



Computer Hardware Tuning

Hank Cullen

First Edition, 2012

ISBN 978-81-323-4113-0



© All rights reserved.

Published by:

White Word Publications

4735/22 Prakashdeep Bldg,

Ansari Road, Darya Ganj,

Delhi - 110002

Email: info@wtbooks.com

Table of Contents

Chapter 1 - Performance Tuning

Chapter 2 - Dynamic Frequency Scaling and Jumper (Computing)

Chapter 3 - Case Modding

Chapter 4 - Memory Divider and Modchip

Chapter 5 - Modding

Chapter 6 - Overclocking

Chapter 7 - Power Management and SpeedFan

Chapter 8 - SpeedStep and Tweaking

Chapter 9 - Underclocking

Chapter 10 - Random-Access Memory

Chapter 11 - Graphics Processing Unit

Chapter 12 - Hard Disk Drive

Chapter 1

Performance Tuning

Performance tuning is the improvement of system performance. This is typically a computer application, but the same methods can be applied to economic markets, bureaucracies or other complex systems. The motivation for such activity is called a performance problem, which can be real or anticipated. Most systems will respond to increased load with some degree of decreasing performance. A system's ability to accept higher load is called scalability, and modifying a system to handle a higher load is synonymous to performance tuning.

Systematic tuning follows these steps:

1. Assess the problem and establish numeric values that categorize acceptable behavior.
2. Measure the performance of the system before modification.
3. Identify the part of the system that is critical for improving the performance. This is called the bottleneck.
4. Modify that part of the system to remove the bottleneck.
5. Measure the performance of the system after modification.

This is an instance of the measure-evaluate-improve-learn cycle from quality assurance.

A performance problem may be identified by slow or unresponsive systems. This usually occurs because high system loading, causing some part of the system to reach a limit in its ability to respond. This limit within the system is referred to as a bottleneck.

A handful of techniques are used to improve performance. Among them are code optimization, load balancing, caching strategy, distributed computing and self-tuning.

Performance analysis

Performance analysis, commonly known as profiling, is the investigation of a program's behavior using information gathered as the program executes. Its goal is to determine which sections of a program to optimize.

A profiler is a performance analysis tool that measures the behavior of a program as it executes, particularly the frequency and duration of function calls. Performance analysis tools existed at least from the early 1970s. Profilers may be classified according to their output types, or their methods for data gathering.

Performance engineering

Performance engineering is the discipline encompassing roles, skills, activities, practices, tools, and deliverables used to meet the non-functional requirements of a designed system, such as increase business revenue, reduction of system failure, delayed projects, and avoidance of unnecessary usage of resources or work.

Several common activities have been identified in different methodologies:

- Identification of critical business processes.
- Elaboration of the processes in use cases and system volumetrics.
- System construction, including performance tuning.
- Deployment of the constructed system.
- Service management, including activities performed after the system has been deployed.

Code optimization

Enhancing performance by rewriting specific portions of a program to run faster is one form of code optimization. The term code optimization can refer to improving the implementation of a particular algorithm for performing a task (code tuning). It can also refer to utilizing a better algorithm. Examples of code optimization include improving the code so that work is done once before a loop rather than inside a loop or replacing a call to a simple selection sort with a call to the more complicated algorithm for a quicksort.

Caching strategy

Caching is a fundamental method of removing performance bottlenecks that are the result of slow access to data. Caching improves performance by retaining frequently used information in high speed memory, which reduces access time and thus improves performance. Caching is an effective manner of improving performance in situations where the principle of locality of reference The methods used to determine which data is stored in progressively faster storage are collectively called **caching strategies**.

Load balancing

A system can consist of independent components, each able to service requests. If all the requests are serviced by one of these systems (or a small number) while others remain idle then time is wasted waiting for used system to be available. Arranging so all systems are used equally is referred to as load balancing and can improve over-all performance.

Load balancing is often used to achieve further gains from a distributed system by intelligently selecting which machine to run an operation on based on how busy all potential candidates are, and how well suited each machine is to the type of operation that needs to be performed.

Distributed computing

Distributed computing is used to increase the performance of operations that can be performed in parallel, by concurrently executing multiple operations. Operations may be distributed across multiple processes on a single CPU, taking advantage of multitasking, multiple processes across multiple CPUs, or across multiple machines. As operations are executed concurrently, ensuring synchronization between processes is essential to ensure correct results.

As the trend of increasing the potential for parallel execution on modern CPU architectures continues, the use of distributed systems is essential to achieve performance benefits from the available parallelism. High performance cluster computing is a well known use of distributed systems for performance improvements.

Distributed computing and clustering can negatively impact latency while simultaneously increasing load on shared resources, such as database systems. To minimize latency and avoid bottlenecks, distributed computing can benefit significantly from distributed caches.

Self-tuning

A self-tuning system is capable of optimizing its own internal running parameters in order to maximize or minimize the fulfillment of an objective function; typically the maximization of efficiency or error minimization. Self-tuning systems typically exhibit non-linear adaptive control. Self-tuning systems have been a hallmark of the aerospace industry for decades, as this sort of feedback is necessary to generate optimal multi-variable control for nonlinear processes.

Bottlenecks

The bottleneck is the part of a system which is at capacity. Other parts of the system will be idle waiting for it to perform its task.

In the process of finding and removing bottlenecks, it is important to prove their existence, such as by sampling, before acting to remove them. There is a strong temptation to *guess*. Guesses, by definition, are often wrong, and investing time in them is itself a bottleneck.

Chapter 2

Dynamic Frequency Scaling and Jumper (Computing)

Dynamic frequency scaling

Dynamic frequency scaling (also known as **CPU throttling**) is a technique in computer architecture whereby the frequency of a microprocessor can be automatically adjusted "on-the-fly", either to conserve power or to reduce the amount of heat generated by the chip. Dynamic frequency scaling is commonly used in laptops and other mobile devices, where energy comes from a battery and thus is limited. It is also used in quiet computing settings and to decrease energy and cooling costs for lightly loaded machines. Less heat output, in turn, allows the system cooling fans to be throttled down or turned off, reducing noise levels and further decreasing power consumption. It is also used for reducing heat in insufficiently cooled systems when the temperature reaches a certain threshold, such as in poorly cooled overclocked systems.

The dynamic power (*switching power*) dissipated by a chip is $C \cdot V^2 \cdot f$, where C is the capacitance being switched per clock cycle, V is voltage, and f is the switching frequency (as a unitless quantity). As frequency changes, the dynamic power will change linearly with it. Dynamic power does not account the total power of the chip however, as there is also static power, which is primarily due to various leakage currents. Leakage current has become more and more important as feature sizes has become smaller and threshold levels lower. In state-of-the-art deep submicrometre technologies in 2008, dynamic power accounts for approximately two-thirds of the total chip power, which limits the effectiveness of frequency scaling.

Dynamic voltage scaling is another power conservation technique that is often used in conjunction with frequency scaling, as the frequency that a chip may run at is related to the operating voltage.

The efficiency of some electrical components, such as voltage regulators, decreases with increasing temperature, so the power used may increase with temperature. Since

increasing power use may increase the temperature, increases in voltage or frequency may increase system power demands even further than the CMOS formula indicates, and vice-versa.

Performance impact

Dynamic frequency scaling reduces the number of instructions a processor can issue in a given amount of time, thus reducing performance. Hence, it is generally used when the workload is not CPU-bound.

Dynamic frequency scaling by itself is rarely worthwhile as a way to conserve switching power. Saving the most power requires dynamic voltage scaling too, because of the V^2 component and the fact that modern CPUs are strongly optimized for low power idle states. In most constant-voltage cases it is more efficient to run briefly at peak speed and stay in a deep idle state for longer (called "race to idle"), than it is to run at a reduced clock rate for a long time and only stay briefly in a light idle state. However, reducing voltage along with clock rate can change those tradeoffs.

A related-but-opposite technique is overclocking, whereby processor performance is increased by ramping the processor's (dynamic) frequency beyond the manufacturer's design specifications.

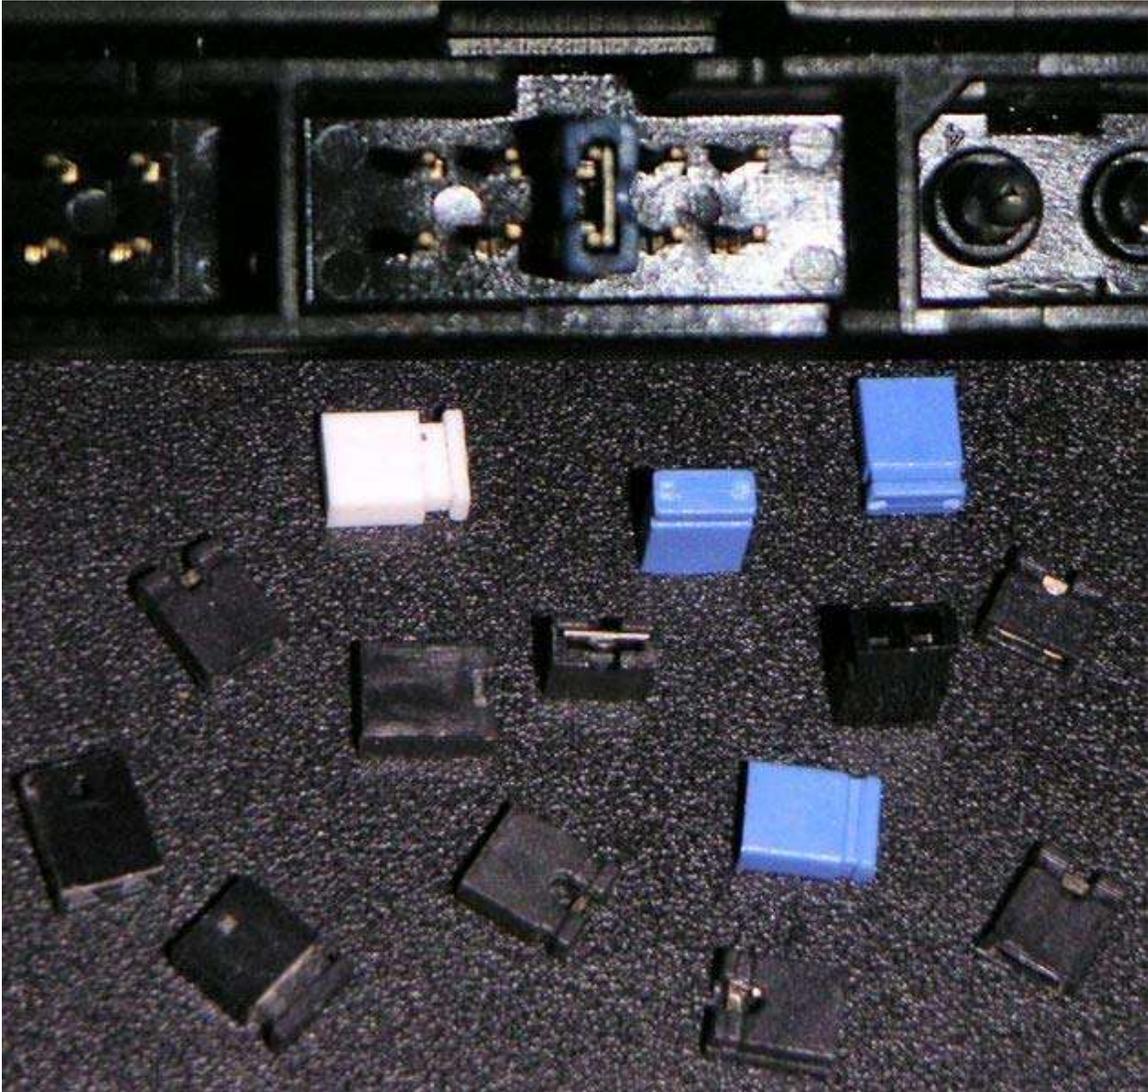
One major difference between the two is that in modern PC systems overclocking is mostly done over the Front Side Bus (mainly because the multiplier is normally locked), but dynamic frequency scaling is done with the multiplier. Moreover, overclocking is often static, while dynamic frequency scaling is always dynamic. Software can often incorporate overclocked frequencies into the frequency scaling algorithm, if the chip degradation risks are allowable.

Implementations

Intel's CPU throttling technology, SpeedStep, is used in its mobile CPU line. AMD employs two different CPU throttling technologies. AMD's Cool'n'Quiet technology is used on its desktop and server processor lines. The aim of Cool'n'Quiet is not to save battery life, as it is not used in AMD's mobile processor line, but instead with the purpose of producing less heat, which in turn allows the system fan to spin down to slower speeds, resulting in cooler and quieter operation, hence the name of the technology. AMD's PowerNow! CPU throttling technology is used in its mobile processor line, though some supporting CPUs like the AMD K6-2+ can be found in desktops as well.

According to the ACPI Specs, the C0 working state of a modern-day CPU can be divided into the so called "P"-states (performance states) which allow clock rate reduction and "T"-states (throttling states) which will further throttle down a CPU (but not the actual clock rate) by inserting STPCLK (stop clock) signals and thus omitting duty cycles.

Jumper (computing)



Top: jumper block on IDE hard drive with jumper; **bottom:** assorted jumpers

In electronics and particularly computing, a **jumper** is a short length of conductor used to close a break in or bypass part of an electrical circuit. Jumpers are typically used to set up or adjust printed circuit boards, such as the motherboards of computers.

Description

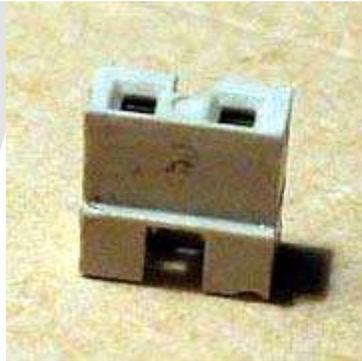
Jumper pins (points to be connected by the jumper) are arranged in groups called *jumper blocks*, each group having at least one pair of contact points and often more. Sometimes these groups are referred to as *headers*. In general, each contact in a jumper block terminates in a small metal pin. An appropriately sized conductive sleeve called a jumper, or more technically, a jumper shunt, is slipped over the pins to complete the circuit.

Jumpers must be electrically conductive; they are usually encased in a non-conductive block of plastic for convenience. This also avoids the risk that an unshielded jumper will accidentally short out something critical (particularly if it is dropped on a live circuit).

When a jumper is placed over two or more jumper pins, an electrical connection is made between them, and the equipment is thus instructed to activate certain settings accordingly. For example, with older PC systems, CPU speed and voltage settings were often made by setting jumpers. Informally, technicians often call setting jumpers "strapping". To adjust the SCSI ID jumpers on a hard drive, for example, is to "strap it up".

Jumper blocks and jumpers are also often used on motherboards to clear the CMOS information, resetting the BIOS configuration settings. This allows the computer to boot if a recent BIOS setting made it unable to boot, or if the CMOS boot password was forgotten.

Move to reduce jumpers



a jumper

Early generations of any given computer hardware technology usually have many jumper blocks, often laid out in a way that is poorly documented and difficult to set correctly. Often, designers find ways to streamline and simplify the jumper layout. For example, a typical early model Intel 386 motherboard might have 30 or 40 jumper pairs, while the last production models typically had just a handful, or sometimes only one. Typically, each jumper block is assigned and labeled with a number, which is documented in an instructional list printed on the motherboard or in the manual.

The recent trend has been to try to eliminate jumpers entirely from hardware devices by the use of auto-configuration or software-controlled configuration. Configurations may be stored in NVRAM, loaded by a host processor, or negotiated at system initialization time. In some cases, hot swappable devices may be able to renegotiate their configuration while the system is running. Jumperless designs have the advantage that they are usually fast and easy to set up, often require little technical knowledge, and can be adjusted without having physical access to the circuit. With newer PCs, the most common use of jumpers is in setting the operating mode for ATA drives (master, slave, or cable select).

Systems using boards with physical jumpers, on the other hand, tend to be configured correctly by end users as, in general, non-technical people are less willing to physically alter hardware settings than they are to experiment with settings from the keyboard. They also have the advantage that they usually only need to be set once; while firmware settings can be easily lost or corrupted by a careless user, a virus, or a power failure, the only way to alter a correct jumper setting is to physically change it.

WWT

Chapter 3

Case Modding



Modded PC case with CCFL lights and switches, LED fans, perspex window, custom fan hole, and customized Windows XP.

Case modification (commonly referred to as **case modding** where an individual project is referred to as a **case mod**) is the modification of a computer chassis (often just referred to as the case), or a video game console chassis. Modifying a computer case in any non-standard way is considered a case mod. Many people, particularly hardware enthusiasts, use case mods to illustrate a computer's power (by showing off the internal hardware),

and also for aesthetic purposes. Cases may also be modified to improve a computer's performance.

