particular, read the product reviews in computer magazines. These are good places to find out about a system's quirks or nonstandard design that can lead to hardware conflicts down the road. You can also find these product reviews on the Internet at such sites as ZDNet (<http://www.zdnet.com>), C|Net (<http://www.cnet.com>), and TechWeb (<http://www.techweb.com>).
- If you are installing a new component, read the manual first. Reading manuals on motherboards, add-in boards, or CD-ROMs may be like reading a foreign language, but they're better than hours of guessing. The more you familiarize yourself with the terminology in these manuals, the easier it is to diagnose and solve problems as they arise.

TIP: Experienced users save the boxes and packing materials for their equipment, in case it has to be shipped back to the manufacturer.

Such components as motherboards, hard drives, and CD-ROMs use jumpers and DIP switches that must be set correctly to work with the other equipment in your system. On a new computer the manufacturer will set the jumpers, although sometimes an incorrect setting may be overlooked, causing glitches. Or, in adding a new device to your system--for example, if you add a second hard drive to your computer--you may set some jumpers incorrectly, and the machine can't

detect the drive.

How do you know which jumpers or DIP switch settings to use? Start with the manuals for your motherboard and the component. These will show the proper settings for most circumstances.



Avoiding Problems with Device Drivers

Device drivers are another frequent source of hardware foul-ups. A device driver is a software program that acts as an interface between software and hardware (or between two software programs). Video display drivers and printer drivers are among those that most commonly cause frustration for new users. New users often aren't aware that drivers are necessary for these and other components to work, so they think they have a catastrophe on their hands when the video goes haywire, or when the printer doesn't work correctly.






All the drivers you'll ever need aren't automatically installed when you set up Windows 95. You may have to manually install driver upgrades or reinstall drivers if the driver files have been corrupted. You can change device drivers by way of the Add New Hardware Wizard or the Device Manager.

Most devices requiring drivers come with a disk that has the necessary drivers on it. However, if you need a driver, you can quickly find it on the Internet. Go to the Windows Sources page at <http://finders.zdnet.com/> and choose DriverFinder to search for the driver for the component you're using. Or, go to the device manufacturer's Web site.



Dealing with Interrupt Conflicts (IRQ)

Interrupt (IRQ) conflicts are another hardware problem that has wasted millions of man hours of effort. When a component needs a chunk of time from the computer's processor, it sends out its request on an interrupt line. Each device has its own interrupt, but there are a limited number of IRQs available on a PC. If two devices are competing for the same interrupt, conflicts can develop. As you cram your system with more devices, juggling hard



Upgrading Your Hardware

For most of us, cost is a major consideration when deciding upon upgrading versus buying a new system. As you look at the cost of the new parts that you want to upgrade in your system, keep a running tab and compare that to the cost of buying a whole new system.

The cost of buying every part in your system and assembling it yourself (upgrading the whole thing) is much higher than the cost of a new system. PC manufacturers buy in bulk and pass their savings along. In addition, a new system will usually come with a complete warranty and software. If you upgrade your parts, the warranty will depend on the warranty for each part.

In general, if the cost of your upgrade starts to approach \$1,000, you'll want to look carefully at the option of getting a new PC. There's always a fast new system that can be bought in the \$1,000 to \$2,000 range. If you decide a new PC is the way to go, be sure to read the later section "What to Look for in a New PC."

Are You Uncomfortable Opening Your Computer Case?

If you aren't comfortable opening the case of your PC and adding or replacing things inside it, you won't want to consider upgrades to items such as hard drives, video cards, or RAM on your own. However, for a small service charge, most computer stores will perform these upgrades for you. So, if you choose to upgrade and have the service done for you, factor in that cost of the service to the total of the upgrade.

TIP: Professional service should come with a warranty and a guarantee. If you perform your own video card upgrade and you break it or the motherboard, you'll have to pay for the replacement. Hiring a pro with a guarantee should mean that anything that goes wrong is fixed at their cost.

If you don't want to open the case and you aren't going to pay someone else to upgrade it, there still are a few items you can upgrade or add on to increase the life of your PC and what you can do with it:

- *Monitors.* Replacing a monitor is as easy as unplugging the old and plugging in the new.
- *Printers.* A printer is another item you can add or replace by just plugging it in to the back
- *Modems.* An external modem can be added or replaced by hooking it up to an external serial port if you have a serial port that is not being used.
- *Keyboards, mice, and joysticks.* Keyboards can be replaced just by plugging them in to the back of the PC. If you buy a new mouse or joystick with the same connection time as the one on your PC, these can be replaced without opening the PC case.
- *Tape drives.* Many tape drive models can be connected to your parallel port, and most even have a cable to have your printer and tape drive connected to the port at the same time.

