

The BIOS Survival Guide

Version 5.1

Co-Edited by Jean-Paul Rodrigue (rodriguj@ere.umontreal.ca) and Phil Croucher

29-06-1996

Portions Copyright 1996, Jean-Paul Rodrigue

Portions Copyright 1996, Phil Croucher

Contributors

Major Contributors

I would like to thank the contributors to have taken some of their time to write varied topics and provide enlightening feedback. Some parts of this document are still incomplete and some information may be inaccurate. **Your feedback** will help this document to be **as accurate and up to date as possible**. I am sorry if I cannot answer to everyone or add everything that is sent to me.

- Gordon L. Burditt (gordon@sneaky.lonestar.org)
- Phil Croucher (100427.2731@compuserve.com)
- Peter Herweijer (pieterh@sci.kun.nl)
- Kajetan Hinner (uf341ea@sunmail.lrz-muenchen.de)
- Piotr Karocki (yskarock@cyf-kr.edu.pl)
- Donald Kenney (72630.3560@compuserve.com)
- Brian Lee (blee@bart.conestogac.on.ca)
- Aad Offerman (offerman@einstein.et.tudelft.nl)
- Keith Rohrer (rohrer@fnrcrd8.fnal.gov)
- Jerome Schneider (jls@atg.com)
- Will Spencer (will@spiff.gnu.ai.mit.edu)
- Cameron Spitzer (cls@truffula.sj.ca.us)

Minor Contributors

I would also like to thank the following persons for providing information on specific topics: Mark Bergman, Michiel de Vries, Andy Eskilsson, Doug Hogarth, Reinhard Kirchner, Mirek Komon, Jim Kozma, Juha Laiho, Alain Lavoie, Erik Mouw, Chris Pollard, Dietrich Schmidt, Hans Schrader, Loren Schwiebert, Dan Sobel, Dave Spensley, Dmitry Stefankov, Andy Walton and Robert W. Wong Jr.

Special Thanks to Phil Croucher

Special thanks to Phil Croucher (100427.2731@compuserve.com) for sharing with us information in *The BIOS Companion*, a guide he publishes, among other technical manuals such as *The Hard Disc Database*. These are a GREAT manuals, and every computer technician should have them. He can be contacted at P.O. Box 52083, Winnipeg, Canada, R2M 5P9. Phone: (204) 257-9721. Fax: (204) 253-2936.

Also, special thanks to Alain Lavoie (lavoiea@ere.umontreal.ca) for providing the server for this document.

Definitions

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 "i?s 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).

The Motherboard

This is a large circuit board to which are fixed the **Central Processor** (CPU), the data bus, memory and various support chips, such as those that control speed and timing the keyboard, etc. The CPU does all the thinking, and is told what to do by instructions contained in memory, so there will be a direct two-way connection between them. The data bus is actually a part of the CPU, although it's treated separately.

Extra circuitry in the form of **expansion cards** is placed into **expansion slots** on the data bus, so the basic setup of the computer can be changed easily (for example, you can connect more disk drives or a modem there).

A math co-processor is frequently fitted alongside the main processor, which is specially built to cope with floating point arithmetic (e.g. decimal points). The main processor has to convert decimals and fractions to whole numbers before calculating them, and then has to convert them back again.

The Central Processor

The chip that was the brains of the original IBM PC was called the 8088, manufactured by Intel. Although it was classified as being 16-bit, it spoke to the data bus and memory with 8 bits in order to keep the costs down and keep in line with the capabilities of the support chips. Then when it wanted to send two characters to the screen over the data bus, it had to send them one at a time, rather than both together, so there was an idle state where nothing was done every time the data was sent.

In addition, it could only talk to 1 Mb of memory at any time; there were 20 physical connections between it and the CPU. On a binary system this represents 2^{20} , or 1024Kb.

The 80286

The 80286 was introduced in response to competition from manufacturers who were cloning the IMB PC. The connections between the various parts of the motherboard became 16-bit throughout, thus increasing efficiency 4 times. It also has 24 memory address lines, so it could talk to 16Mb of physical memory. However, DOS could not use it since it had to be addressed in protected mode. DOS can only run in real mode, which is restricted to the 1 Mb that can be seen by the 8088. Therefore, a Pentium running DOS is just a fast XT.