History

When personal computers first became available to the public, the majority were produced in simple, beige-colored cases. This design is sometimes referred to be as a *beige box*. Although this met the purpose of containing the components of the personal computer, many users considered their computers as "tacky" or "dull", and some began modifying their existing chassis, or building their own from scratch.

A new market for third-party computer cases and accessories began to develop, and today cases are available in a wide variety of colors and styles. Today the business of "modding" computers and their cases is a hugely profitable endeavor, and modding competitions are everywhere.

WWT

Common modifications



A computer modded with lights

- **Window mods:** This refers to a window placed within one of the panels of a computer case. This is most often done to the left hand side panel, and less often to the top panel. This modification is so popular that many of the major case manufacturers now offer cases with the windows pre-installed, or replaceable side panels with a window installed. Some companies even offer entire cases made out of transparent materials. A window kit may be modified to hold an LCD screen. Laser engraving can be done on acrylic windows to add a distinct look to a modded case.

- **Lighting mods:** A lighting mod refers to lighting in or on the computer cases. This is usually achieved with cold cathode lights (CCLs), LED case fans, or electroluminescent wire lights. The lights are sometimes paired with sound controllers that make the lights pulse in time to sound. CCLs come in long tubes and generally produce a little bit of heat. LEDs come in many sizes and forms, most often seen in bars similar to CCLs or within fans, called LED fans. Electroluminescent wire, which takes the form of a small light rope, is often embedded in cables such as SATA cables.

Lighting modifications are often paired with window mods to help show off the components. Although not as common, they are also placed in cases without a window, allowing light to shine through any holes or gaps of the case which add subtle aesthetics to an otherwise plain looking case; this is also done for practical purposes such as to make a computer double as a night light. Internal components such as case fans, CPU heatsink fans, and power supplies are often available with built in lighting.

- **Cooling mods:** There are many modifications that can fall into this category. The most common one is simply drilling a mount for a new fan, or removing a restrictive fan grill. Others include air ducts, water cooling, filtering, sealing openings to promote better air flow, or even the adding of a tank of pressurized carbon dioxide or liters of mineral oil to the case. These modifications are often performed by overclockers either looking for better cooling for hot components or sound reduction. Some fan modifications are merely a show of modding skill or talent and have no true functional purpose.
- **Spray paint:** Painting a case is another method of distinguishing your work from others. Spray paint is the common method preferred among amateur modders. There are many spray painting guides for amateur modders. The finish cannot be compared to automotive paint or powder coating, but is a simple way to change the look of a case.
- **Vinyl dye:** Re-colouring the plastics of a case, keyboard, mouse and speakers is another method of highlighting your system and making it different than the rest. Vinyl dye has the advantages of ease of use like spray paint, but is much more durable as it doesn't chip or scratch off, it also doesn't create a 'layer' like paint.

Less common modifications



A computer modded with different coloured CCFLs and LED fans.

- **Automotive paint & other finishes:** Automotive paint refers to the paint typically seen on cars and trucks. This type of finish requires a compressed air source, such as an air compressor or CO2 tank, and a spray gun. It is more expensive than a finish using spray cans, but when done skillfully it can be better looking and much more durable. Other methods of painting can include powder coating which is highly durable though not quite as aesthetically pleasing to many modders as automotive paint. Electroplating can also be done on steel computer cases and parts. Aluminum cases can be plated or anodized as well, and processes are available to plate plastic cases. Plated coatings can range from nickel to chrome and even gold. More elaborate finishes can be crafted by using a combination of techniques, such as chrome plating with a transparent color coat.
- **Body filler:** Body filler (or Bondo) is a two-part putty often used to fix dents in automobiles. Case modders use it to fill and sculpt their own creations. When mixed with a paste catalyst the filler hardens in a short period of time and can be sanded, ground or cut to a desired shape. An alternative system uses fiberglass resin (catalyzed with liquid hardener) and either fiberglass cloth or mat to fill holes and form shapes. Lacquer based Spot Putty is often used to fill smaller imperfections before the application of primer. Typically, a case modder uses a combination of these materials to obtain the desired result. This method is usually used on the front plastic bezel of a computer case to give it a new look.

Types of case mods



Computer built into a Microwave oven

- **Peripheral mods:** Peripherals like the keyboard, mouse, and speakers are sometimes painted or otherwise modified to match the computer. Some system builders, in an effort to make their system more portable and convenient, install speakers and small LCD screens into the case.
- **Unusual cooling mods:** Hardcore overclockers often install cooling systems for the sole purpose of achieving performance records. Such systems frequently include phase change, thermoelectric/Peltier and liquid nitrogen. However, some of these systems are noisy and expensive. They are rarely used for extended periods of time.
- **Case building:** Sometime modders even build entire cases from scratch. Some may attempt to treat the case as a work of art. Others make it appear to be something else, like a teddy bear, wooden cabinet, or a shelf mounted on a wall. Still others pursue a retro look, like a Macintosh Plus or an old Atari 2600 video game console. Case modders (or case builders) who create their computer cases from scratch are few and far between. These people sometimes put hundreds of hours into their work. The WMD case, Project Nighthawk, and Dark Blade case are a few examples of professional cases built from scratch.
- **Component modding:** This type of modding, as the name suggests, involves modifying the PC components themselves. An example is the relocation of buttons on optical drives. This is often done in combination with "stealthing", which hides the drive's visibility by masking it with a blank face. A riskier modification involves installing hard drive windows. This is done in a clean room

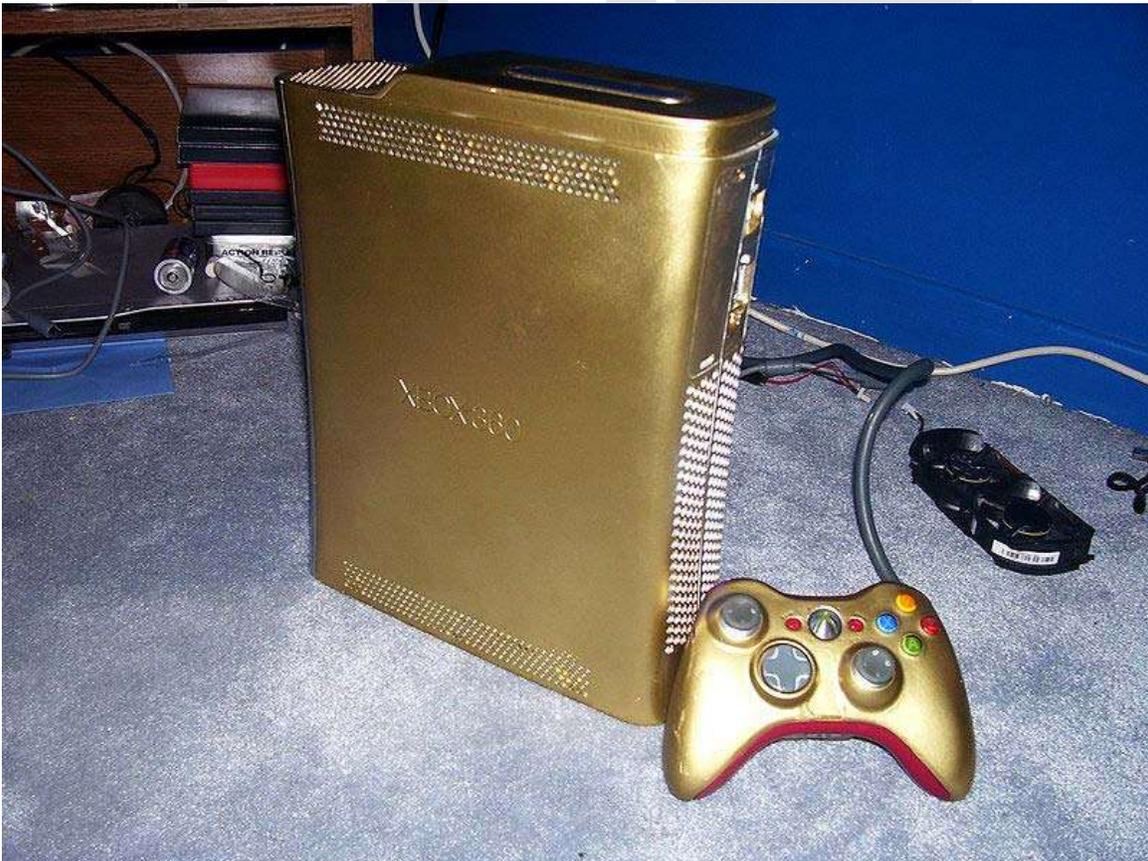
where there is little to no dust. Few people have attempted it and results seem to vary. Some hard drives, including the WD Raptor, now come with a window as standard.

- **Laptop modding:** Laptops can be modified much like a typical computer case. While most laptop mods consist of new paint or other finishes, others have chosen to engrave or cut out designs into their laptop cover (behind the screen). Laptops may also be turned into digital photo frames. These types of mods will typically void the warranty of the device. To avoid warranty issues, skins or stickers can be purchased that are easily removable from the casing.

Case modding contests

Many websites and companies run contests for case modders, awarding prizes and accolades to the winners. Some of these contests are sponsored by computer enthusiast magazines and others by computer retailers.

Console case modding



An example of an Xbox 360 case mod

Console case modding refers to gamers who modify the case of their game consoles. The most common consoles to modify are the Xbox and Xbox 360, because there is much

more room inside to customize them with items such as lights, fans, etc. These consoles and their controllers are also relatively easy to take apart. For those who do not wish to scratch-build mods, there are several companies that sell transparent Xbox cases and various cooling/lighting equipment for them.

Console case modding started in the late 1980s when the NES and Sega Genesis, were released; many customers simply put pictures or stickers on them until the PlayStation was released. Many case modders started to change hardware, for example by altering them to play copied games (known as 'chipping' the games console). The most common modification for the PlayStation was the 'chipping' process (mentioned above). When the Nintendo 64, Dreamcast and PlayStation 2 was released, many chipped them, styled them, and added additional cooling; some went as far as changing the hardware itself. When the Xbox and Xbox 360 were released, many modders personally customized them further, using neon lights, transparent cases, fans, and PC hard drives (as opposed to Xbox-branded drives). Many modders found that altering the interior of Xbox 360s was difficult due to absence of a power cable (normally in a PC, this cord attaches the hard disk drive to its motherboard). Despite shortcomings, modders also found a way to power neon lighting and other powered equipment by using the DVD-ROM power supply; however, due to insufficient power to the hard disk drive, it often caused freezing during disk access. Fortunately, the most recent products available for internal case modding uses the power outlet for the internal fan by splitting the cord with a "Y" connector.

Chapter 4

Memory Divider and Modchip

Memory divider

A **memory divider** is a ratio which is used to determine the operating clock frequency of computer memory in accordance with front side bus (FSB) frequency, if the memory system is dependent on FSB clock speed. Along with memory latency timings, memory dividers are extensively used in overclocking memory subsystems to find stable, working memory states at higher FSB frequencies. A memory divider is also commonly referred to as "DRAM:FSB ratio".

Memory dividers are only applicable to those chipsets in which memory speed is dependent on FSB speeds. Certain chipsets like nVidia 680i have separate memory and FSB lanes due to which memory clock and FSB clock are asynchronous and memory dividers are not used there. Setting memory speeds and overclocking memory systems in such chipsets are different issues which do not use memory dividers.

Overview

Memory Dividers allows system memory to run slower than or faster than the actual FSB (Front Side Bus) speed. Ideally, Front Side Bus and system memory should run at the same clock speed because FSB connects system memory to the CPU. But, it is sometimes desired to run the FSB and system memory at different clock speeds. It is possible to run FSB and memory clock at different clock speeds, within certain limits of the motherboard and corresponding chipset. So, settings termed as Memory Divider or FSB/DRAM settings are available and are expressed in a "ratio" which control the difference in memory clock rate and FSB speed.

Entry Level motherboards usually do not provide memory dividers to be changed and the memory dividers are managed by Memory Controller (if chipset supports memory dividers). High end motherboards meant for overclocking provide facilities to change memory dividers (if chipset support memory dividers). However, in certain chipsets

memory dividers are not used, because in those systems memory speed is independent of FSB speed.

Description & Application

Usually $(\text{Memory Divider}) \times (\text{Front Side Bus Frequency})$ gives I/O Bus clock of the memory. Memory clock then determines the final operating frequency or effective clock speed of memory system depending upon DRAM types (DDR, DDR2 and DDR3 SDRAM).

By default, FSB speed and memory are usually set to a 1:1 ratio, meaning that increasing FSB speed (by overclocking) increases memory speed by the same amount. Normally system memory is not built for overclocking and thus may not be able to take the level of overclocking that the processor or motherboard can achieve. The memory divider allows users to mitigate this problem by reducing the speed increase of the memory relative to that of the FSB and the processor.

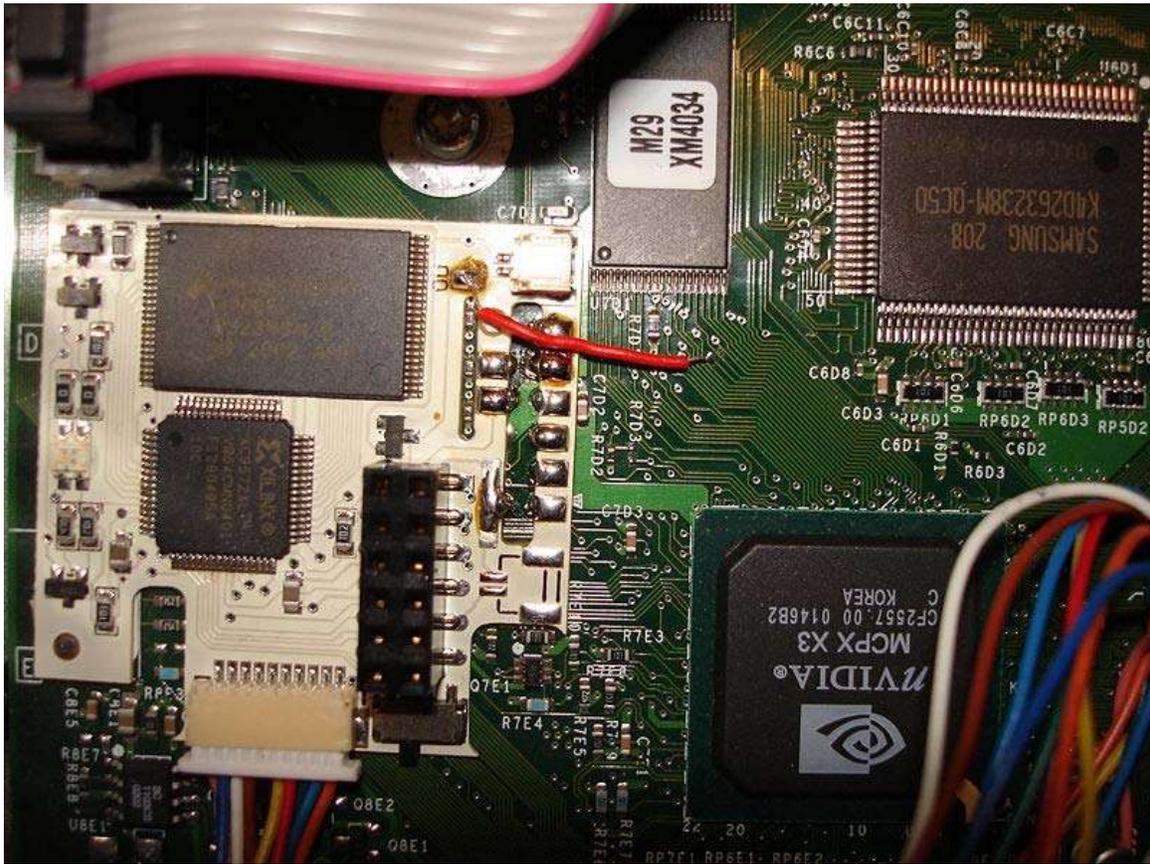
Example

Suppose a computer system has DDR memory, a Memory Divider of 1:1, a FSB operating at 200 MHz and a CPU multiplier of 10x. Then, the base memory clock will operate at $(\text{Memory Divider}) \times (\text{FSB}) = 1 \times 200 = 200$ MHz and the effective memory clock would be 400 MHz since it's a DDR system ("DDR" stands for Double Data Rate; the effective memory clock speed is double the actual clock speed). The CPU will operate at $10 \times 200 \text{ MHz} = 2.0 \text{ GHz}$. Using I/O bus frequency of 200 MHz various types of DRAM will operate as:

DDR SDRAM at 400 MHz (DDR-400 or PC-3200)
DDR2 SDRAM at 800 MHz (DDR2-800 or PC2-6400)
DDR3 SDRAM at 1600 MHz (DDR3-1600 or PC3-12800)

Now suppose that we overclock FSB to 250 MHz so that CPU operates at $10 \times 250 \text{ MHz} = 2.5 \text{ GHz}$ and memory clock operates at 250 MHz (Memory Divider \times FSB). Since DDR-400 RAM is used then effective memory clock (Actual Memory Frequency) will be 500 MHz. A normal DDR-400 SDRAM will fail to work at 500 MHz since it is designed to work at 400 MHz and system will become unstable. But a modern CPU (having overclocking potential) can work at 2.5 GHz (even if it is designed to work at 2 GHz) flawlessly without giving any problem of stability. To keep running overclocked CPU at 2.5 GHz or even at higher speeds (by increasing FSB) we need to slow down memory clock so as to achieve a stable system. For this if we decrease DRAM:FSB ratio to say 4:5 then resulting memory clock speed is $(4/5) \times 250 \text{ MHz} = 200 \text{ MHz}$ resulting effective clock speed of 400 MHz on DDR-400. So, we are able to operate upon a stable overclocked CPU at 2.5 GHz from 2 GHz without increasing effective memory clock.