- *Zip drives.* If you need more disk storage space, Zip drives come in a version that can be connected to your existing parallel port. This is an easy way to add storage space without opening the PC.
- *Speakers.* If you already have a sound card, adding and replacing speakers is an easy upgrade.

Upgrading Parts Inside Your PC

If you want to upgrade anything inside your PC, it is highly recommended that you read the book *Upgrading PCs Illustrated* by Que. This book has many drawings and photographs to show you step-by-step procedures for adding and replacing parts in your PC. Although most of the parts inside a PC are very easy to replace, they can also be very easy to damage if you don't follow the right steps. This book can save you a lot of time and frustration with those processes.

What type of upgrades can you make inside your PC?

- **Memory.** This is one of the easiest internal upgrades to make. It usually just involves snapping a new RAM or SIMM(s) into place. Your motherboard and BIOS should recognize the new memory when you restart the PC.
- **Processor.** This is another relatively easy but more delicate process. Usually you need to unplug the old processor and plug in the new one. Most PCs have a small lever that helps lift out the old processor and secure the new one in place. You may have to change settings on your motherboard by moving small connectors called *jumper*s, or making changes in the system BIOS.
- **Hard drive.** Adding a new hard drive can be more complicated. In addition to physically mounting the drive in place, you need to connect cables to the drive and configure it to work with an existing drive. Replacing your C drive is even more complicated because you need to have a way to get all of your old files onto the new drive. Once the new drive is installed, you have to set up the BIOS to recognize and format it before storing any data on it.
- **Video card.** Upgrading the video card means removing the old one from its slot and plugging in the new one. You also have to select the correct video card drivers in Windows.
- **Sound card.** Sound cards simply plug into any empty expansion slot.

- *Other expansion cards.* Internal modems, an interface card for a SCSI drive or scanner, a video capture card, and any other type of interface card are added by installing them into an open expansion slot and installing any software that came with them.
- *BIOS.* Most computers built in the last few years have BIOSes that can be upgraded by installing a new software called a *flash* upgrade. However, if you ever have to replace the BIOS chip itself, be sure the replacement comes with good directions and that all of your data on your hard drives is backed up first.
- *System battery.* There's a small battery on your motherboard that runs the clock and helps your BIOS remember settings for your hard drives. These batteries usually last three years or more. There is usually a small clip that holds the battery in place. When replacing these, be sure you have a printed record of your BIOS settings (there's usually a BIOS utility to help with this). Also, make sure to get the right voltage battery and to get it in the correct direction.

IDE and SCSI Drives

Hard drives and CD-ROM drives come in two different varieties. If you want to add a second hard drive, a new CD-ROM drive, or replace a hard drive, you'll need to be sure to get one of the same type that is already in your system.

IDE drives are most common. Unless you specifically asked for SCSI drives when you bought your PC, you probably got it with IDE drives. Most PCs are capable of supporting four IDE drives at one time. So you could have any combination of four IDE hard drives and CD-ROM drives.

There was a time when IDE drives were slower than most SCSI drives. But IDE drives have evolved quickly, and for most uses are now just as fast as SCSI.

NOTE: Some older IDE drives have compatibility issues working with other older IDE drives. If you have an IDE hard drive that is 250M or less and more than two years old, you'll want to consider replacing it with a newer drive rather than adding a second drive. If you do decide to add a second drive, your best bet to avoid incompatibilities is to buy a drive of the same brand as the first.

If your PC does have a SCSI drive connector (either built-in or as an adapter card), be sure to buy SCSI drives. These generally cost more than IDE drives of the same brand and size.

Matching Memory

When it's time to upgrade your computer's RAM, there are several memory type issues you need to be aware of. It's important to buy the right type of memory for your PC, or you will have all sorts of odd problems. Or, the PC may just refuse to boot.

A PC motherboard has slots for memory. (These are different than the expansion slots for cards.) The most common of these slots are called *SIMM (Single Inline Memory Module) slots*. SIMMs come in several different sizes; the most common are 72-pin. These are most common on any PC with a Pentium or faster processors. Some older PCs have 36- or 32-pin SIMMs.