The 80386

Compaq was the first company to use the 80386 (the DX version, as opposed to the SX-see below), which uses 32 bits between itself and memory, but 16 towards the data bus, which has not really been developed in tandem with the rest of the machine. This is partly to ensure backwards compatibility and partly due to the plumbing arrangements of running a fast CPU with faster memory and a slow bus (8 MHz).

The 386 can run multiple copies of real mode (that is, it can create several virtual 8088s). It uses paging to remap memory so that these machines are brought to the attention of the CPU when the programs in them require it; this is done on a timeslice basis, around 60 times a second, which is how we get multitasking in Windows or OS/2 (in OS/2 the slice is every 20 ns).

The 386 can also switch out of protected mode on the fly, or at least in a more elegant way than the 286; in order to get the hard disk and other parts of the computer, protected-mode software has to get DOS to perform real mode services, so the CPU has to switch in and out of protected mode continually. The goal is therefore to use real mode as little as possible and to run in protected mode. Windows does this by using 32-bit instructions.

The 386 uses pipelining to help streamline memory access-the idea is that they are done independently of each other (at the same time) while other units get on with their jobs; a form of primitive parallel processing. The 386 also has a pre-fetch unit for instructions, that tries to speed things up by guessing which ones the processor will use next.

Although the 386 is 32-bit and has certain benefits, like the ability to manipulate memory and switch in and out of protected mode more readily, replacing a 286 with a 386 does not automatically give you performance benefits if you are running 16-bit 80286 code (most DOS programs).

The 80386SX

The 80386Sx is a 32-bit chip internally, but 16-bit externally to both memory and the data bus, so you get bottlenecks. It is a cut-down version of the 80386DX, created to both cut costs and give the impression the 286 was obsolete (true), because at the time other manufacturers could make the 286 under license. Although it can run 386-specific software, it looks like a 286 to the machine it is in, so existing 286 motherboards could be used to plug in 386SX CPUs. At the same clock speed, a 386SX is around 25% slower than a 386DX.

The 80486

To non-technical people, the 80486 is a fast 80386DX with an on-board math co-processor and 8K of cache memory. It is not really a newer technology as such (only second generation), but better use is made of its facilities. For example, it takes fewer instruction cycles to do the same job, and is optimized to keep as many operations inside the chip as possible. The 386 pre-fetch unit was replaced by 8K of SRAM cache, and pipelining was replaced by burst mode, which works on the theory that most of the time spent getting data concerns its address. Burst allows a device to send large amounts of data in a short time without interruption. Pipelining on the 386 requires 2 clocks per transfer; only one is needed with 486 Burst Mode. Memory parity checks also take their own path at the same time as the data they relate to. The 486 has an on-board clock, and both edges of the square wave signal are used to calculate the clock signal, so the motherboard runs at the same speed as the CPU. In addition, the bus system uses a single pulse cycle. Generally speaking, at the same clock speed, a 486 will deliver between 2-3 times the performance of a 386.

The 80486SX

The 486SX is as above, but with the math co-processor facility disabled, therefore you should find no significant difference between it and a 386; a 386/40 is broadly equivalent to a 486/25.

Clock Doubling

The DX/" chip runs at double speed of the original, but it is not the same as having a proper high speed motherboard because the bus will still be running at the normal speed. Unfortunately, high speed motherboards are more expensive because of having to design out RF emissions, and the like.

Actual performance depends on how many accesses are satisfied from the chip's cache, which is how the CPU is kept busy, rather than waiting for the rest of the machine. If the CPU has to go outside the cache, effective speed is the same as the motherboard or, more properly, the relevant bus (memory or data), so best performance is obtained when all the CPU's needs are satisfied from inside itself. However, performance is still good if it has to use cache, as the hit rate is around 90%. The DX4 has a larger cache (16K) to cope with the higher speed.

The Pentium

Essentially two 486s in parallel (or rather an SX and a DX), so more instructions are processed at the same time; typically two at once. This, however, depends on whether software can take advantage of it, and get the timing of the binary code just right. It has separate 8K caches, for instructions and data, split into banks which can be accessed alternately. It has a 64-bit bus, to cope with 2 32-bit chips.