Modchip



A Modchip, the white circuit board, attached to an Xbox's PCB

A **modchip** (short for *modification chip*) is a small electronic device used to modify or disable built-in restrictions and limitations of computers, specifically videogame consoles. It introduces various modifications to its host system's function, including the circumvention of region coding, digital rights management, and copy protection checks for the purpose of running software intended for other markets, copied game media, or unlicensed third-party (homebrew) software.

Modchips are mainly used in CD/DVD-based videogame systems due to the availability and low cost of user-writable media. In addition to games consoles, modchips are also available for some DVD players to circumvent region code enforcement and user operation prohibitions.

Function and construction

Modchips operate by replacing or overriding a system's protection hardware or software. They achieve this by either exploiting existing interfaces in an unintended or

undocumented manner, or by actively manipulating the system's internal communication, sometimes to the point of re-routing it to substitute parts provided by the modchip.

Most modchips consist of one or more integrated circuits (microcontrollers, FPGAs, or CPLDs), often complemented with discrete parts, usually packaged on a small PCB to fit within the console system it is designed for. Although there are modchips that can be reprogrammed for different purposes, most modchips are designed to work within only one console system or even only one specific hardware version.

Modchips typically require some degree of technical acumen to install. Modchips must be connected to a console's circuitry, most commonly by soldering wires to select traces or chip legs on a system's circuit board. Some modchips allow to be installed by directly soldering the modchip's contacts to the console's circuit ("quicksolder"), by the precise positioning of electrical contacts ("solderless"), and, in rare cases, by plugging it into a system's internal or external connector.

Memory cards or cartridges that offer functions similar to modchips work on a completely different concept, namely by exploiting flaws in the system's handling of media. Such devices are not referred to as modchips, even if they are frequently traded under this umbrella term.

History

Cartridge-based console systems did not have modchips produced for them. They usually implemented copy protection and regional lockout with game cartridges, both on physical and software level. Converters or passthrough devices have been used to circumvent the restrictions, while flash memory devices (game backup devices) were widely adopted in later years to copy game media. Early in the transition from solid-state to optical media, CD-based console systems did not have regional market segmentation or copy protection measures due to the rarity and high cost of user-writable media at the time.

Modchips started to surface with the PlayStation system, due to its popularity and the increasing availability and affordability of CD writers. At the time, a modchip's sole purpose was to allow the use of imported and copied game media.

Today, modchips are available for practically every current console system, often in a great number of variations. In addition to circumventing regional lockout and copy protection mechanisms, modern modchips may introduce more sophisticated modifications to the system, such as allowing the use of user-created software (homebrew), expanding the hardware capabilities of its host system, or even installing an alternative operating system to completely re-purpose the host system (*e.g.* for use as a home theater PC).

Anti-modchip measures

Most modchips open the system to copied media, therefore the availability of a modchip for a console system is undesirable for console manufacturers. They react by removing the intrusion points exploited by a modchip from subsequent hardware or software versions, changing the PCB layout the modchips are customized for, or by having the firmware or software detect an installed modchip and refuse operation as a consequence. Since modchips oftentimes hook into fundamental functions of the host system that cannot be removed or adjusted, these measures may not completely prevent a modchip from functioning but only prompt an adjustment of its installation process or programming, e.g. to include measures to make it undetectable ("stealth") to its host system.

With the advent of online services to be used by video game consoles, some manufacturers have executed their possibilities within the service's license agreement to ban consoles equipped with modchips from using those services.

Some console manufacturers included the option to run homebrew software or even an alternative operating system on their consoles, however, these features have mostly been withdrawn again. An argument can be made that a console system remains largely untouched by modchips as long as their manufacturers provide a legitimate way of running unlicensed third-party software.

Legality

One of the functions of many modchips—the circumvention of copy protection mechanisms—is outlawed by many countries' copyright laws such as the DMCA in the USA, the EU CD and its various implementations by the EU member countries, and the Australian Copyright Act. However, due to the many diversified functions of a modchip, other laws may apply to a modchip as well, allowing specific functions (e.g. the circumvention of region coding under Australian law).

The nonuniform interpretation of applicable law by the courts and constant profound changes and amendments to copyright law do not allow for a definitive statement on the legality of modchips. Because of the ambiguity of applicable laws, a modchip's legality under a country's legislature may only be individually asserted in court.

Most of the very few cases that have been brought before a court have ended with the conviction of the modchip merchant or the manufacturer under the countries' anti-circumvention laws. A small number of cases in the United Kingdom and Australia have been dismissed under the argument that a system's copy protection mechanism would not be able to prevent the actual infringement of copyright (which would be the actual process of copying game media), therefore it cannot be considered an effective TPM protected by anti-circumvention laws. Australian copyright law has since been amended to effectively close this legal loophole.

Chapter 5

Modding

Modding is a slang expression that is derived from the verb "modify". Modding refers to the act of modifying a piece of hardware or software or anything else for that matter, to perform a function not originally conceived or intended by the designer. The term modding is often used within the computer game community, particularly in regard to creating new or altered content and sharing that via the web. It may also be applied to the overclocking of computers in order to increase the frequency at the which the CPU operates. Case modding is also a popular activity amongst many computer enthusiasts which involves the customization of a computer chassis or the installation of water cooling technology. In connection with automobiles, modding often refers to engine tuning, remapping of a vehicle's engine control unit or customization of the bodywork.

Computers and digital equipment

Legal issues

Modding may sometimes infringe the legal rights of the copyright owner. Some nations have laws prohibiting modding and accuse modders of attempting to overcome copy protection schemes. In the United States, the DMCA has set up stiff penalties for mods that violate the rights of intellectual property owners. In the European Union, member states have agreed the EU Copyright Directive and are transposing it into national law. A man was convicted in the United Kingdom in July 2005 for selling a modded Xbox with built in software and games. However it is also worthy of note that some other European countries have not interpreted the legal issues in the same way. In Italy a judge threw out a Sony case saying it was up to owners of a console what they did with it. Similarly in Spain, mod chips are seen as legal despite the EU copyright legislation. Modding may be an unauthorized changed made to a software or hardware to a platform in gaming. Case mods are modifications to a device with the altering of certain styles. For example, people who mod a Microsoft Xbox 360 can alter the led lights on the controller to glow different colors.

Multi-user licensing

Computer systems, hardware, and software are often sold or licensed to one home machine, to a business that has many computers, a government agency, or a non-profit organization. When the software license says that it is for a specific person, then it is not legal for that software to be used by some other person on that same computer, even a member of the same family, or another employee of the same company. But this strict licensing is only one approach. In this form of licensing, for more than one person to be using that software or hardware, they need to have a multi-user license that usually dictates how many different people may use it.

Derivative software

Some software is licensed with a copy of the program source code supplied along with the executable code in which the license specifically authorizes changes to the supplied software. This is a common standard in business software packages. Hundreds of thousands of computer programmers in some nations have jobs because businesses want the purchased software tailored to the specific needs of the individual businesses. Most every major city has want ads in the newspaper where there are job openings for people to modify some company's computer systems, where the ad specifies what programming languages or operating systems the applicant needs to know.

Derivative software is licensed copyrighted software used in whole or part in new programs written by programmers at businesses using the licensed software. Programmers copy the copyright notices into the source code where the code was copied, and track all such places, because if the license is permitted to expire, then the business loses software use rights, including any place to which it was copied. An annual fee is typically paid to keep the license in effect, and over time, the software supplier can increase the fees to the point that the business chooses to convert to some other commercial software that seems to be more cost effective. It is not unusual in business software to find programs that have many different copyright notices, each referring to different sources of the derived source code.

Video game consoles

A common example of one kind of modding is video game console mod chips, which can allow users to play homemade games, games legitimately purchased in other regions, or legal backup copies, but can also allow illegal unauthorized copies by allowing the player to play personally recorded CD or DVD copies of video games. Modchips, in their current form, were first available for the Sony PlayStation (and later the PlayStation 2). Various other types of copyright circumvention systems also existed for the Nintendo 64 and the older Game Boy consoles (though neither include actual modding, but instead backup devices).

Types of modding

There are two different ways of running unsigned code on a game console. One is through soft modding (modifying software, normally using a softmod) to allow the user to change data contained on its hard drive in the case of the Xbox. Another type of modding, known as hard modding, exploits the BIOS of the console to run unsigned code, or games. This form of 'modding' (more correctly termed as hacking) is very popular as it is able to 'run' many different types of software. But soft modding is even more popular because of its ease of installation and its relatively low price (it can even be done for free with the right tools).

Another type of console modding is about appearances, much like computer case modification. Which includes, adding lights (most likely LEDs, cathodes or other electro-luminescent lighting). Cutting the game system case, to fit hardware and/or expose the internal systems. Cooling is a large part of console hard 'modding', including: heat sink upgrades, more powerful or quieter fans, some even go so far as to abandon common heat exchange to air all together by liquid a console (most notably in the Xbox 360, which initially had some heat problems).

Game software

On the other side, some companies actively encourage modding of their products. In cases such as TiVo and Google, there has been an informal agreement between the modders and the company in which the modders agree not to do anything that destroys the company's business model and the company agrees to support the modding community by providing technical specifications and information. Some commercial video games thrive through a modding community. In the case of *Half-Life*, a mod called *Counter-Strike* drove sales of the original software for years.

Many games, such as *The Elder Scrolls series*, come with a mod editing tool that allows users to create original content for themselves and others. Other games provide the source code for users to use in experimenting and creating. Still others, like *18 Wheels of Steel*, will provide the non-programmed data (images, small codepieces and the like) in a simple archive, which can often be opened by renaming it to a .zip file. Often modders will take the game in directions that the developers never anticipated or didn't have time/funding to include. Generally, a small percent of game players will spend much time mod making, but those who do usually develop communities around modding a particular game. Communities are generally connected via a web forum where new modders can ask questions of more experienced ones, and everyone can find inspiration in the work of others. Some games, like *Neverwinter Nights*, could never have been as successful as they are without a thriving user community. And as more people have been more connected via the web, this has become a vital and dynamic creative phenomena where users become content creators not just content consumers.

Skilled computer users who are able to crack data formats and reverse engineer a game can modify them to their heart's content, because the creator of the software has copyright

authority over who may use it, or change it. Software is sold with a license that spells out what guarantees, if any, come with the software, and what rights the purchaser has to change the software. Many people do not read these contracts, or store them in such a way to be able to prove what contract came with what purchase, so some computer users are ignorant to what their rights are with respect to backing up software, modifying software, and sharing it with other potential users.

On August 5, 2009 Matthew Crippen, a student at California State University, was arrested for modifying game consoles such as Xbox, Playstation, and Nintendo Wii for profit. According to him it was so that the owners could play their backup discs of games they legally own. However, according to the DMCA, it is illegal to circumvent copyright protection software, even for non-pirating uses such as backing up legally owned games.

Device drivers

Modded drivers are made for improved performance which official versions of drivers do not offer or in cases where there are no official versions of drivers for new hardware designed for older operating systems such as Windows 98.

Computer hardware

Case modding may range from simple case painting to full blown case mod with cooling mods and fabricated pieces.

Overclocking may also be termed as 'modding', and the overclocking of a graphics card using driver software to gain the performance of a more expensive model is known as 'soft-modding'. Volt modding is a term in which jumpers and rheostats are used to mod a hardware's voltages to yield better overlocks.

Cars and vehicles

Orthopedic

Ortho-modding is the car adaptation (seats, pedals, etc.) to help drivers to prevent, correct and diminish light orthopedic and backbone/spine problems.

Eco-modding

Eco-modding is the reduction of drag, petroleum car adaptation to use renewable energy (generally, changing or adding a new engine or motor), generally hydrogen or electricity. Occasionally, it has been known to run a Diesel engine on plant and animal oils.

Performance tuning

Car and engine tuning are modification of a car for enhanced performance, in engine power output, handling, aerodynamics, or cosmetic purposes.

Industrial machines

Factories get rather expensive machines that are used to mass produce specialized parts. These machines can be altered to make parts other than how the manufacturers of the machines designed or intended them. The legality of doing this depends on who owns the machines, and whether the agreement, that supplied the machines to the factories, said anything about this, and what the laws are in the nation where this is being done.

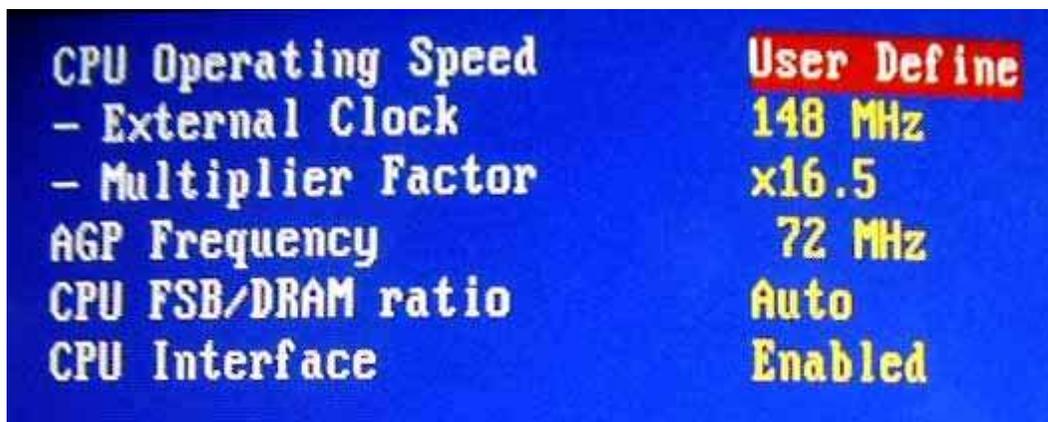
For example, the machines might be leased from the manufacturer of the machines. If they are ever to be returned, they need to go back in the same kind of condition and engineering shape as when they were first delivered. There is an annual physical inventory to make sure the factory has everything that they are leasing. This audit might be done by representatives of the leasing company, who are looking to see recognizable machines, that match their models and safety rules.

Pens

Pen-modding is the act of combining many pen parts either to help with pen spinning, in which a perfect balance is desired to create an ideal spinning pen, or simply for decoration. These pen mods can either be made by combining parts from different pens and/or mechanical pencils, or by buying modded pens online. In some cases, pen mods can exceed over \$30–40 USD per pen. Recently, the practice of pen modding has grown dramatically in popularity, with several mod brands appearing, and multiple online communities dedicated to pen modding and spinning.

Chapter 6

Overclocking



CPU Operating Speed	User Define
- External Clock	148 MHz
- Multiplier Factor	x16.5
AGP Frequency	72 MHz
CPU FSB/DRAM ratio	Auto
CPU Interface	Enabled

AMD Athlon XP overclocking BIOS setup on ABIT NF7-S. Front side bus frequency (external clock) has increased from 133 MHz to 148 MHz, and the clock multiplier factor has changed from 13.5 to 16.5

Overclocking is the process of running a computer component at a higher clock rate (more clock cycles per second) than it was designed for or was specified by the manufacturer, usually practiced by enthusiasts seeking an increase in the performance of their computers. Some purchase low-end computer components which they then overclock to higher clock rates, or overclock high-end components to attain levels of performance beyond the specified values. Others overclock outdated components to keep pace with new system requirements, rather than purchasing new hardware.