Most PCs have somewhere between four and eight SIMM slots. To decide what memory to buy to upgrade your PC, you should open the case and see how many memory slots are free. If you have no free slots, you will have to replace memory to upgrade. Unfortunately, this situation is common, because many PCs with four slots were sold with all of the slots full.

The most important question in dealing with memory is: "How many megabytes?" Regardless of the number of pins your memory needs to have, the SIMMs can be found in many megabyte sizes. These typically range from 4 to 32M for new RAM. Older RAM that you are replacing may be as low as 1M.

Buying RAM gets tricky at this stage. Even though a SIMM slot may have the right number of pins to accept a PIN, all motherboards have a limit to the number of megabytes that can be on any single SIMM. So, if your motherboard is limited to using 16M SIMMs, don't buy 32M even if they look like they'll fit.

Some motherboards require the SIMMs to be added in pairs called *banks*. If your motherboard has this limitation (check your computer's specifications or contact the vendor to determine), be sure to buy two matching SIMMs to add the RAM you want. For example, to add 16M, buy two 8M SIMMs.

The next complicating factor in buying RAM is choosing parity or non-parity RAM. Most new machines work with *non-parity RAM*. *Parity RAM* is used to check the RAM for internal errors. Again, check your computer specifications to determine which type of RAM your machine requires--then make sure you buy

the right type.

Another RAM term you'll encounter is *EDO* (*Extended Data Output*). This is a faster type of SIMM that has become very standard in PCs. If your motherboard supports it, buy it. It usually sells for the same price as standard (non-EDO) RAM.

If all of this isn't confusing enough, there are several new RAM types and terms you may encounter. Many new PCs are now shipping with slots for *DIMM* memory (*Dual Inline Memory Module*). DIMMs pack more pins (168 instead of 72) and more RAM into just a slightly larger package. If you want to buy DIMMs, be sure to consult an expert first, usually a qualified technician where you bought your PC. DIMMs have a dizzying array of voltages and other specifications in addition to memory amounts, and buying the wrong type will cause problems.

CAUTION: Don't purchase "cheaper" slower additional RAM. Additional RAM must be at least the same speed or faster than the machine's original RAM. Mismatched speeds can cause system problems.

Another new type of RAM is called *SDRAM*. Like EDO, this is a faster kind of RAM. Unlike EDO, SDRAM is currently much more expensive than standard RAM; unless you are a hard-core computer speed addict, it isn't worth the extra money.

ISA, PCI, and Other Expansion Slots

If you are going to install a new card into your PC, you need to know what kind of slot it will go in.

There are two slots that you will find in most current PCs. The first is an *ISA slot*. This is an older standard and is slower than the other type. But, for many uses, it is still more than fast enough. Cards that you'll probably want to buy ISA versions of include:

- Internal modems
- Sound cards

NOTE: If you have an older PC, you may find that you have two different types of ISA slots. One is a short 16-bit slot and the other is a longer 16-bit slot. If you have a system that has 8-bit slots, you can only use 8-bit cards in these. Eight-bit cards will work in either 8- or 16-bit slots, but you'll want to save your 16-bit slots for 16-bit cards.

The other kind of slot is a *PCI slot*. Most PCs now have three of these. These slots are much faster at transferring information from the card to the processor than ISA. They also have the advantage of working better with Plug and Play. Cards that can take advantage of PCI slots include:

- Video cards
- Hard drive interfaces (SCSI or IDE), although most motherboards have IDE built in
- Network adapters

Before buying any card, be sure to check the PC to see that there is a free slot of the right type. If the card is a replacement, remember that you will be able to use

the slot currently filled by the old card.



NOTE: Before there was PCI, there were two other types of fast slots. One was MCA and the other was Vesa Local Bus (VLB). You won't find these on any new PCs, and if you need to upgrade a card in one of these, you may have trouble finding them. If you are considering upgrading a machine with either of these, replacement is probably a better option.



Understanding Plug and Play Technology

Adding and replacing expansion cards and peripherals in your PC is not easy. Most cards have small connectors called *jumper*s that determine what system resources the card will use to communicate with the PC. But no card can use all system resources. These resources are limited, and configuring more than two cards to work together can drive even experienced PC upgraders crazy.

So, Microsoft and several hardware vendors established a new system called *Plug and Play* (available with the release of Windows 95). The theory is that you plug the new card into a slot or plug in the new peripheral, and when you restart the computer, Windows and the motherboard BIOS find the new card and configures it for you.