The Pentium Pro

This is a RISC chip with a 486 hardware emulator on it. Several techniques are used by this chip to produce more performance than its predecessors; speed is achieved by dividing processing into more stages, and more work is done within each clock cycle; three instructions can be decoded in each one, as opposed to two for the Pentium.

In additions, instruction decoding and execution are decoupled, which means that instructions can still be executed if one pipeline stops (such as when one instruction is waiting for data from memory; the Pentium would stop all processing at this point). Instructions are sometimes executed out of order, that is, not necessarily as written down in the program, but rather when information is available, although they won't be much out of sequence; just enough to make things run smoother.

It has a 8K cache for programs and data, but it will be a two chip set, with the processor and a 256K L2 cache in the same package. It is optimized for 32-bit code, so will run 16-bit code no faster than a Pentium.

Summing up

In principle, the faster the CPU the better, but only if your applications do a lot of logical operations and calculation (where the work is centered around the chip) rather than writing to disk. For example, when a typical word processing task, replacing a 16 MHz 386 with a 33 MHz one (doubling the speed) will only get you something like a 5-10% increase in practical performance, regardless of what the benchmarks might say. **It is often a better idea to spend money on a faster hard disk.**

Also, with only 8 Mb RAM in your computer, you won't see much performance increase form a DX2/66 until you get a Pentium 90 (none at all between a DX4/100 and a Pentium 75). With Windows, this is because the hard disk is used a lot for virtual memory (swap files), which means more activity over the data bus. Since motherboards below the 90 run at 33Mhz (only the chips run faster), the bottleneck is the disk I/O, running at much the same speed on them all. This is especially true if you use Programmed I/O (PIO), where the CPU must scrutinize every bit to and from the hard drive (although Multi-sector I/O or EIDE will improve things). As the Pentium 90's motherboard runs faster (66 MHz), the I/O can proceed at a much faster pace, and performance will more than double (a more sophisticated chipset helps).

With 16 Mb, on the other hand, performance will be almost double anyway, regardless of the processor, because the need to go to the hard disk is so much reduced, and the processor can make a contribution to performance. The biggest jump is from a DX2/66 to a DX/4, with the curve flattening out progressively up to the Pentium 90.

Memory

Memory Types

Unlike the good old days of mainframe computers when there were only two types of memory to deal with, the PC uses five, again for historical reasons. The memory contains the instructions that tell the Central Processor what to do, as well as the data created by its activities. Since the computer works with the binary system, memory chips work by keeping electronic switches in one state or the other for however long they are required. Actually, they consist of a capacitor and a transistor; the capacitor stores a charge (data), which represents a 1, and the transistor acts as a switch that turns the charge on or off. Where these states can be changed at will, it is called **Random Access Memory**, or RAM. The term derives from when magnetic tapes were used for data storage, and the information could only be accessed sequentially. A ROM, on the other hand, is a memory chip with its electronic switches permanently on or off, so they can't be changed, hence **Read Only Memory**.

- **Static RAM (SRAM)** is the fastest available, with a typical access time of 25 nanoseconds. Static RAM is expensive and can only store a quarter of the data that DRAM is able to, as it uses two transistors to store a bit against DRAM's one, although it does retain it for as long as the chip is powered. The transistors are connected so that only one is either in or out at any time; whichever one is in stands for a 1 bit. Synchronous SRAM allows a faster data stream to pass through it; which is needed when used for cacheing on 90 and 100 MHz Pentium.
- **Dynamic RAM (DRAM)** uses internal capacitors to store data (a single transistor turns it on or off) which lose their charge over time, so they need constant refreshing to retain information, otherwise 1s will turn to 0s. The end result is that between every memory access is sent an electrical charge that refreshes the chip's capacitors to keep data in a fit state, which cannot be reached whilst recharging is going on. Reading a DRAM discharges its contents, so they have to be written back to immediately to keep the same information.
- **Enhanced DRAM (EDRAM)** replaces standard DRAM and the SRAM in the level 2 cache on the motherboard, typically combining 256 bytes of 15ns SRAM inside 35ns DRAM. Since the SRAM can take a whole 256 byte page of memory at once, it gives an effective 15ns access speed when you get a hit (35ns otherwise). The level 2 cache is replaced with an SIC chip to sort out chipset vs. memory requirements. System performance is increased by around 40%. EDRAM has a separate write path that accepts and completes requests without the rest of the chip.
- **WRAM (Windows RAM)**, created by Samsung, is dual ported, but costs about 20% less than VRAM and is 50% faster. It runs at 50 MHz and is optimized for acceleration and can transfer blocks and supports text and pattern fills. Mostly used for video cards.
- **Synchronous DRAM (SDRAM)** takes memory access away from the CPU's control; internal registers in the chips accept the request, and let the CPU do something else while the data requested is assembled for the next time the CPU

talks to the memory. As they work on their own clock, the rest of the system can be clocked faster. There is a version optimized for video cards.