People who overclock their components mainly focus their efforts on processors, video cards, motherboard chipsets, and RAM. It is done through manipulating the CPU multiplier and the motherboard's front-side bus (FSB) clock rate until a maximum stable operating frequency is reached, although with the introduction of Intel's new X58 chipset and the Core i7 processor, the front side bus has been replaced with the QPI (Quick Path Interconnect); often this is called the Baseclock (BCLK). While the idea is simple, variation in the electrical and physical characteristics of computing systems complicates the process. CPU multipliers, bus dividers, voltages, thermal loads, cooling techniques

and several other factors such as individual semiconductor clock and thermal tolerances can affect it.

$$\text{FSB} \times \text{Multiplier} = \text{Frequency}$$

Considerations

There are several considerations when overclocking. First is to ensure that the component is supplied with adequate power to operate at the new clock rate. However, supplying the power with improper settings or applying excessive voltage can permanently damage a component. Since tight tolerances are required for overclocking, only more expensive motherboards—with advanced settings that computer enthusiasts are likely to use—have built-in overclocking capabilities. Motherboards with fewer features, such as those found in original equipment manufacturer (OEM) systems, often do not support overclocking.

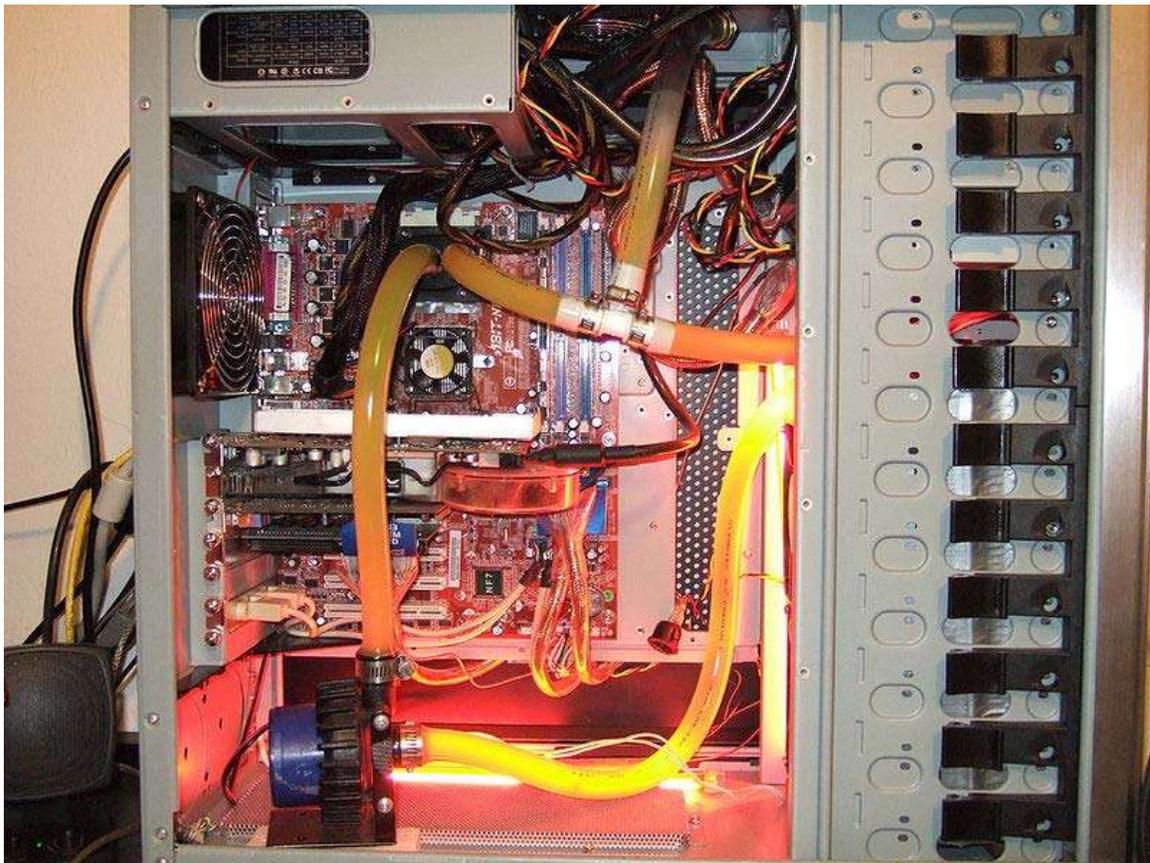
Cooling



High quality heat sinks are often made of copper.

All electronic circuits produce heat generated by the movement of electrical current. As clock frequencies in digital circuits and voltage applied increase, the heat generated by components running at the higher performance levels also increases. The relationship

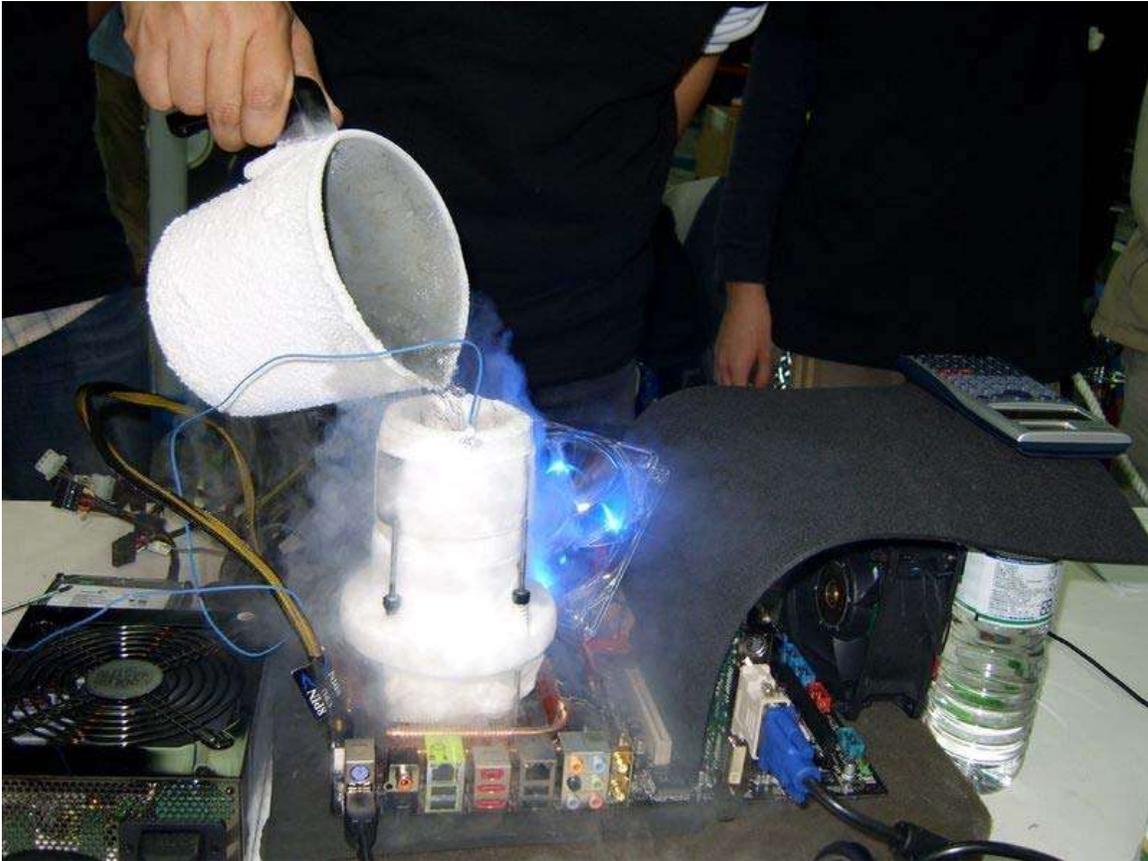
between clock frequencies and thermal design power (TDP) are linear. However, there is a limit to the maximum frequency which is called a "wall". To overcome this issue, overclockers raise the chip voltage to increase the overclocking potential. The relationship between chip voltage and TDP is exponential due to the fact that as the chip warms, the resistance increases. This increased heat requires effective cooling to avoid damaging the hardware. In addition, some digital circuits slow down at high temperatures due to changes in MOSFET device characteristics. Because most stock cooling systems are designed for the amount of power produced during non-overclocked use, overclockers typically turn to more effective cooling solutions, such as powerful fans, larger heat sinks, heat pipes and water cooling. Size, shape, and material all influence the ability of a heatsink to dissipate heat. Efficient heatsinks are often made entirely of copper, which has high thermal conductivity, but is expensive. Aluminium is more widely used; it has poorer thermal conductivity, but is significantly cheaper than copper. Heat pipes are commonly used to improve conductivity. Many heatsinks combine two or more materials to achieve a balance between performance and cost.



Interior of a water-cooled computer, showing CPU water block, tubing, and pump.

Water cooling carries waste heat to a radiator. Thermoelectric cooling devices, also known as Peltier devices, are recently popular with the onset of high thermal design power (TDP) processors made by Intel and AMD. Thermoelectric cooling devices create temperature differences between two plates by running an electric current through the

plates. This method of cooling is highly effective, but itself generates significant heat. For this reason, it is often necessary to supplement thermoelectric cooling devices with a convection-based heatsink or a water-cooling system.



Liquid nitrogen may be used for cooling an overclocked system, when an extreme measure of cooling is needed.

Other cooling methods are forced convection and phase change cooling which is used in refrigerators and can be adapted for computer use. Liquid nitrogen, liquid helium, and dry ice are used as coolants in extreme cases, such as record-setting attempts or one-off experiments rather than cooling an everyday system. In June 2006, IBM and Georgia Institute of Technology jointly announced a new record in silicon-based chip clock rate above 500 GHz, which was done by cooling the chip to 4.5 K ($-268.7\text{ }^{\circ}\text{C}$; $-451.6\text{ }^{\circ}\text{F}$) using liquid helium. These extreme methods are generally impractical in the long term, as they require refilling reservoirs of vaporizing coolant, and condensation can be formed on chilled components. Moreover, silicon-based junction gate field-effect transistors (JFET) will degrade below temperatures of roughly 100 K ($-173\text{ }^{\circ}\text{C}$; $-280\text{ }^{\circ}\text{F}$) and eventually cease to function or "freeze out" at 40 K ($-233\text{ }^{\circ}\text{C}$; $-388\text{ }^{\circ}\text{F}$) since the silicon ceases to be semiconducting so using extremely cold coolants may cause devices to fail.

Submersion cooling, used by the Cray-2 supercomputer, involves sinking a part of computer system directly into a chilled liquid that is thermally conductive but has low

electrical conductivity. The advantage of this technique is that no condensation can form on components. A good submersion liquid is Fluorinert made by 3M, which is expensive and can only be purchased with a permit. Another option is mineral oil, but impurities such as those in water might cause it to conduct electricity.

Stability and functional correctness

As an overclocked component operates outside of the manufacturer's recommended operating conditions, it may function incorrectly, leading to system instability. Another risk is silent data corruption by undetected errors. Such failures might never be correctly diagnosed and may instead be incorrectly attributed to software bugs in applications, device drivers, or the operating system. Overclocked use may permanently damage components enough to cause them to misbehave (even under normal operating conditions) without becoming totally unusable.

In general, overclockers claim that testing can ensure that an overclocked system is stable and functioning correctly. Although software tools are available for testing hardware stability, it is generally impossible for any private individual to thoroughly test the functionality of a processor. Achieving good fault coverage requires immense engineering effort; even with all of the resources dedicated to validation by manufacturers, faulty components and even design faults are not always detected.

A particular "stress test" can verify only the functionality of the specific instruction sequence used in combination with the data and may not detect faults in those operations. For example, an arithmetic operation may produce the correct result but incorrect flags; if the flags are not checked, the error will go undetected.

To further complicate matters, in process technologies such as silicon on insulator (SOI), devices display hysteresis—a circuit's performance is affected by the events of the past, so without carefully targeted tests it is possible for a particular sequence of state changes to work at overclocked rates in one situation but not another even if the voltage and temperature are the same. Often, an overclocked system which passes stress tests experiences instabilities in other programs.

In overclocking circles, "stress tests" or "torture tests" are used to check for correct operation of a component. These workloads are selected as they put a very high load on the component of interest (e.g. a graphically-intensive application for testing video cards, or different math-intensive applications for testing general CPUs). Popular stress tests include Prime95, Everest, Superpi, OCCT, IntelBurnTest/Linpack/LinX, SiSoftware Sandra, BOINC, Intel Thermal Analysis Tool and Memtest86. The hope is that any functional-correctness issues with the overclocked component will show up during these tests, and if no errors are detected during the test, the component is then deemed "stable". Since fault coverage is important in stability testing, the tests are often run for long periods of time, hours or even days. An overclocked computer is sometimes described using the number of hours and the stability program used, such as "prime 12 hours stable".

Factors allowing overclocking

Overclockability arises in part due to the economics of the manufacturing processes of CPUs and other components. In most cases components with different rated clock rates are manufactured by the same process, and tested after manufacture to determine their actual ratings. The clock rate that the component is rated for is at or below the clock rate at which the CPU has passed the manufacturer's functionality tests when operating in worst-case conditions (for example, the highest allowed temperature and lowest allowed supply voltage). Manufacturers must also leave additional margin for reasons discussed below. Sometimes manufacturers produce more high-performing parts than they can sell, so some are marked as medium-performance chips to be sold for medium prices. Pentium architect Bob Colwell calls overclocking an "uncontrolled experiment in better-than-worst-case system operation".

Measuring effects of overclocking

Benchmarks are used to evaluate performance. The benchmarks can themselves become a kind of 'sport', in which users compete for the highest scores. As discussed above, stability and functional correctness may be compromised when overclocking, and meaningful benchmark results depend on correct execution of the benchmark. Because of this, benchmark scores may be qualified with stability and correctness notes (e.g. an overclocker may report a score, noting that the benchmark only runs to completion 1 in 5 times, or that signs of incorrect execution such as display corruption are visible while running the benchmark). A widely used test of stability is Prime95 as this has in-built error checking and the computer fails if unstable.

Given only benchmark scores it may be difficult to judge the difference overclocking makes to the overall performance of a computer. For example, some benchmarks test only one aspect of the system, such as memory bandwidth, without taking into consideration how higher clock rates in this aspect will improve the system performance as a whole. Apart from demanding applications such as video encoding, high-demand databases and scientific computing, memory bandwidth is typically not a bottleneck, so a great increase in memory bandwidth may be unnoticeable to a user depending on the applications used. Other benchmarks, such as 3DMark attempt to replicate game conditions.

Variance

The extent to which a particular part will overclock is highly variable. Processors from different vendors, production batches, steppings, and individual units will all overclock with varying degrees of success.

Manufacturer and vendor overclocking

Commercial system builders or component resellers sometimes overclock to sell items at higher profit margins. The retailer makes more money by buying lower-value

components, overclocking them, and selling them at prices appropriate to a non-overclocked system at the new clock rate. In some cases an overclocked component is functionally identical to a non-overclocked one of the new clock rate, however, if an overclocked system is marketed as a non-overclocked system (it is generally assumed that unless a system is specifically marked as overclocked, it is not overclocked), it is considered fraudulent.

Overclocking is sometimes offered as a legitimate service or feature for consumers, in which a manufacturer or retailer tests the overclocking capability of processors, memory, video cards, and other hardware products. Several video card manufactures now offer factory overclocked versions of their graphics accelerators, complete with a warranty, which offers an attractive solution for enthusiasts seeking an improved performance without sacrificing common warranty protections. Such factory-overclocked products may cost a little more than standard components, but may be more cost-effective than a product with a higher specification.

Naturally, manufacturers would prefer enthusiasts to pay additional money for profitable high-end products, in addition to concerns of less reliable components and shortened product life spans affecting brand image. It is speculated that such concerns are often motivating factors for manufacturers to implement overclocking prevention mechanisms such as CPU locking. These measures are sometimes marketed as a consumer protection benefit, which typically generates a negative reception from overclocking enthusiasts.

Advantages

- The user can, in many cases, purchase a lower performance, cheaper component and overclock it to the clock rate of a more expensive component.
- Higher performance in games, encoding, video editing applications, and system tasks at no additional expense, but at an increased cost for electrical power consumption. Particularly for enthusiasts who regularly upgrade their hardware, overclocking can increase the time before an upgrade is needed.
- Some systems have "bottlenecks," where small overclocking of a component can help realize the full potential of another component to a greater percentage than the limiting hardware is overclocked. For instance, many motherboards with AMD Athlon 64 processors limit the clock rate of four units of RAM to 333 MHz. However, the memory performance is computed by dividing the processor clock rate (which is a base number times a CPU multiplier, for instance 1.8 GHz is most likely 9×200 MHz) by a fixed integer such that, at a stock clock rate, the RAM would run at a clock rate near 333 MHz. Manipulating elements of how the processor clock rate is set (usually lowering the multiplier), one can often overclock the processor a small amount, around 100–200 MHz (less than 10%), and gain a RAM clock rate of 400 MHz (20% increase), releasing the full potential of the RAM.
- Overclocking can be an engaging hobby in itself and supports many dedicated online communities. The PCMark website is one such site that hosts a leaderboard for the most powerful computers to be bench-marked using the program.