In order for this to work, your PC has to support Plug and Play and you have to buy hardware add-ons that are labeled as Plug and Play (PnP) compatible. The bad news is that if your system is more than a couple of years old, it definitely won't support this. And while there are a lot of PnP devices available, there are also a lot of devices that aren't PnP. If you have several older non-PnP cards in your PC, these may cause Windows trouble when trying to configure PnP devices.

If you are shopping for new cards, there are a couple of things you'll want to be sure to look for PnP varieties of. First is sound cards. Sound cards can be the most difficult addition to a PC, and Plug and Play can be a huge help with them.

You should also look for PnP varieties of internal modems. If your PC has built-in serial ports and supports PnP, a PnP modem is much easier to install. There are also so many modem varieties that it can be hard to find the exact correct model in the Windows modem setup if you don't have a PnP modem that Windows can identify for you.

Some devices such as hard drives, keyboards, and mice don't need to be PnP. Printers, monitors, video cards, hard drive interface cards, CD-ROM drives, network cards, and scanners all come in PnP varieties.

What to Look for in a New PC

Unless you enjoy the work of upgrading your PC and adding on to it, be sure you get everything that you want in your new PC when you buy it. Here are some guidelines for what to look for.

All of the items in this first list are essential must-have items for your new PC:

- *Processor.* Get the fastest processor you can afford. Although Intel makes the processors sold in the majority of PCs, don't shy away from PCs with AMD or Cyrix processors. I'm writing this on a PC with a nice, new AMD processor, and it's every bit as good as PCs I've had with Intel brains.
- *Memory.* Consider 16M to be the minimum for a new machine. If you can afford the upgrade to 32M, get it now. If not, you'll probably end up buying more later anyway.
- *Hard drive.* You won't find many new PCs with drives smaller than 1G. It is common now to have 2-4G drives. Although it may seem that you would never use all of this space, it will fill up fast, especially if you plan to load a lot of software, games, or files from the Internet.
- Most machines will come with IDE drives. Unless you plan to use the machine for copying or mastering CD-ROMs (with a CD-ROM recorder), an IDE drive will suit you fine. If you do plan to make your own CDs, go ahead and upgrade to a SCSI hard drive now.
- *Monitor.* Don't buy anything less than a 15-inch monitor unless you can't afford it. In fact, you're better off skimping elsewhere to be able to afford a better monitor. For example, the cost difference between a 14-inch and 15-inch monitor is probably about \$75-100. That's usually the same difference between 16M of RAM and 32M. If that difference is what is breaking your budget, get the better monitor. You can always upgrade the RAM later if you have more money and that will still only cost you around \$75-100. But to upgrade the monitor, you'll have to replace it, which will cost around \$300. Your eyes will thank you.

- *CD-ROM drive.* Every new system should come with a CD-ROM drive. Almost all new software that you want to install will come on CDs. And any multimedia game or educational software will require a CD. If you want to make your own CDs, see if you can find a PC that comes with a CD recorder.
- *Keyboard.* If you will be using your PC a lot, invest in a good ergonomic keyboard and avoid any chance that you will develop a medical problem like carpal tunnel syndrome.
- *Mouse.* The new Microsoft IntelliMouse with the roller for scrolling Web pages is a nice addition, but I wouldn't pay extra for it. Look for a mouse that fits your hand comfortably.
- *Floppy drive.* You still need one of these. It should be a 3 1/2-inch model.
- *Case.* Your computer needs a case. Most makers make systems with desktop (flat) and tower (upright) cases. The choice is really a matter of personal preference and the layout of the space where you use the PC.
- *Windows 95.* Your new system should come with Windows 95 preinstalled and ready to run. It should also include the installation CD and documentation. Some vendors save a few dollars by not including a printed manual.

In addition, here's what to look for in other optional equipment:

- *Tape drive or other removable drive.* You probably won't buy a tape drive until after the first time you have a bad system crash and lose data, but you should. Many new systems include these. The Zip and Jaz drives are also good alternatives for backing up data or moving large files. Zip drives cost about \$200, but I've seen new PCs with these as an option for about \$100. You'll get similar savings by buying the PC with a tape drive instead of adding one later.
- *Sound card and speakers.* Although not essential for your PC to operate, almost every new PC comes with a sound card. Be sure it is Sound Blaster-compatible. If you want booming bass, upgrade to speakers with a subwoofer.
- *Modem.* Don't buy anything less than a 33.6K modem. If you want to use

the Internet a lot, consider an upgrade to a 56K modem. If buying a 56K modem, be sure that you can get 56K service and that your brand modem will work with the provider.