- **EDO** (Extended Data Output) is an advanced version of fast page mode (often called Hyper Page Mode), which can be up to 30% better and only cost 5% more. Single-cycle EDO will carry out a complete memory transaction in 1 clock cycle; otherwise, each sequential RAM access inside the same page takes 2 clock cycles instead of 3, once the page has been selected. As it replaces level 2 cache and doesn't need a separate controller, space on the motherboard is saved, which is good for notebooks. It also saves battery power. In short, EDO gives and increased bandwidth due to shortening of the page mode cycle, but it doesn't appear to be that much faster in practice.

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).

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 of caching 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.

Base Memory

The first 640K available, which traditionally contains DOS, device drivers, TSRs and any programs to be run, plus their data, so the less room DOS takes up, the more there is for the rest. Different versions of DOS were better or worse in this respect. In fact, under normal circumstances, you can expect the first 90K or so to consist of:

- An *Interrupt Vector Table*, which is 1K in size, including the name and address of the program providing the interrupt service. Interrupt vectors point to routines in the BIOS or DOS that programs can use to perform low level hardware access. DOS uses **io.sys** and **msdos.sys** for the BIOS and DOS, respectively.
- ROM BIOS tables, which are used by system ROMs to keep track of what's going on. This will include I/O addresses and possibly user-defined hard disk data.
- DOS itself, plus any associated data files it needs to operate with (e.g. buffers, etc.).

DOS was written to run applications inside the bottom 640K block simply because the designers of the original IBM PC decided to-memory at the time was expensive, and most CP/M machines only used 64K anyway (the PC with 128K was \$10,000!). Other machines of the same era used more; the Sirius allowed 896K for programs. Contrary to popular belief, Windows uses memory below 1Mb, for administration purposes; although it pools all memory above and below 1 Mb (and calls it the *Global Heap*), certain essential Windows structures must live below 1 Mb, such as the *Task DataBase* (TDB) which is necessary for starting new tasks.

Every Windows application needs 512 bytes of memory below 1 Mb to load, but some will take much more if they can, even all that's available, thus preventing others from loading, which is one source of "Out of Memory" messages. There are programs that will purposely fragment base memory so it can't be hogged by any one program.

Rather than starting at 0 and counting upwards, memory addressing on the PC uses a two-step **segment:offset** addressing scheme. The *segment* specifies a 16-byte paragraph of RAM; the *offset* identifies a specific byte within it. The CPU finds a particular byte in memory by using two registers. One contains the starting segment value and the other the offset. The maximum that can be stored in one is

65,535 (FFFF in hex). The CPU calculates a physical address by taking the contents of the segment register, shifting it one character to the left, and adding the two together (see High Memory).

Sometimes, you will see both values separated by a colon, as with FFFF:000F, meaning the sixteenth byte in memory segment FFFF; this can also be represented as the effective address 0FFFFFFh. When referring only to 16-byte paragraph ranges, the offset value is often left out. The 1025KB of DOS memory is divided into 16 segments of 64K each. Conventional memory contains the ten segments from 0000h to 9FFFh (bytes 0 to 6555,167), and Upper memory contains the six segments ranging from A000h to FFFFh.

Upper Memory

The next 384K is reserved for private use by the computer, so that any expansion cards with their own memory or ROMs can operate safely there without interfering with programs in base memory, and *vice versa*. Typical examples include network interface cards or graphics adapters. **There is no memory in it;** the space is simply reserved. This is why the memory count on older machines with only 1 Mb was 640 + 384K of *extended memory*; the 384K was remapped above 1 Mb so it could be used. When upper memory blocks are needed, that memory is remapped back again, so you lose a bit of extended memory.