- A new overclocker with proper research and precaution or a guiding hand can gain useful knowledge and hands-on experience about their system and PC systems in general.

Disadvantages

Many of the disadvantages of overclocking can be mitigated or reduced in severity by skilled overclockers. However, novice overclockers may make mistakes while overclocking which can introduce avoidable drawbacks and which are more likely to damage the overclocked components (as well as other components they might affect).

General

- The lifespan of a processor may be reduced by higher operating frequencies, increased voltages and heat, although processors rapidly become obsolete in performance due to technological progress.
- Increased clock rates and/or voltages result in higher power consumption.
- While overclocked systems may be tested for stability before use using programs that "burn" the computer, these programs create an artificial strain that pushes individual or many components to their maximum (or beyond it). Some common stability programs are Prime95, Super PI (32M), Intel TAT, LinX, PCMark, FurMark and OCCT. Stability problems may surface after prolonged usage due to new workloads or untested portions of the processor core. Aging effects previously discussed may also result in stability problems after a long period of time. Even when a computer appears to be working normally, problems may arise in the future. For example, Windows may appear to work with no problems, but when it is re-installed or upgraded, error messages may be received such as a "file copy error" during Windows Setup. Microsoft says this of errors in upgrading to Windows XP: "Your computer [may be] over-clocked. Because over-clocking is very memory-intensive, decoding errors may occur when files are extracted from the Windows XP CD-ROM".
- High-performance fans used for extra cooling can be noisy. Older popular models of fans used by overclockers can produce 50 decibels or more. However, nowadays, manufacturers are overcoming this problem by designing fans with aerodynamically optimized blades for smoother airflow and minimal noise (around 20 decibels at approximately 1 metre). The noise is not always acceptable, and overclocked machines are often much noisier than stock machines. Noise can be reduced by utilizing strategically-placed larger fans, which are inherently less noisy than smaller fans; by using alternative cooling methods (such as liquid and phase-change cooling); by lining the chassis with foam insulation; and by installing a fan-controlling bus to adjust fan speed (and, as a result, noise) to suit the task at hand. Now that overclocking is of interest to a larger target audience, this is less of a concern as manufacturers have begun researching and producing high-performance fans that are no longer as loud as their predecessors. Similarly, mid- to high-end PC cases now implement larger

fans (to provide better airflow with less noise) as well as being designed with cooling and airflow in mind.

- Even with adequate CPU cooling, the excess heat produced by an overclocked processing unit increases the ambient air temperature of the system case; consequently, other components may be affected. Also, more heat will be expelled from the PC's vents, raising the temperature of the room the PC is in - sometimes to uncomfortable levels.
- Overclocking has the potential to cause component failure ("heat death"). Most warranties do not cover damage caused by overclocking. Some motherboards offer safety measures that will stop this from happening (e.g. limitations on FSB increase) so that only voltage control alterations can cause such harm.
- Some motherboards are designed to use the airflow from a standard CPU fan in order to cool other heatsinks, such as the northbridge. If the cpu heatsink is changed on such boards, other heatsinks may receive insufficient cooling.
- Overclocking a PC component may void its warranty (depending on the conditions of sale).
- Changing the Heatsink on a Graphics Card often voids its warranty

Incorrectly performed overclocking

- Increasing the operation frequency of a component will usually increase its thermal output in a linear fashion, while an increase in voltage usually causes heat to increase quadratically. Excessive voltages or improper cooling may cause chip temperatures to rise almost instantaneously, causing the chip to be damaged or destroyed.
- More common than hardware failure is functional incorrectness. Although the hardware is not permanently damaged, this is inconvenient and can lead to instability and data loss. In rare, extreme cases entire filesystem failure may occur, causing the loss of all data.
- With poor placement of fans, turbulence and vortices may be created in the computer case, resulting in reduced cooling effectiveness and increased noise. In addition, improper fan mounting may cause rattling or vibration.
- Improper installation of exotic cooling solutions like liquid cooling may result in failure of the cooling system, which may result in water damage.
- With sub-zero cooling methods such as phase-change cooling or liquid nitrogen, extra precautions such as foam or spray insulation must be made to prevent water from condensing upon the PCB and other areas. This can cause the board to become "frosted" or covered in frost. While the water is frozen it is usually safe, however once it melts it can cause shorts and other malignant issues.
- Sometimes products claim to be intended specifically for overclocking and may be just decoration. Novice buyers should be aware of the marketing hype surrounding some products. Examples include heat spreaders and heat sinks designed for chips (or components) which do not generate enough heat to benefit from these devices (capacitors, for example).

Limitations

The utility of overclocking is limited for a few reasons:

- Personal computers are mostly used for tasks which are not computationally demanding, or which are performance-limited by bottlenecks outside of the local machine. For example, web browsing does not require a high performance computer, and the limiting factor will almost certainly be the bandwidth of the Internet connection of either the user or the server. Overclocking a processor will also do little to help increase application loading times as the limiting factor is reading data off the hard drive. Other general office tasks such as word processing and sending email are more dependent on the efficiency of the user than on the performance of the hardware. In these situations any performance increases through overclocking are unlikely to be noticeable.
- It is generally accepted that, even for computationally-heavy tasks, clock rate increases of less than ten percent are difficult to discern. For example, when playing video games, it is difficult to discern an increase from 60 to 66 frames per second (FPS) without the aid of an on-screen frame counter. Overclocking of a processor will rarely improve gaming performance noticeably, as the frame rates achieved in most modern games are usually bound by the GPU at resolutions beyond 1024×768. One exception to this rule is when the overclocked component is the bottleneck of the system, in which case the most gains can be seen.

Graphics cards



The BFG GeForce 6800GSOC ships with higher memory and clock rates than the standard 6800GS.

Graphics cards can also be overclocked, with utilities such as EVGA's Precision, RivaTuner, ATI Overdrive (on ATI cards only), MSI Afterburner, Zotac Firestorm on Zotac cards, or the PEG Link Mode on Asus motherboards. Overclocking a GPU will often yield a marked increase in performance in synthetic benchmarks, and usually will improve game performance too. Sometimes, it is possible to see that a graphics card is pushed beyond its limits before any permanent damage is done by observing on-screen distortions known as artifacts. Two such discriminated "warning bells" are widely understood: green-flashing, random triangles appearing on the screen usually correspond to overheating problems on the GPU itself, while white, flashing dots appearing randomly (usually in groups) on the screen often mean that the card's RAM is overheating. It is common to run into one of those problems when overclocking graphics cards. Showing both symptoms at the same time usually means that the card is severely pushed beyond its heat/clock rate/voltage limits. If seen at normal clock rate, voltage and temperature, they may indicate faults with the card itself. However, if the video card is simply clocked too high and doesn't overheat then the artifacts are a bit different. There are many different ways for this to show up and any irregularities should be considered but usually if the core is pushed too hard black circles or blobs appear on the screen and overclocking the video memory beyond its limits usually results in the application or the

entire operating system crashing. Luckily though, after the computer is restarted the settings is reset to stock (stored in the video card firmware) and the maximum clock rate of that specific card has been found.

Some overclockers use a hardware voltage modification where a potentiometer is applied to the video card to manually adjust the voltage. This results in much greater flexibility, as overclocking software for graphics cards is rarely able to freely adjust the voltage. Voltage mods are very risky and may result in a dead video card, especially if the voltage modification ("voltmod") is applied by an inexperienced individual. A pencil volt mod refers to changing a resistor's value on the graphics card by drawing across it with a graphite pencil. This results in a change of GPU voltage. It is also worth mentioning that adding physical elements to the video card immediately voids the warranty.

Alternatives

Flashing and Unlocking are two popular ways to gain performance out of a video card, without technically overclocking.

Flashing refers to using the firmware of another card, based on the same core and design specs, to "override" the original firmware, thus effectively making it a higher model card; however, 'flashing' can be difficult, and sometimes a bad flash can be irreversible. Sometimes stand-alone software to modify the firmware files can be found, i.e. NiBiTor, (GeForce 6/7 series are well regarded in this aspect). It is not necessary to acquire a firmware file from a better model video card (although it should be said that the card in which firmware is to be used should be compatible, i.e. the same model base, design and/or manufacture process, revisions etc.). For example, video cards with 3D accelerators (the vast majority of today's market) have two voltage and clock rate settings - one for 2D and one for 3D - but were designed to operate with *three* voltage stages, the third being somewhere in the middle of the aforementioned two, serving as a fallback when the card overheats or as a middle-stage when going from 2D to 3D operation mode. Therefore, it could be wise to set this middle-stage prior to "serious" overclocking, specifically because of this fallback ability - the card can drop down to this clock rate, reducing by a few (or sometimes a few dozen, depending on the setting) percent of its efficiency and cool down, without dropping out of 3D mode (and afterwards return to the desired high performance clock and voltage settings).

Some cards also have certain abilities not directly connected with overclocking. For example, Nvidia's GeForce 6600GT (AGP flavor) features a temperature monitor (used internally by the card), which is invisible to the user in the 'vanilla' version of the card's firmware. Modifying the firmware can allow a 'Temperature' tab to become visible in the card driver's advanced menu.

Unlocking refers to enabling extra pipelines and/or pixel shaders. The 6800LE, the 6800GS and 6800 (AGP models only) and Radeon X800 Pro VIVO were some of the first cards to benefit from unlocking. While these models have either 8 or 12 pipes enabled, they share the same 16x6 GPU core as a 6800GT or Ultra, but may not have

passed inspection when all their pipelines and shaders were unlocked. In more recent generations, both ATI and Nvidia have laser cut pipelines to prevent this practice..

It is important to remember that while pipeline unlocking sounds very promising, there is absolutely no way of determining if these 'unlocked' pipelines will operate without errors, or at all (this information is solely at the manufacturer's discretion). In a worst-case scenario, the card may not start up ever again, resulting in a 'dead' piece of equipment. It is possible to revert to the card's previous settings, but it involves manual firmware flashing using special tools and an identical but original firmware chip.

WWT

Chapter 7

Power Management and SpeedFan

Power management

Power management is a feature of some electrical appliances, especially copiers, computers and computer peripherals such as monitors and printers, that turns off the power or switches the system to a low-power state when inactive. In computing this is known as PC power management and is built around a standard called ACPI. This supersedes APM. All recent (consumer) computers have ACPI support.

Motivations

PC power management for computer systems is desired for many reasons, particularly:

- Reduce overall energy consumption
- Prolong battery life for portable and embedded systems
- Reduce cooling requirements
- Reduce noise.
- Reduce operating costs for energy and cooling.

Lower power consumption also means lower heat dissipation, which increases system stability, and less energy use, which saves money and reduces the impact on the environment.

Processor level techniques

The power management for microprocessors can be done over the whole processor, or in specific areas.

With dynamic voltage scaling and dynamic frequency scaling, the CPU core voltage, clock rate, or both, can be altered to decrease power consumption at the price of potentially lower performance. This is sometimes done in real time to optimize the power-performance tradeoff.

Examples:

- AMD Cool'n'Quiet
- AMD PowerNow!
- IBM EnergyScale
- Intel SpeedStep
- Transmeta LongRun and LongRun2
- VIA LongHaul (PowerSaver)

Additionally, processors can selectively power off internal circuitry (power gating). For example:

- Newer Intel Core processors support ultra-fine power control over the functional units within the processors.
- AMD CoolCore technology get more efficient performance by dynamically activating or turning off parts of the processor.

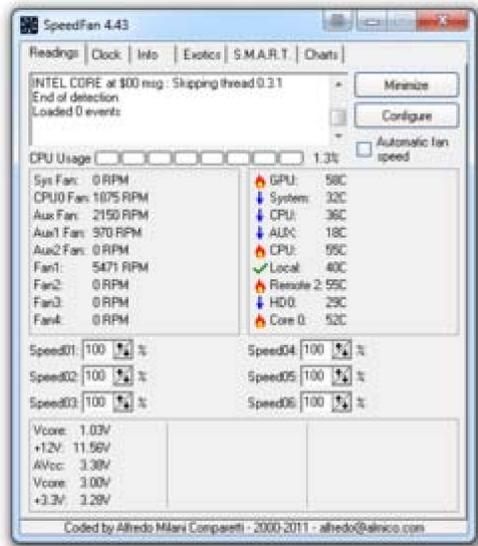
Intel VRT technology split the chip into a 3.3V I/O section and a 2.9V core section. The lower core voltage reduces power consumption.

Operating system level: Hibernation

When a computer system hibernates it saves the contents of the RAM to disk and powers down the machine. On startup it reloads the data. This allows the system to be completely powered off while in hibernate mode. This requires a file the size of the installed RAM to be placed on the hard disk, potentially using up space even when not in hibernate mode. Hibernate mode is enabled by default in some versions of Windows and can be disabled in order to recover this disk space.

SpeedFan

SpeedFan



SpeedFan 4.43 in Windows 7

Original author(s)	Alfredo Milani Comparetti
Developer(s)	Alfredo Milani Comparetti
Initial release	?
Stable release	4.43 (March 17, 2011; 22 days ago) [+/-]
Preview release	4.44 beta 4 (April 7, 2011; 1 day ago) [+/-]
Written in	Delphi, C++, C
Operating system	Windows 95 and later
Available in	Multilanguage
Type	System hardware monitor
License	Freeware

SpeedFan is software that can read temperatures, voltages and fan speeds of computer components. It can change computer fan speeds depending on the temperature of various components. The program can display system variables as a chart and as an indicator in the system tray. Fully configurable user events can be defined to execute specific actions based on system status. As of version 4.37, temperatures in the event section must be specified in °C no matter if the display setting is for °F.

Hard disk support

SpeedFan also monitors S.M.A.R.T. readings for EIDE, SATA and SCSI hard disks. Starting with version 4.35, SpeedFan fully supports Areca RAID controllers. Version 4.38 added full support for AMCC/3ware SATA and RAID controllers.

Hard disk in-depth online analysis

SpeedFan offers a feature named "in-depth online analysis" that compares your hard disk's S.M.A.R.T. data to a comprehensive database with statistical models of every known hard disk allowing early detection of potentially degraded hard disks. Full text messages inform the user of specific situations and problems, as if a human expert had looked at the data.

Risks

SpeedFan crashes some systems when it is launched, possibly causing registry corruption or loss of recent changes to the registry. This can cause recently-installed programs to disappear from Add/Remove Programs, programs to no longer function correctly, or loss of operating-system stability.

Chapter 8

SpeedStep and Tweaking

SpeedStep

SpeedStep is a trademark for a series of dynamic frequency scaling technologies (codenamed **Geyserville** and including SpeedStep, SpeedStep II, and SpeedStep III) built into some Intel microprocessors that allow the clock speed of the processor to be dynamically changed (to different *P-states*) by software. This allows the processor to meet the instantaneous performance needs of the operation being performed, while minimizing power draw and heat dissipation. *Enhanced Intel SpeedStep* is sometimes abbreviated as *EIST*.

Explanation

Running a processor at high clock speeds allows for better performance. However, when the same processor is run at a lower frequency (speed), it generates less heat and consumes less power. In many cases, the core voltage can also be reduced, further reducing power consumption and heat generation. This can conserve battery power in notebooks, extend processor life, and reduce noise generated by variable-speed fans. By using SpeedStep, users can select the balance of power conservation and performance that best suits them, or even change the clock speed dynamically as the processor burden changes.

The power consumed by a CPU with a capacitance C , running at frequency f and voltage V is approximately

$$P = CV^2f.$$

For a given processor, C is a fixed value. However, V and f can vary considerably. For example, for a 1.6 GHz Pentium M, the clock frequency can be stepped in 200 MHz increments over the range from 0.6 to 1.6 GHz. At the same time, the voltage requirement decreases from 1.484 V to 0.956 V. The result is that the power consumption theoretically goes down by a factor of 6.4. In practice, the effect may be smaller because

some CPU instructions use less energy per tick of the CPU clock than others. For example, when an operating system is not busy, it tends to issue halt instructions, which suspend operation of parts of the CPU for a time period, so it uses less energy per tick of the CPU clock than when executing productive instructions in its normal state. For a given rate of work, a CPU running at a higher clock rate will execute a greater proportion of HLT instructions. The simple equation which relates power, voltage and frequency above also does not take into account the static power consumption of the CPU. This tends not to change with frequency, but does change with temperature and voltage. Hot electrons, and electrons exposed to a stronger electric field are more likely to migrate across a gate as "gate leakage" current, leading to an increase in static power consumption.

Older processors, using older versions of the SpeedStep technology, have fewer increments, such as the Pentium 4-M. For example, a 1.7 GHz Pentium 4M can run at 1.6 GHz, at 1.2 GHz, and at 786 MHz.

SpeedStep technology is partly responsible for the reduced power consumption of Intel's Pentium M processor, part of the Centrino brand.