- **Printer.** These are easy enough to add on your own, but you can usually get a good deal on one if you buy it with the PC. If you just need to print black and white business documents, an inexpensive laser printer is your best bet. For color printing, ink jets and bubble jet color printers have great output quality and low prices. But be prepared to spend money on refills for the ink.
- **Joysticks, steering wheels, and other game input hardware.** Buying these amounts to matters of personal taste and what games you like to play. If you want to be a master of *Quake* or *Duke Nukem*, don't expect to do it with a bargain-basement \$10 joystick. Likewise, there's a reason real racecars have steering wheels and gas pedals. If you live to play *Indy Car Racing* or *Sega Rally Championship* on your PC, buy a PC steering wheel with pedals if your games support it.
- **Optional software.** If you use Microsoft Office (Word, Excel, PowerPoint, Access) at work, look for a PC that has these preinstalled if you want to use your PC to work on these types of files. Many vendors offer lots of preinstalled software on PCs. If the software is something you'll use, it's a great deal. If not, you'll just end up uninstalling it at some later date.

Where to Buy New PCs and Upgrades

There are two main options to look at when buying PCs or parts. You can buy them from any local store that sells them, or you can buy them from a mail-order (direct) outlet.

The advantages of buying from a local store usually involve convenience. If you want the security of seeing the PC or part before you buy it, or knowing that you can return it or have it serviced by simply driving to the local store at the mall, buying locally is for you.

TIP: When shopping for a new monitor, I almost always recommend buying it locally rather than mail-order. You'll want to see what the picture looks like on-screen as monitors with all of the same technical specifications can have dramatically different picture quality. You can't see that in a magazine ad or over the phone.

Buying locally usually gives you a chance to test-drive the system in the store and ask a helpful sales representative questions about the hardware. A little common sense goes a long way here. A lot of stores sell PCs now. Chances are the salesperson at the local PC store will be able to answer more of your PC questions than the salesperson at an electronics store who also sells refrigerators and TVs, or the salesperson at a department store selling everything from underwear to computers. The bottom line is, shop where you feel most comfortable.

The big advantage of buying from a mail-order or direct outlet is usually price. You will be hard pressed to find a local retail store that can sell you a new system with all of the same features as a mail-order outlet for the same price.

Buying direct also has the advantage of having a much larger selection. If you like to have a lot of choices, call several mail-order outlets and compare their offerings. You'll always find a better selection than at the local stores.

A final advantage that most major mail-order retailers have is a better trained

sales staff. For example, when you call Gateway or Dell's order lines, the person that answers the phone does nothing but sell computers. Before they ask you technical questions about what type of processor you need or how much RAM you want, they do a good job of finding out how you plan to use the system so that they understand your needs. Gateway and Dell aren't the only ones; any of the major PC vendors who sell direct can help you a lot when buying a new system.

When buying parts through a direct channel, I recommend picking up a copy of *Computer Shopper* magazine before buying. Any computer store carries this massive magazine, and you can find dozens of reputable dealers listed there. For about \$5, this is a great investment in the most up-to-date information about who's selling what for how much.


Special Tips for Buying a Notebook Computer

There are several things about buying a new notebook that are different than buying a PC. Here's a quick list of additional or different features you'll want to look at in a notebook.

- **Processor.** The processor should be a special *mobile* processor. These are made to run cooler and use less power so your battery will last longer. Some notebook makers put desktop processors in their notebook to save money. I recommend avoiding notebooks that use desktop processors.
- When buying a notebook, buy one with the processor you intend to use for the life of the notebook. You'll never upgrade the processor.
- **Video card and display.** These are built into the system in the notebook and can't be upgraded. Be sure to get ones you'll be happy with as long as you own the notebook.
- **PC Card slots.** This is a different type of expansion slot for notebooks. Be sure to get one with two PC Card slots. The newest technology for these is Cardbus and Zoomed Video. Get one that supports these standards if you can.
- **Battery.** Battery life is critical if you travel with your notebook. Lithium Ion (LiO) batteries are the best for life. Nickel-Metal Hydride (NmH) are also good. Nickel-Cadmium is an old battery type that you should avoid. Look for a notebook that uses standard battery types so replacements will be inexpensive and easy to find. If possible, get one with *hot swappable* batteries so you can buy a second battery and switch without rebooting.
- **Power management.** Only buy a notebook that supports power management to save battery life.
- **CD-ROM drive.** Most notebooks now have internal CD-ROM drives standard or at least as an option. Many times, these are removable and can be swapped with the floppy drive and second battery. You won't be able to install new CD-based software without a CD-ROM drive, so your

notebook really should have one.