This area is split into regions, A-F, which in turn are split into areas numbered from 0000 to FFFF hexadecimally (64K each). With the right software, this area can be converted in *Upper Memory* for use by TSR (memory-resident programs) to make more room downstairs. The amount of upper memory available varies between computers, and depends on the amount of space taken up by the System BIOS and whether you have a separate VGA BIOS (on board video sometimes has its BIOS integrated in the system BIOS). It also depends on the number of add-in cards you have, e.g. disk controllers, that normally take up around 16K.

Some chipsets will always reserve this 384K area for shadowing, so it will not appear in the initial memory count on power-up, the system configuration screen, or when using MEM. Other chipsets have a *Memory Relocation* option which will re-address it above 1 Mb as extended memory. Occasionally, some ROM space is not needed once the machine has booted, and you might be able to use it. A good example is the first 32K of the System BIOS, at F000 in ISA machines. It's only used in the initial stages of booting up, that is, before DOS gets to set up device drivers, so this area is often useable.

Extended Memory

Memory above 1 Mb is known as **extended memory**, and is not normally useable under DOS, except to provide RAM disks or caches, because DOS runs in *real mode*, and it can't access extended memory in protected mode which OS/2 and Windows 95 do. Some programs are able to switch the CPU from one to the other by using the *DOS Protected Mode Interface* (DPMI). Although extended memory first appeared on the 286, and some software was written to take advantage of it, the 286 was used mostly as a fast XT, because DOS wasn't rewritten. It wasn't until the 386, with its memory paging capability, that extended memory came to be used properly.

High Memory

The first 64K (less 16 bytes) of extended memory, which is useable only by 286, 386 or 486 based computers that have more than 1Mb of memory. It's a quirk in the chip design (or a bug!) that can be exploited by DOS to use that portion of extended memory as if it were below 1Mb, leaving yet more available for programs in base memory. In other words, it is extended memory that can be accessed in real mode. It is activated with **himem.sys**.

HMA access is possible because of the **segment:offset** addressing scheme of the PC. Memory addresses on a PC are 20 bits long, and are calculated by shifting the contents of a 16-bit register 4 bits to the left, and adding it to a 16-bit offset. The 8088, with only 20 address lines, cannot handle the address carry bit, so the processor simply wraps around to address 0000:0000 after FFFF:000F; in other words, the upper 4 bits are discarded.

On a 286 or later, there is a 21st memory address line that was left open by accident and which can be operated by software, which gives you a dirty bit. If the system activates this bit while in 8088 (real) mode, the wraparound doesn't happen, and the high memory area becomes available.

Expanded Memory

This is the most confusing one of all, because it sounds so much like expansion memory, which is what extended memory is sometimes called. Once the PC was on the market, it wasn't long before 640K wasn't enough, particularly for people using Lotus, the top-selling application of the time, who were creating large spreadsheets and not having enough memory to load them, especially when version 2 needed 60K more memory than the original. It wasn't entirely their fault; Lotus itself in its early days was very inefficient in its use of memory anyway.

Users got onto Lotus, Intel and Microsoft for a workaround, and they came up with LIM memory, also known as **Expanded**. It's a system of physical bank-switching, where several extra banks of memory can be allocated to a program, but only one will be in the address space of the CPU at any time, as that bank switched, or **paged**, as required. In other words, the program code stays in the physical cells, but the electronic address of those cells is changed, either by software or circuitry.

In effect, LIM (4.0) directly swaps the contents of any 16K block of expanded memory with a similar one inside upper memory; no swapping takes place, but the pages have their address changed to look as if it does. Once the page frame is mapped to a page on the card, the data of that page can be seen by the CPU. Points to note about LIM:

- It's normally only available for data (no program code).
- Programs need to be specially written to use it.

In theory, LIM 4 doesn't need a page frame, the programs you run may well expect to see one. In addition, there could be up to 64 pages, so you could bank switch up to a megabyte at a time, effectively doubling the address space of the CPU, and enabling program code to be run and multitasking. This was called *large-frame EMS*, but it still used only four pages in upper memory; the idea was to remove most of the memory on the motherboard. The memory card *backfilled* conventional memory and used the extra pages for banking.