Problems when using SpeedStep

Microsoft reports there may be problems previewing video files when SpeedStep (or the AMD equivalent PowerNow!) is enabled. It also may decrease reliability when overclocking.

Operating system support

Older versions of Microsoft Windows, Windows 2000 and earlier, need a special driver and dashboard application to access the SpeedStep feature. Intel's website specifically states that such drivers *must* come from the computer manufacturer; there are no generic drivers supplied by Intel which will enable SpeedStep for older Windows versions if one cannot obtain a manufacturer's driver.

Under Microsoft Windows XP, SpeedStep support is built into the power management console under the control panel. In Windows XP a user can regulate processor speed indirectly by changing power schemes. The "Home/Office Desk" setting disables SpeedStep, the "Portable/Laptop" power scheme enables SpeedStep, and the "Max Battery" uses SpeedStep to slow the processor to minimal power levels as the battery weakens. The SpeedStep settings for power schemes, either built-in or custom, cannot be modified from the control panel's GUI, but can be modified using the POWERCFG.EXE command-line utility.

In contrast, AMD continues to supply and support drivers for its competing PowerNow! technology that will work on Windows 2000, ME, 98, and NT.

Linux has full SpeedStep support integrated into the kernel version 2.6.

Mac OS also has SpeedStep built into the kernel, since the release of the Intel version of Mac OS X 10.4 and is already enabled. It cannot be controlled in the System Preference "Energy Saver." To disable this feature, and set a specific clock speed (full speed or reduced) requires a third party application, such as coolbook

Solaris has supported SpeedStep since OpenSolaris SXDE 9/07.

The BSD kernels have full SpeedStep support integration.

Versions

V1.1 is used by second generation Pentium III processors. It enables the CPU to switch between two modes: high and low frequency. This is done by modifying the CPU's multiplier. A 1 GHz Pentium III consuming about 20 watts could be reduced to 600 MHz which reduces the power consumption to about 6 watts.

V2.1 (Enhanced SpeedStep) is used in Pentium III-Mobile processors and is similar to the previous version, but in the low frequency mode the CPU also uses a different voltage than the high frequency mode.

V2.2 is adapted for Pentium 4-Mobile processors. With this, a 1.8 GHz Pentium 4-M consuming about 30 watts can lower its frequency to 1.2 GHz, thus reducing power consumption to about 20 watts.

V3.1 (EIST) is used with the first and second generation of Pentium M processors (Banias and Dothan cores, used in Centrino platforms). With this technology, the CPU varies its frequency (and voltage) between about 40% and 100% of its base frequency in increments of 100 MHz (for Banias core) or 133 MHz (for Dothan core). With this technology, Intel also introduces realtime Level 2 cache capacity variation, further improving power savings.

Tweaking

Tweaking refers to fine-tuning or adjusting a complex system, usually an electronic device. Tweaks are any small modifications intended to improve a system.

In electronics, it is a synonym for "trimming." Analog circuit boards often have small potentiometers or other components on them that are used to calibrate or adjust the board as a service procedure: the small insulated screwdriver used to turn them is often called a "tweaker."

This use was echoed in the name of the product *Tweek*, a popular but controversial audio product during the 1980s, which was claimed to improve the electrical characteristics of audio switch contacts.

Hardware

Hardware tweaking is a process of modifying certain parts of a hardware such as changing the cables, cleaning the heads of a VHS player with a branded cleaning fluid or oiling the moving parts of an engine with the best possible oil.

Computer hardware

Computer hardware tweaking is an extension of hardware tweaking, specifically geared towards the components of a PC. They include: changing voltage and clock rates of processing units, modifying RAM unit timing, improving cooling systems to reduce chance of overheating, etc.

Tweaks specifically designed to allow a processor to operate at a higher clock speed than normal are known as overclocking.

Modifications of computer systems not aimed at increasing performance, such as quieter fans, external controls, and decorations such as lights or windows, are known as modding.

Software

Software tweaking is the process of improving the performance of an application or the quality of its output. There can be two ways of accomplishing this: manually (that is, if one is familiar with programming; though it may be irrelevant if the source of the application is closed, and there are no built-in means to adjust its performance) or using another piece of software specialized for that purpose (such as Red Button and SysTool, semi-automatic tweaking utilities used to improve the performance of Microsoft Windows XP. Tweaking of this kind generally increases usability, in terms of personal configuration preferences, rather than objective performance of the system overall). Linux and other open source products are designed to facilitate the tweaking process as much as possible, as opposed to Microsoft Windows which limits tweaking but allows it with provision, and Mac OS which largely prohibits or strongly discourages tweaking.

Some very precise applications need constant and thorough tweaking to stay up to date and deliver best possible results. One of the most obvious examples of such a fine tuning is LAME MP3 encoder, whose 3.9x branch is not only considered as the state-of-the-art MP3 encoder, but also continues to shape the boundaries of the MP3 codec and stay competitive with its successors.

Chapter 9

Underclocking

"**Underclocking**" also known as "**Downclocking**" is the practice of modifying a synchronous circuit's timing settings to run at a lower clock rate than it was specified to operate at. It may be said to be the computer equivalent to drive a car at a speed below the speed limit. Usually, underclocking is used to reduce a computer's power consumption and heat emission, sometimes also to increase the system's stability and compatibility. Underclocking may be implemented by the factory, but many computers and components are end user underclockable.

Types of underclocking

CPU underclocking

For microprocessors, the purpose is generally to decrease the need for heat dissipation devices or decrease the electrical power consumption. This can provide increased system stability in high-heat environments, or can allow a system to run with a lower airflow (and therefore quieter) cooling fan or without one at all. For example, a Pentium 4 processor clocked at 2.4 GHz can be "underclocked" to 1.8 GHz and can then be safely run with reduced fan speeds. This invariably comes at the expense of some system performance. However, the performance usually is reduced less than the reduction of clock speed because the performance often is limited by other bottlenecks: The hard disk, the disk controller, the Internet, the network, etc.

Graphics cards

"Underclocking" can also be performed on graphics card processors (graphics processing units; GPUs), usually with the aim of reducing heat output. For instance, it is possible to set a GPU to run at lower clock rates when performing everyday tasks (e.g. internet browsing and word processing), thus allowing the card to operate at lower temperature and thus lower, quieter fan speeds. The GPU can then be overclocked for more graphically intense applications, such as games. "Underclocking" a GPU will reduce performance, but this decrease will probably not be noticeable except in graphically intensive applications.

Memory "underclocking"

Newer and faster RAM may be "underclocked" to match older systems as an inexpensive way to replace rare or discontinued memory. This might also be necessary if stability problems are encountered at higher settings, especially in a PC with several memory modules of different clock speed. If you underclock a PC processor, and do not change the clock factor or multiplier (the ratio between the processor and the memory clock speed), the memory will also be underclocked.

When used

Dynamic frequency scaling (automatic "underclocking") is very common on laptop computers and is beginning to emerge on desktop computers as well. In laptops, the processor is usually "underclocked" automatically whenever the computer is operating on batteries. Most newer notebook and some desktop processors (like AMD's Cool'n'Quiet and PowerNow!) will also underclock themselves automatically when under a light processing load. Intel has also used this method on their Core 2 Duo and later processors, through a feature called SpeedStep.

Some processors "underclock" automatically as a defensive measure, to prevent overheating which could cause permanent damage. When such a processor reaches a temperature level deemed too high for safe operation, the *thermal control circuit* activates, automatically decreasing the clock and CPU core voltage until the temperature has returned to a safe level. In a properly cooled environment, this mechanism should trigger rarely (if ever).

There are several different "underclocking" competitions similar in format to overclocking competitions, except the goal is to have the lowest clocked computer, as opposed to the highest.

Advantages

- Reduced electrical power consumption, especially when combined with undervolting (i.e., reduce the component's voltage below the nominal). For instance, by underclock an Athlon XP 1700+ processor from 1466 to 1000 MHz and reduce the core voltage from 1.75 to 1.15V, a computer user reduced the power consumption from 64.0 to 21.6W, i.e. 66 % power reduction and only 26 % less performance. The same is true for newer processors: When a single-core Intel CPU was 20 % underclocked, the PC's performance was down only 13 % with a 49 % power reduction.

In general, the power consumed by a CPU with a capacitance C , running at frequency f and voltage V is approximately

$$P = CV^2f.$$

- Reduced heat generation, which is exactly proportional to the power consumption.
- Less noise because the cooling fans may be slowed down, or even eliminated. A cooling fan's efficiency is proportional to its rotation speed, but its noise grows much more.
- Longer hardware lifespan.
- Increased stability.
- Reduced noise from cooling parts due to reduced heat dissipation requirements.
- Increased battery life.
- Better compatibility with old applications.

In practice

Linux

Linux kernel supports CPU frequency modulation. In supported processors, using *cpufreq* to gain access to this feature gives the system administrator a variable level of control over the CPU's clock rate. The kernel includes five governors by default: Conservative, Ondemand, Performance, Powersave, and Userspace. The Conservative and Ondemand governors adjust the clock rate depending on the CPU load, but each with different algorithms. The Ondemand governor jumps to maximum frequency on CPU load and decreases the frequency step by step on CPU idle, whereas the Conservative governor increases the frequency step by step on CPU load and jumps to lowest frequency on CPU idle. The Performance, Powersave and Userspace governors set the clock rate statically: Performance to the highest available, Powersave to the lowest available, and Userspace to a frequency determined and controlled by the user.

Windows

"Underclocking" can be done manually in the BIOS or with Windows applications, or dynamically using features such as Intel's SpeedStep or AMD's Cool'n'Quiet.

Asus Eee PC

Some versions of the Asus Eee PC uses a 900 MHz Intel Celeron M processor underclocked to 630 MHz.

Smartphones and PDAs

Most smartphones and PDAs such as the Motorola Droid, Palm Pre, and Apple's iPhone, iPhone 3G, and iPhone 3GS use the "underclocking" of a more powerful processor, rather than the full clocking of a less powerful processor, to maximize battery life. The designers for such mobile devices often discover that a slower processor gives worse battery life than a more powerful processor at a lower clock rate. They select a processor on the basis of the performance per watt of the processor.

Performance

The performance of an "underclocked" machine will often be better than might be expected. Under normal desktop use, the full power of the CPU is rarely needed. Even when the system is busy, a large amount of time is usually spent waiting for data from memory, disk, or other devices. Such devices communicate with the CPU through a bus which operates at a much lower bandwidth. Generally speaking, the lower the CPU multiplier (and thus clockrate of a CPU), the closer its performance will be to that of the bus, and the less time it will spend waiting.

WWT

Chapter 10

Random-Access Memory



Example of writable volatile random-access memory: Synchronous Dynamic RAM modules, primarily used as main memory in personal computers, workstations, and servers.

Random-access memory (RAM) is a form of computer data storage. Today, it takes the form of integrated circuits that allow stored data to be accessed in any order in a constant time, regardless of its physical location and whether it is related to the previous piece of data. RAM is often associated with volatile types of memory (such as DRAM memory

modules), where its stored information is lost if the power is removed. Many other types of non-volatile memory are RAM as well, including most types of ROM and a type of flash memory called *NOR-Flash*. The first RAM modules to come into the market were created in 1951 and were sold until the late 1960s and early 1970s. However, other memory devices (magnetic tapes, disks) can access the storage data in a predetermined order, because mechanical designs only allow this.

History

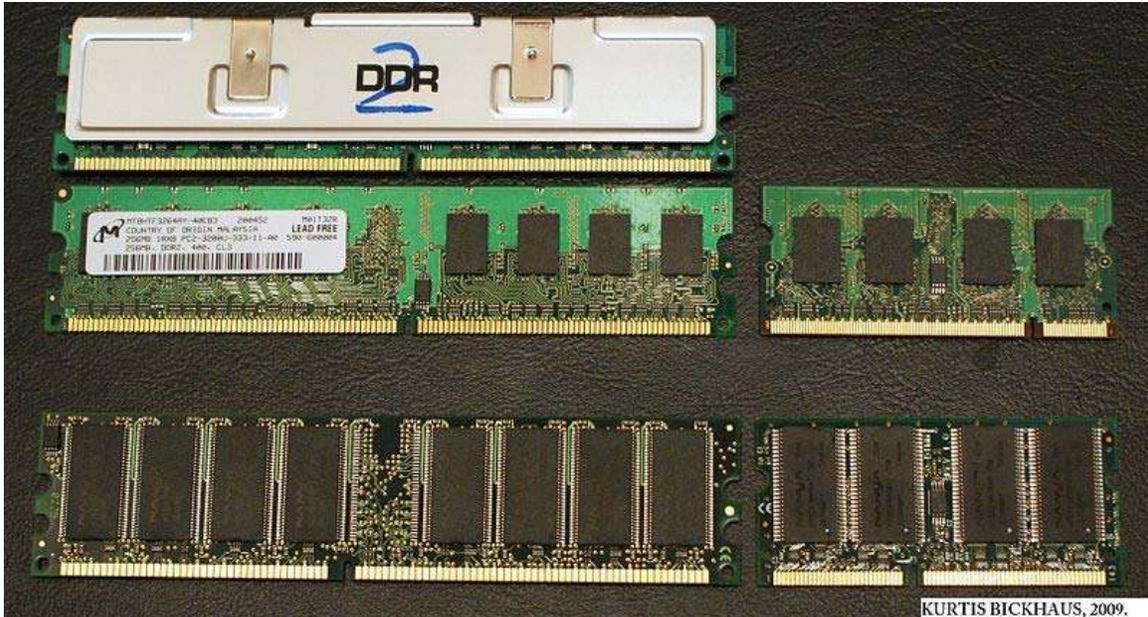


Old RAM modules

An early type of widespread *writable* random-access memory was the magnetic core memory, developed from 1949 to 1952, and subsequently used in most computers up until the development of the static and dynamic integrated RAM circuits in the late 1960s and early 1970s. Before this, computers used relays, delay line/delay memory, or various kinds of vacuum tube arrangements to implement "main" memory functions (i.e., hundreds or thousands of bits), some of which were *random access*, some not. Drum memory could be expanded at low cost but retrieval of non-sequential memory items required knowledge of the physical layout of the drum to optimize speed. Latches built out of vacuum tube triodes, and later, out of discrete transistors, were used for smaller and faster memories such as random-access register banks and registers. Prior to the development of integrated ROM circuits, *permanent* (or *read-only*) random-access

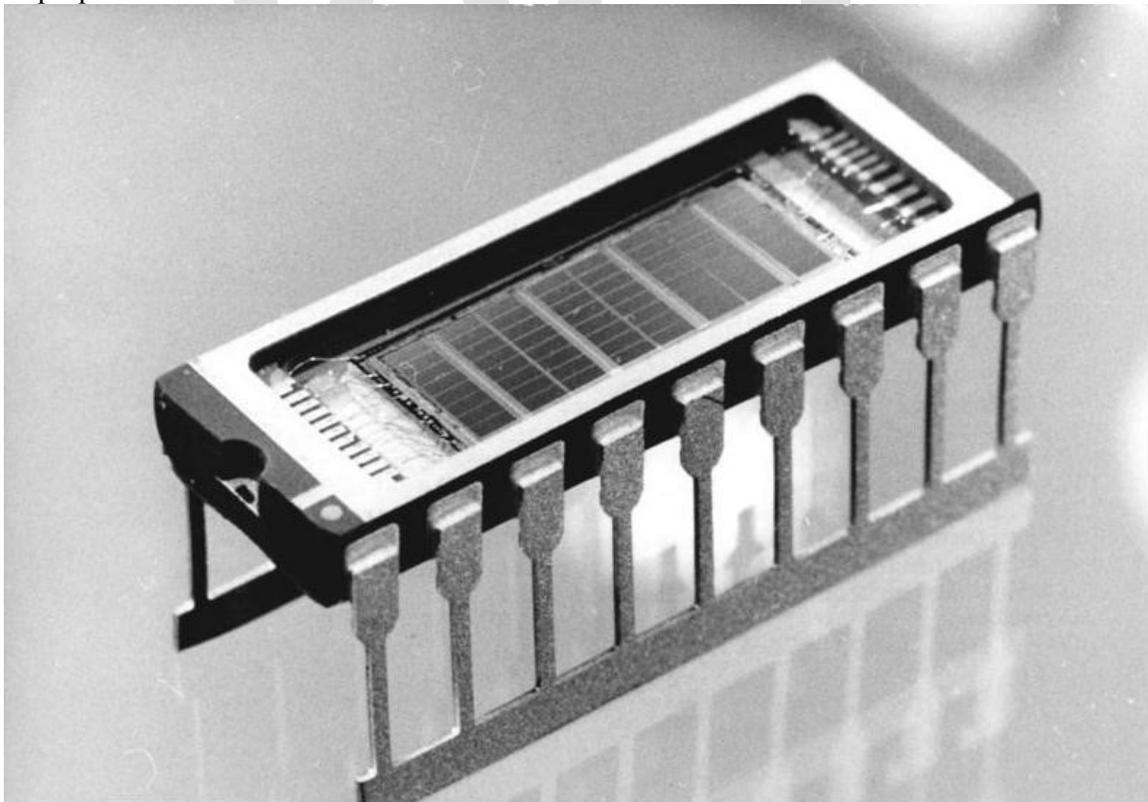
memory was often constructed using semiconductor diode matrices driven by address decoders, or specially wound core memory planes.