- **RAM.** Expect to pay a little more for RAM in a notebook than in a desktop system. Most notebooks use proprietary RAM modules, but a few use standard desktop SIMMs. In some notebooks, upgrading RAM will be more difficult than in a desktop system, so you may want to buy it with all the RAM you think you'll ever need.
- **Keyboard.** The closer to full size and regular key spacing, the easier it will be to use.
- **Pointing device.** Most of the time, you won't use a mouse with the notebook. Most notebooks now come with a small, built-in *track pad* that you drag your finger across for mouse options. I prefer the pointing stick that you find on IBM ThinkPads and a few off-brands, but these are hard to find outside of IBM. Trackballs are very rare now.
- **Hard drives.** Expect to pay a little more for your notebook hard drive, too. Capacities are a little smaller than desktop drives and they can be harder to replace. Most notebooks can't take a second drive.
- **Sound card and speakers.** Most notebooks include a basic Sound Blaster-compatible card and speakers as a standard feature.
- **Weight.** If you travel with a notebook, buy the lightest-weight model you can with all of the features you need. Consider the weight of a carry case, extra battery, charger, and anything else you plan to carry.
- **External ports.** There should be ports for connecting a mouse, keyboard, printer, and external monitor. Optional ports may include speakers and serial plugs (for an external modem).
- **Port replicator.** If you plan to use the notebook as a replacement for a desktop system, you'll want to see if a port replicator is available. The port replicator has all of the same external ports as the notebook. You plug the monitor, keyboard, and mouse into the replicator and the notebook snaps into the replicator. Leave the external devices plugged in to the replicator,



and you can plug and unplug the notebook without having to attach the monitor, keyboard, and mouse individually.

- **Docking station.** This is the big brother to the port replicator. A docking station has all of the same features as the replicator with additional features such as bays for added drives and slots for expansion cards.

AGP - PCI Which is Better !!

Technical Background of AGP

AGP has several advantages over PCI. It offers a minor data transfer rate advantage when it comes to moving the geometry stream from the CPU to the graphics card. When it comes to managing large texture databases, AGP's GART table allows the OS to manage textures in off screen memory as well as in system memory, and allows the graphics card to access them directly in either location. Prior to AGP, the game developer had three options for methods to manage textures:

1. Limit the texture database to whatever would fit in off-screen memory only. This usually delivers outstanding frame rates, but memory size constraints can limit artistic creativity. Depending on the graphics card, texture space could be as low as only one megabyte or possibly as high as five or six megabytes.
2. Use the OS to manage textures in main memory, and require the CPU to copy textures from main memory to graphics memory as needed. This is PAINFULLY slow. In order to make AGP look as good as possible, Intel likes to compare AGP to this mode. This is what DirectX Retained Mode does, but game developers do not generally use this method.
3. Place the most frequently used textures in off screen memory (like #1 above), then lock down a few additional megabytes of system memory for the remainder of the texture database. If the graphics chip needs a texture that is not in graphics memory, then the accelerator must use PCI master mode (DMA) to copy the needed texture, on demand, into texture swap space in graphics memory. Performance is very good. This has been the preferred approach for game developers, and can be programmed under Direct X Immediate Mode.

AGP is a modified approach to option 3. Memory management is a little more flexible, and the data transfer rate is better because of AGP's faster clock speed and bus pipelining.

Intel is always trying to stir up more ways to persuade users to dump their PCI

Pentium 233 systems ASAP, and go buy a more costly P3 AGP system. Intel's Mission is not to "Accelerate 3D Graphics", but rather to "Accelerate Obsolescence".

The other side of the performance equation is the CPU. What happens to the CPU when AGP texturing is activated? When AGP texturing is turned on, the graphics card takes control of the main memory bus in order to access texture data. When this happens, the CPU is locked out of main memory. If the CPU is not very busy, this may not be a problem. But if the CPU is engaged in a computationally challenging task (such as a game) it is highly possible that the CPU may stall, waiting for its turn to access main memory. The main memory conflicts between the CPU and the graphics controller will result in a CPU performance reduction of more than 10%. This assumes the use of a Pentium II with the back side cache intact. The cacheless Covington would be brought to its knees under these circumstances.