On an 8086 or 286-based machine, expanded memory is usually provided by circuitry on an expansion card, but there are some software solutions. 386 (and 486) -based machines have memory management built in to the central processor, so all that's needed is the relevant software to emulate LIM (**emm386.exe** or similar).

Virtual Memory

"Virtual" in the computer industry is a word meaning that something is other than what it appears to be. In view of that, Virtual Memory isn't memory at all, but hard disk space made to look like it; the opposite of a RAM disk. Windows (and System 7) uses virtual memory for *swap files*, used when physical memory runs out (you need protected mode on a PC to do that). Like disk cacheing, VM was used on mainframes for some time before migrating to the PC; VMS, the OS used on DEC VAXes, actually stands for *Virtual Memory System*. There is a speed penalty, of course, as you have to access the hard disk to use it, but Virtual Memory is a good stopgap when you're running short.

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 card 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 transfer 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 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 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! For more information about this scam (mostly occurring in England), please consult this Web site: .
- 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). For more information (notably, to value to choose depending on your memory speed) go to the section.

- Make sure to **shadow** your Video and System ROM. On older systems, this may improve performance significantly. For more information, go to the section.

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 **CRTL-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 .
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).

BIOS Identification Line

There may be up to three identification lines in the BIOS startup screen.

First Identification Line

(1) Processor Type

- ? 8088 or 8086
- ? 80286
- ? 80386
- X 80386SX
- 4 80486
- 5 Pentium

(2) Size of BIOS

- 0 64K
- 1 128K

(4-5) Major Version Number

(6-7) Minor Version Number

(9-14) Reference Number

(16) Halt on POST Error

- 0 Off
- 1 On

(17) Initialize CMOS in every boot

- 0 Off
- 1 On

(18) Block pins 22 and 23 of the keyboard controller

- 0 Off
- 1 On

(19) Mouse support in BIOS/keyboard controller

- 0 Off
- 1 On

(20) Wait for <F1> if error found

- 0 Off
- 1 On

(21) Display floppy error during POST

- 0 Off
- 1 On

(22) Display video error during POST

- 0 Off
- 1 On

(23) Display keyboard error during POST

- 0 Off
- 1 On

(25-26) BIOS Date Month

(27-28) BIOS Date Day

(29-30) BIOS Date Year

(32-39) Chipset Identification / BIOS Name

(41) Keyboard controller version number

Second Identification Line

(1-2) Pin number for clock switching through keyboard controller

(3) Indicates high signal on pin switches clock speed to high or low

- H High
- L Low

(5) Clock switching through chipset registers

- 0 Off
- 1 On

(7-10) Port address to switch clock high

(12-13) Data value to switch clock high

(15-16) Mask value to switch clock high

(18-21) Port address to switch clock low

(23-24) Data value to switch clock low

(26-27) Mask value to switch clock low

(29-31) Pin number for turbo switch

Third Identification Line

(1-2) Keyboard Controller Pin for cache control

(3) Indicates whether high signal on cache control pin enables or disables cache

- H Enable
- L Disable

(5) Indicates if the high signal if used on the keyboard controller pin

- 0 False
- 1 True

(7-9) Cache control through chipset registers

- 0 Cache control off
- 1 Cache control on

(11-12) Port address to enable cache through special port

(14-15) Data value to enable cache through special port

(17-20) Mask value to enable cache through special port

(22-23) Port address to disable cache through special port

(25-26) Data value to disable cache through special port

(28-29) Mask value to disable cache through special port

(31) Pin number for resetting the 82335 memory controller

(33) BIOS modified flag (Incremented each time BIOS is modified)

Standard CMOS Setup

Remember young Jedi...

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 of the quality of the motherboard expect to loose (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.
- **Cyln:** 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 Cyln (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: $(Hds * Cyl * Sect * 512) / 1048$.

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 (Extended 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 (1 Gigabyte), but does not support LBA.

- **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 .

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). 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

Remember...

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 be 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 know 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 0x0FFFFFF 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

Remember...

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 an 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 an 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. Here are some appropriate settings:

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 at a time. 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.), Mode 4 (up to 16,7 Mb/sec.) and Mode 5 (up to 20 Mb/sec.). If this setting is not set properly, communication with COM ports may not work properly.
- **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.
- **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.
- **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 your 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.
 - **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 strobes. 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.
-

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 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 you 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 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 depend 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).