Types of RAM



KURTIS BICKHAUS, 2009.

Top L-R, DDR2 with heat-spreader, DDR2 without heat-spreader, Laptop DDR2, DDR, Laptop DDR



1 Megabit chip - one of the last models developed by VEB Carl Zeiss Jena in 1989

Modern types of *writable* RAM generally store a bit of data in either the state of a flip-flop, as in SRAM (static RAM), or as a charge in a capacitor (or transistor gate), as in DRAM (dynamic RAM), EPROM, EEPROM and Flash. Some types have circuitry to detect and/or correct random faults called *memory errors* in the stored data, using parity bits or error correction codes. RAM of the *read-only* type, ROM, instead uses a metal mask to permanently enable/disable selected transistors, instead of storing a charge in them. Of special consideration are SIMM and DIMM memory modules. Because of economic reasons, personal computers, workstations and game consoles do not control large bits of data stored in DRAM. The computer's cache memory uses static RAM disks of data buffers.

SRAM and DRAM are *volatile*. Other forms of computer storage, such as disks and magnetic tapes, have been used as persistent storage. Many newer products instead rely on flash memory to maintain data when not in use, such as PDAs or small music players. Certain personal computers, such as many rugged computers and netbooks, have also replaced magnetic disks with flash drives. With flash memory, only the NOR type is capable of true random access, allowing direct code execution, and is therefore often used instead of ROM; the lower cost NAND type is commonly used for bulk storage in memory cards and solid-state drives. A memory chip is an integrated circuit (IC) made of millions of transistors and capacitors. In the most common form of computer memory, dynamic random access memory (DRAM), a transistor and a capacitor are paired to create a memory cell, which represents a single bit of data. The capacitor holds the bit of information — a 0 or a 1. The transistor acts as a switch that lets the control circuitry on the memory chip read the capacitor or change its state.

Memory hierarchy

Many computer systems have a memory hierarchy consisting of CPU registers, on-die SRAM caches, external caches, DRAM, paging systems, and virtual memory or swap space on a hard drive. This entire pool of memory may be referred to as "RAM" by many developers, even though the various subsystems can have very different access times, violating the original concept behind the *random access* term in RAM. Even within a hierarchy level such as DRAM, the specific row, column, bank, rank, channel, or interleave organization of the components make the access time variable, although not to the extent that rotating storage media or a tape is variable. The overall goal of using a memory hierarchy is to obtain the higher possible average access performance while minimizing the total cost of the entire memory system (generally, the memory hierarchy follows the access time with the fast CPU registers at the top and the slow hard drive at the bottom).

In many modern personal computers, the RAM comes in an easily upgraded form of modules called memory modules or DRAM modules about the size of a few sticks of chewing gum. These can quickly be replaced should they become damaged or when changing needs demand more storage capacity. As suggested above, smaller amounts of RAM (mostly SRAM) are also integrated in the CPU and other ICs on the motherboard, as well as in hard-drives, CD-ROMs, and several other parts of the computer system.

Swapping

If a computer becomes low on RAM during intensive application cycles, many CPU architectures and operating systems are able to perform an operation known as "swapping". Swapping uses a *paging file*, an area on a hard drive temporarily used as additional working memory. Excessive use of this mechanism is called thrashing and is generally undesirable because it lowers overall system performance, mainly because hard drives are far slower than RAM. However, if a program attempts to allocate memory and fails, it may crash.

Other uses of the "RAM" term

Other physical devices with read–write capability can have "RAM" in their names: for example, DVD-RAM. "Random access" is also the name of an indexing method: hence, disk storage is often called "random access" because the reading head can move relatively quickly from one piece of data to another, and does not have to read all the data in between. However the final "M" is crucial: "RAM" (provided there is no additional term as in "DVD-RAM") always refers to a solid-state device.

Often, RAM is shorthand referring to the computer's main working memory.

RAM disks

Software can "partition" a portion of a computer's RAM, allowing it to act as a much faster hard drive that is called a RAM disk. A RAM disk loses the stored data when the computer is shut down, unless memory is arranged to have a standby battery source.

Shadow RAM

Sometimes, the contents of a relatively slow ROM chip are copied to read/write memory to allow for shorter access times. The ROM chip is then disabled while the initialized memory locations are switched in on the same block of addresses (often write-protected). This process, sometimes called *shadowing*, is fairly common in both computers and embedded systems.

As a common example, the BIOS in typical personal computers often has an option called "use shadow BIOS" or similar. When enabled, functions relying on data from the BIOS's ROM will instead use DRAM locations (most can also toggle shadowing of video card ROM or other ROM sections). Depending on the system, this may not result in increased performance, and may cause incompatibilities. For example, some hardware may be inaccessible to the operating system if shadow RAM is used. On some systems the benefit may be hypothetical because the BIOS is not used after booting in favor of direct hardware access. Free memory is reduced by the size of the shadowed ROMs.

Recent developments

Several new types of *non-volatile* RAM, which will preserve data while powered down, are under development. The technologies used include carbon nanotubes and approaches utilizing the magnetic tunnel effect. Amongst the 1st generation MRAM, a 128 KiB (128×2^{10} bytes) magnetic RAM (MRAM) chip was manufactured with 0.18 μm technology in the summer of 2003. In June 2004, Infineon Technologies unveiled a 16 MiB (16×2^{20} bytes) prototype again based on 0.18 μm technology. There are two 2nd generation techniques currently in development: Thermal Assisted Switching (TAS) which is being developed by Crocus Technology, and Spin Torque Transfer (STT) on which Crocus, Hynix, IBM, and several other companies are working. Nantero built a functioning carbon nanotube memory prototype 10 GiB (10×2^{30} bytes) array in 2004. Whether some of these technologies will be able to eventually take a significant market share from either DRAM, SRAM, or flash-memory technology, however, remains to be seen.

Since 2006, "Solid-state drives" (based on flash memory) with capacities exceeding 256 gigabytes and performance far exceeding traditional disks have become available. This development has started to blur the definition between traditional random access memory and "disks", dramatically reducing the difference in performance.

Some kinds of random-access memory, such as "EcoRAM", are specifically designed for server farms, where low power consumption is more important than speed.

Memory wall

The "memory wall" is the growing disparity of speed between CPU and memory outside the CPU chip. An important reason for this disparity is the limited communication bandwidth beyond chip boundaries. From 1986 to 2000, CPU speed improved at an annual rate of 55% while memory speed only improved at 10%. Given these trends, it was expected that memory latency would become an overwhelming bottleneck in computer performance.

Currently, CPU speed improvements have slowed significantly partly due to major physical barriers and partly because current CPU designs have already hit the memory wall in some sense. Intel summarized these causes in their Platform 2015 documentation (PDF)

“First of all, as chip geometries shrink and clock frequencies rise, the transistor leakage current increases, leading to excess power consumption and heat... Secondly, the advantages of higher clock speeds are in part negated by memory latency, since memory access times have not been able to keep pace with increasing clock frequencies. Third, for certain applications, traditional serial architectures are becoming less efficient as processors get faster (due to the so-called Von Neumann bottleneck), further undercutting any gains that frequency increases might otherwise buy. In addition, partly due to limitations in the means of producing inductance within solid state devices, resistance-

capacitance (RC) delays in signal transmission are growing as feature sizes shrink, imposing an additional bottleneck that frequency increases don't address.”

The RC delays in signal transmission were also noted in Clock Rate versus IPC: The End of the Road for Conventional Microarchitectures which projects a maximum of 12.5% average annual CPU performance improvement between 2000 and 2014. The data on Intel Processors clearly shows a slowdown in performance improvements in recent processors. However, Intel's Core 2 Duo processors (codenamed Conroe) showed a significant improvement over previous Pentium 4 processors; due to a more efficient architecture, performance increased while clock rate actually decreased.

WWT

Chapter 11

Graphics Processing Unit



GeForce 6600GT (NV43) GPU

A **graphics processing unit** or **GPU** (also occasionally called **visual processing unit** or **VPU**) is a specialized circuit designed to rapidly manipulate and alter memory in such a way so as to accelerate the building of images in a frame buffer intended for output to a

display. This specialization allows a GPU to perform such tasks more rapidly than a less specialized and more general central (micro-)processor. GPUs are used in embedded systems, mobile phones, personal computers, workstations, and game consoles. Modern GPUs are very efficient at manipulating computer graphics, and their highly parallel structure makes them more effective than general-purpose CPUs for algorithms where processing of large blocks of data is done in parallel. In a personal computer, a GPU can be present on a video card, or it can be on the motherboard, or in certain CPUs, on the CPU die. More than 90% of new desktop and notebook computers have integrated GPUs, which are usually far less powerful than those on a dedicated video card.

The term was defined and popularized by Nvidia in 1999, who marketed the GeForce 256 as "the world's first 'GPU', or Graphics Processing Unit, a single-chip processor with integrated transform, lighting, triangle setup/clipping, and rendering engines that is capable of processing a minimum of 10 million polygons per second." Rival ATI Technologies coined the term visual processing unit or VPU with the release of the Radeon 9700 in 2002.

History

Graphics accelerators

A GPU (Graphics Processing Unit) is a graphics coprocessor that offloads graphics tasks from the CPU with the intention of performing them faster than the CPU can perform them. A graphics accelerator incorporates one or more custom microchips which contain special mathematical operations commonly used in graphics rendering. The efficiency and clock speed of the microchips and the algorithms implemented therefore determines the effectiveness of the graphics accelerator. They were once mainly used for playing 3D games or high-end 3D rendering, but now take on more roles and therefore have grown more popular as a result. A GPU implements a number of graphics primitive operations in a way that makes running them much faster than drawing directly to the screen with the host CPU. The most common operations for early 2D computer graphics include the BitBLT operation, combining several bitmap patterns using a raster op, usually in special hardware called a "*blitter*", and operations for drawing rectangles, triangles, circles, and arcs. Modern GPUs also have support for 3D computer graphics, and typically include digital video-related functions. More recent GPUs now support handling highly parallel tasks that CPUs would take a long time to perform.

1980s

The Commodore Amiga was the first personal computer to use a full-featured GPU, called a blitter, and the first US patent filing to use the term "blitter" was "Personal computer apparatus for block transfer of bit-mapped image data," assigned to Commodore-Amiga, Inc. The Amiga was capable of offloading practically all image generation functions to hardware, including line drawing, area filling, and block image transfer and manipulation. Also included was a graphics coprocessor with its own

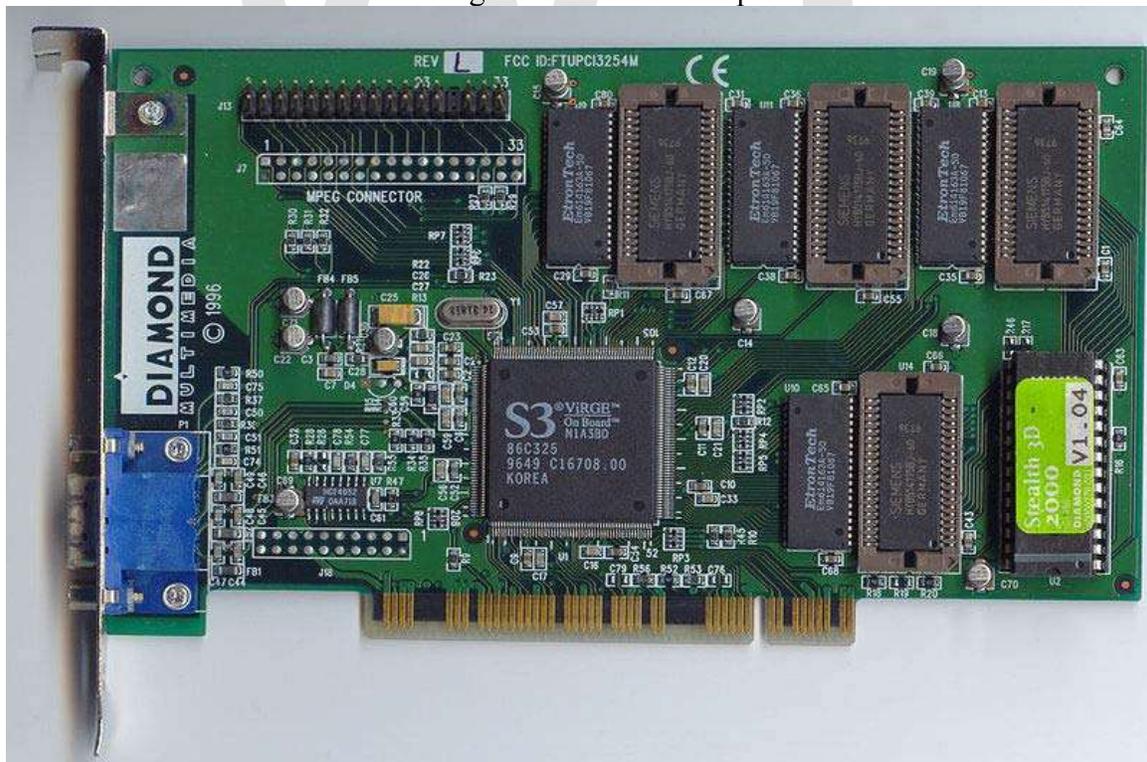
(primitive) instruction set. Prior to this and quite some time after, many other systems required a general purpose CPU to handle every aspect of drawing the display.

In 1987, IBM's 8514 graphics system was released as one of the first video cards for PC compatibles to implement 2D primitives in hardware.

1990s



Tseng Labs ET4000/W32p



S3 Graphics ViRGE



Voodoo3 2000 AGP card

In 1991, S3 Graphics introduced the *S3 86C911*, which its designers named after the Porsche 911 as an indication of the performance increase it promised. The 86C911 spawned a host of imitators: by 1995, all major PC graphics chip makers had added 2D acceleration support to their chips. By this time, fixed-function *Windows accelerators* had surpassed expensive general-purpose graphics coprocessors in Windows performance, and these coprocessors faded away from the PC market.

Throughout the 1990s, 2D GUI acceleration continued to evolve. As manufacturing capabilities improved, so did the level of integration of graphics chips. Additional application programming interfaces (APIs) arrived for a variety of tasks, such as Microsoft's WinG graphics library for Windows 3.x, and their later DirectDraw interface for hardware acceleration of 2D games within Windows 95 and later.

In the early and mid-1990s, CPU-assisted real-time 3D graphics were becoming increasingly common in computer and console games, which led to an increasing public demand for hardware-accelerated 3D graphics. Early examples of mass-marketed 3D graphics hardware can be found in fifth generation video game consoles such as PlayStation and Nintendo 64. In the PC world, notable failed first-tries for low-cost 3D graphics chips were the S3 *ViRGE*, ATI *Rage*, and Matrox *Mystique*. These chips were essentially previous-generation 2D accelerators with 3D features bolted on. Many were even pin-compatible with the earlier-generation chips for ease of implementation and minimal cost. Initially, performance 3D graphics were possible only with discrete boards

dedicated to accelerating 3D functions (and lacking 2D GUI acceleration entirely) such as the 3dfx *Voodoo*. However, as manufacturing technology again progressed, video, 2D GUI acceleration, and 3D functionality were all integrated into one chip. Rendition's *Verite* chipsets were the first to do this well enough to be worthy of note.

OpenGL appeared in the early 90s as a professional graphics API, but originally suffered from performance issues which allowed the Glide API to step in and become a dominant force on the PC in the late 90s. However these issues were quickly overcome and the Glide API fell by the wayside. Software implementations of OpenGL were common during this time although the influence of OpenGL eventually led to widespread hardware support. Over time a parity emerged between features offered in hardware and those offered in OpenGL. DirectX became popular among Windows game developers during the late 90s. Unlike OpenGL, Microsoft insisted on providing strict one-to-one support of hardware. The approach made DirectX less popular as a stand alone graphics API initially since many GPUs provided their own specific features, which existing OpenGL applications were already able to benefit from, leaving DirectX often one generation behind.