The assumption that AGP is a little faster now, but will be a lot faster when it is "really turned on" is completely false. In fact the opposite is true. In most cases, system which actually use AGP for textures will be potentially 40% slower than systems which use local graphics memory for textures instead.

Is PCI better than AGP?

No. AGP is more powerful, more flexible, and has more potential than PCI. If you have the choice, and everything else is equal, AGP should be preferred. However, the power user must be made aware that this flexibility can be used to create resource conflicts that might be unfavorable for ultimate game performance. In contrast, these trade-offs may be perfectly acceptable for the mainstream user. On the whole, the urgency to invest in an upgrade to AGP is not all that it is cracked up to be.

Is AGP faster than PCI?

"Yes, but not by much."

AGP is definitely faster than PCI. According to Read 3d , the theoretical

bandwidths of AGP and PCI are 533 MB/sec and 133 MB/sec, respectively.

With existing systems, we can demonstrate greater than 330 MB/sec through AGP. Due to contention with other devices on the PCI bus and the lack of pipelining, PCI offers roughly 80 MB/sec. Even without using any of the extra features of AGP, the AGP bus runs at 66 MHz and the PCI bus runs at 33 MHz so AGP provides double the data rate. HA HA HA
.....(from Intel docs)

Theoretical peak bandwidth numbers can easily be used for impressing people. However when it comes down to real world applications how much is left from this four time speed.

Does AGP enhance CPU performance?

"No. When AGP texturing is used, CPU performance also decreases."

Accessing local card memory will always be faster, unless the local memory interface was designed very poorly. Local graphics memory gets more and more each year, 16 MB will soon be the rule. High end OpenGL cards use up to 64 MB and more, why would they do that if AGP was a real alternative? Memory gets cheaper by the hour. It's simply untrue to say that local memory will be a problem. It's cheap, it's available and the amount is increasing by 100% every year for the last 4 years. **AGP is a *cost saving* solution, not a high end solution!!!!**

When that time comes, should I buy a faster CPU or more graphics memory?

Intel wants you to think that you need a faster CPU, when in fact all you may need is a few dollars worth of graphics memory...

Don't I need AGP2x or 4x for the future?

All AGP4x does is allow the graphics controller to assert more of its texture bandwidth demand on main memory. This is the source of the performance

conflict, and 4x allows it to become a bigger problem. But there are two sides to examine: Texture demand and geometry transfer rates. We will do that in a minute.

When it is all said and done, the central issue is about graphics memory. The AGP vs. PCI issue gets dragged in because AGP is proposed as an alternative way to access more memory. Let me re-summarize my philosophy on this matter.

RULE #1 - Make sure that you have adequate graphics memory. If you have adequate graphics memory, you will not observe much of a performance difference between AGP and PCI. If you do not have adequate graphics memory, you are eventually going to run into texture related performance problems with PCI or with AGP. There are MANY variables that determine how big those problems are, and how they should best be solved.

Though I don't think AGP is bad, relying on AGP's execute mode texturing to recover from breaking Rule#1, is a less than ideal solution. It may potentially sacrifice MIPS, sacrifice accelerator performance, or both. If you already paid for AGP, and you inadvertently violated rule #1 in the process, AGP becomes a convenient, "no additional cost" way to deal with the problem. Convenient it is, but performance is another matter, and performance is not "free".

So far, texture bandwidth demand has grown to become quite significant, but the future holds further astronomical growth in texture bandwidth requirements for the high performance gamer.

Today and in the future, high quality rendering modes use multitexturing, trilinear mip-mapped texturing, or anisotropic texturing. In these modes, the rendering engine must access 8 to 16 or more texels for each pixel that is created. At a depth complexity of three, the rendering engine will have to do this three times for each pixel on the display in every frame.

Using these assumptions, following is a table that shows the raw texture bandwidth required to deliver content at the resolutions and quality levels described below. The columns define the level of quality that the PC enthusiast expects (or expected) during that year.

It is not AGP or PCI that determines the performance of your 3D

accelerator. Performance will be determined by the accelerator chip and its memory resources. The bottom line is, don't choose your accelerator based on its bus interface. Choose it based on the chip that is on the board and by how its DRAM is configured.