Over time Microsoft began to work more closely with hardware developers, and started to target the releases of DirectX with those of the supporting graphics hardware. Direct3D 5.0 was the first version of the burgeoning API to gain widespread adoption in the gaming market, and it competed directly with many more hardware specific, often proprietary graphics libraries, while OpenGL maintained a strong following. Direct3D 7.0 introduced support for hardware-accelerated transform and lighting (T&L) for Direct3D, while OpenGL already had this capability already exposed from its inception. 3D accelerators moved beyond being just simple rasterizers to add another significant hardware stage to the 3D rendering pipeline. The NVIDIA *GeForce 256* (also known as NV10) was the first consumer-level card on the market with hardware-accelerated T&L, while professional 3D cards already had this capability. Hardware transform and lighting, both already existing features of OpenGL, came to consumer-level hardware in the 90s and set the precedent for later pixel shader and vertex shader units which were far more flexible and programmable.

2000 to present

With the advent of the OpenGL API and similar functionality in DirectX, GPUs added programmable shading to their capabilities. Each pixel could now be processed by a short program that could include additional image textures as inputs, and each geometric vertex could likewise be processed by a short program before it was projected onto the screen. NVIDIA was first to produce a chip capable of programmable shading, the *GeForce 3* (code named NV20). By October 2002, with the introduction of the ATI *Radeon 9700* (also known as R300), the world's first Direct3D 9.0 accelerator, pixel and vertex shaders could implement looping and lengthy floating point math, and in general were quickly becoming as flexible as CPUs, and orders of magnitude faster for image-array operations. Pixel shading is often used for things like bump mapping, which adds texture, to make an object look shiny, dull, rough, or even round or extruded.

As the processing power of GPUs has increased, so has their demand for electrical power. High performance GPUs often consume more energy than current CPUs.

Today, parallel GPUs have begun making computational inroads against the CPU, and a subfield of research, dubbed GPU Computing or GPGPU for *General Purpose Computing on GPU*, has found its way into fields as diverse as oil exploration, scientific image processing, linear algebra, statistics, 3D reconstruction and even stock options pricing determination. Nvidia's CUDA platform is the most widely adopted programming model for GPU computing, with OpenCL also being offered as an open standard.

GPU companies

Many companies have produced GPUs under a number of brand names. In 2008, Intel, NVIDIA and AMD/ATI were the market share leaders, with 49.4%, 27.8% and 20.6% market share respectively. However, those numbers include Intel's integrated graphics solutions as GPUs. Not counting those numbers, NVIDIA and ATI control nearly 100% of the market. In addition, S3 Graphics, VIA Technologies and Matrox produce GPUs.

Computational functions

Modern GPUs use most of their transistors to do calculations related to 3D computer graphics. They were initially used to accelerate the memory-intensive work of texture mapping and rendering polygons, later adding units to accelerate geometric calculations such as the rotation and translation of vertices into different coordinate systems. Recent developments in GPUs include support for programmable shaders which can manipulate vertices and textures with many of the same operations supported by CPUs, oversampling and interpolation techniques to reduce aliasing, and very high-precision color spaces. Because most of these computations involve matrix and vector operations, engineers and scientists have increasingly studied the use of GPUs for non-graphical calculations.

In addition to the 3D hardware, today's GPUs include basic 2D acceleration and framebuffer capabilities (usually with a VGA compatibility mode).

GPU accelerated video decoding

Most GPUs made since 1995 support the YUV color space and hardware overlays, important for digital video playback, and many GPUs made since 2000 also support MPEG primitives such as motion compensation and iDCT. This process of hardware accelerated video decoding, where portions of the video decoding process and video post-processing are offloaded to the GPU hardware, is commonly referred to as "**GPU accelerated video decoding**", "**GPU assisted video decoding**", "**GPU hardware accelerated video decoding**" or "**GPU hardware assisted video decoding**".

More recent graphics cards even decode high-definition video on the card, offloading the central processing unit. The most common API's for GPU accelerated video decoding are DxVA for Microsoft Windows operating-system, and VDPAU, VAAPI, XvMC, and

XvBA for Linux and UNIX based operating-system. All except XvMC are capable of decoding videos encoded with MPEG-1, MPEG-2, MPEG-4 ASP (MPEG-4 Part 2), MPEG-4 AVC (H.264 / DivX 6), VC-1, WMV3/WMV9, Xvid / OpenDivX (DivX 4), and DivX 5 codecs, while XvMC is only capable of decoding MPEG-1 and MPEG-2.

Video decoding processes that can be accelerated

The video decoding processes that can be accelerated by today's modern GPU hardware are:

- Motion compensation (mocomp)
- Inverse discrete cosine transform (iDCT)
 - Inverse telecine 3:2 and 2:2 pull-down correction
- Inverse modified discrete cosine transform (iMDCT)
- In-loop deblocking filter
- Intra-frame prediction
- Inverse quantization (IQ)
- Variable-Length Decoding (VLD), more commonly known as slice-level acceleration
- Spatial-temporal deinterlacing and automatic interlace/progressive source detection
- Bitstream processing (CAVLC/CABAC).

GPU forms

Dedicated graphics cards

The GPUs of the most powerful class typically interface with the motherboard by means of an expansion slot such as PCI Express (PCIe) or Accelerated Graphics Port (AGP) and can usually be replaced or upgraded with relative ease, assuming the motherboard is capable of supporting the upgrade. A few graphics cards still use Peripheral Component Interconnect (PCI) slots, but their bandwidth is so limited that they are generally used only when a PCIe or AGP slot is not available.

A dedicated GPU is not necessarily removable, nor does it necessarily interface with the motherboard in a standard fashion. The term "dedicated" refers to the fact that dedicated graphics cards have RAM that is dedicated to the card's use, not to the fact that *most* dedicated GPUs are removable. Dedicated GPUs for portable computers are most commonly interfaced through a non-standard and often proprietary slot due to size and weight constraints. Such ports may still be considered PCIe or AGP in terms of their logical host interface, even if they are not physically interchangeable with their counterparts.

Technologies such as SLI by NVIDIA and CrossFire by ATI allow multiple GPUs to be used to draw a single image, increasing the processing power available for graphics.

Integrated graphics solutions

Integrated graphics solutions, shared graphics solutions, or Integrated graphics processors (IGP) utilize a portion of a computer's system RAM rather than dedicated graphics memory. Exceptions are AMD's IGPs that use dedicated sideport memory on certain motherboards. Computers with integrated graphics account for 90% of all PC shipments. These solutions are less costly to implement than dedicated graphics solutions, but are less capable. Historically, integrated solutions were often considered unfit to play 3D games or run graphically intensive programs but could run less intensive programs such as Adobe Flash. Examples of such IGPs would be offerings from SiS and VIA circa 2004. However, today's integrated solutions such as AMD's Radeon HD 3200 (AMD 780G chipset) and NVIDIA's GeForce 8200 (nForce 710|NVIDIA nForce 730a) are more than capable of handling 2D graphics from Adobe Flash or low stress 3D graphics. However, most integrated graphics still struggle with high-end video games. Chips like the Nvidia GeForce 9400M in Apple's MacBook and MacBook Pro and AMD's Radeon HD 3300 (AMD 790GX) have an improved performance, but still lag behind dedicated graphics cards. Modern desktop motherboards often include an integrated graphics solution and have expansion slots available to add a dedicated graphics card later.

As a GPU is extremely memory intensive, an integrated solution may find itself competing for the already relatively slow system RAM with the CPU, as it has minimal or no dedicated video memory. System RAM may be 2 Gbit/s to 12.8 Gbit/s, yet dedicated GPUs enjoy between 10 Gbit/s to over 100 Gbit/s of bandwidth depending on the model.

Older integrated graphics chipsets lacked hardware transform and lighting, but newer ones include it.

Hybrid solutions

This newer class of GPUs competes with integrated graphics in the low-end desktop and notebook markets. The most common implementations of this are ATI's HyperMemory and NVIDIA's TurboCache. Hybrid graphics cards are somewhat more expensive than integrated graphics, but much less expensive than dedicated graphics cards. These share memory with the system and have a small dedicated memory cache, to make up for the high latency of the system RAM. Technologies within PCI Express can make this possible. While these solutions are sometimes advertised as having as much as 768MB of RAM, this refers to how much can be shared with the system memory.

Stream Processing and General Purpose GPUs (GPGPU)

A new concept is to use a general purpose graphics processing unit as a modified form of stream processor. This concept turns the massive floating-point computational power of a modern graphics accelerator's shader pipeline into general-purpose computing power, as opposed to being hard wired solely to do graphical operations. In certain applications

requiring massive vector operations, this can yield several orders of magnitude higher performance than a conventional CPU. The two largest discrete GPU designers, ATI and NVIDIA, are beginning to pursue this new approach with an array of applications. Both nVidia and ATI have teamed with Stanford University to create a GPU-based client for the Folding@Home distributed computing project, for protein folding calculations. In certain circumstances the GPU calculates forty times faster than the conventional CPUs traditionally used by such applications.

Furthermore, GPU-based high performance computers are starting to play a significant role in large-scale modelling. Three of the 5 most powerful supercomputers in the world take advantage of GPU acceleration. This includes the current leader as of October 2010, Tianhe-1A, which uses the NVIDIA Tesla platform.

Recently NVidia began releasing cards supporting an API extension to the C programming language CUDA ("Compute Unified Device Architecture"), which allows specified functions from a normal C program to run on the GPU's stream processors. This makes C programs capable of taking advantage of a GPU's ability to operate on large matrices in parallel, while still making use of the CPU when appropriate. CUDA is also the first API to allow CPU-based applications to access directly the resources of a GPU for more general purpose computing without the limitations of using a graphics API.

Since 2005 there has been interest in using the performance offered by GPUs for evolutionary computation in general, and for accelerating the fitness evaluation in genetic programming in particular. Most approaches compile linear or tree programs on the host PC and transfer the executable to the GPU to be run. Typically the performance advantage is only obtained by running the single active program simultaneously on many example problems in parallel, using the GPU's SIMD architecture. However, substantial acceleration can also be obtained by not compiling the programs, and instead transferring them to the GPU, to be interpreted there. Acceleration can then be obtained by either interpreting multiple programs simultaneously, simultaneously running multiple example problems, or combinations of both. A modern GPU (*e.g.* 8800 GTX or later) can readily simultaneously interpret hundreds of thousands of very small programs.

Chapter 12

Hard Disk Drive

Hard disk drive



Interior of a hard disk drive

Date invented December 24, 1954

Invented by An IBM team led by Rey Johnson



An opened hard drive

A **hard disk drive** (HDD) is a non-volatile, random access device for digital data. It features rotating rigid platters on a motor-driven spindle within a protective enclosure.

Data is magnetically read from and written to the platter by read/write heads that float on a film of air above the platters.

Introduced by IBM in 1956, hard disk drives have fallen in cost and physical size over the years while dramatically increasing in capacity. Hard disk drives have been the dominant device for secondary storage of data in general purpose computers since the early 1960s. They have maintained this position because advances in their areal recording density have kept pace with the requirements for secondary storage. Today's HDDs operate on high-speed serial interfaces; i.e., serial ATA (SATA) or serial attached SCSI (SAS).

History

Hard disk drives were introduced in 1956 as data storage for an IBM accounting computer and were developed for use with general purpose mainframe and mini computers.

Driven by areal density doubling every two to four years since their invention, HDDs have changed in many ways, a few highlights include:

- Capacity per HDD increasing from 3.75 megabytes to greater than 1 terabyte, a greater than 270 thousand to 1 improvement.
- Size of HDD decreasing from 87.9 cubic feet (a double wide refrigerator) to 0.002 cubic feet (2½-inch form factor, a pack of cards), a greater than 44 thousand to 1 improvement.
- Price decreasing from about \$15,000 per megabyte to less than \$0.0001 per megabyte (\$100/1 terabyte), a greater than 150 million to 1 improvement.
- Average access time decreasing from greater than 0.1 second to a few thousandths of a second, a greater than 40 to 1 improvement.
- Market application expanding from general purpose computers to most computing applications including consumer applications.

Technology

Magnetic recording

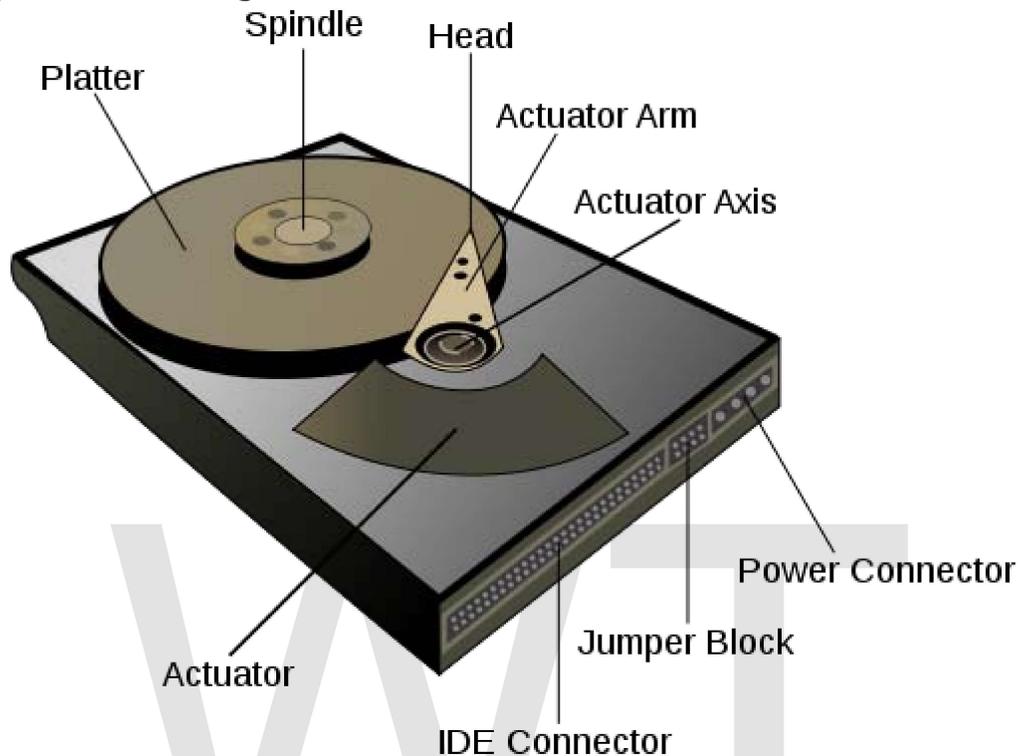
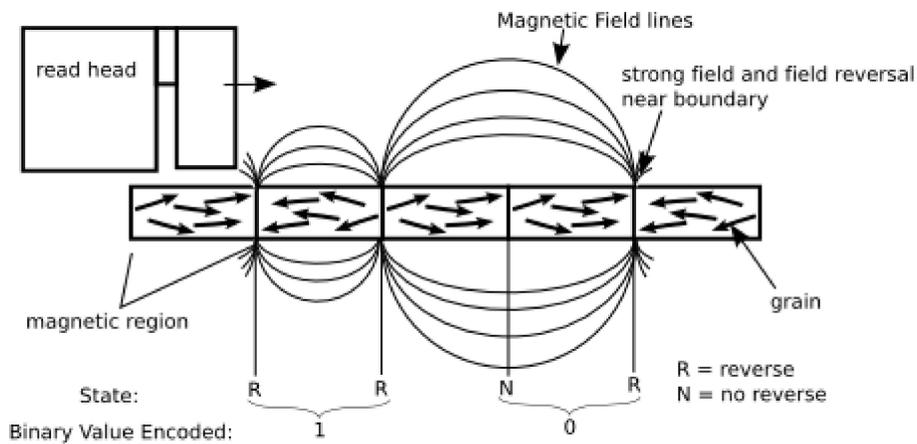


Diagram of a computer hard disk drive

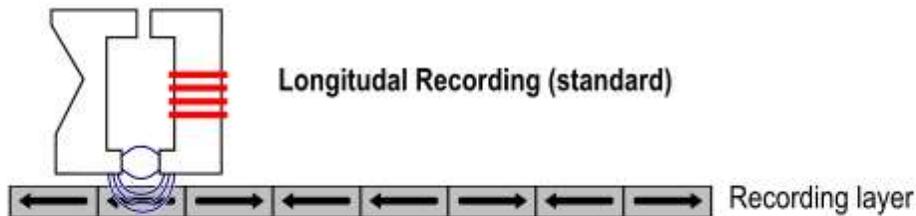
HDDs record data by magnetizing ferromagnetic material directionally. Sequential changes in the direction of magnetization represent patterns of binary data bits. The data are read from the disk by detecting the transitions in magnetization and decoding the originally written data. Different encoding schemes, such as Modified Frequency Modulation, group code recording, run-length limited encoding, and others are used.

A typical HDD design consists of a spindle that holds flat circular disks called platters, onto which the data are recorded. The platters are made from a non-magnetic material, usually aluminum alloy or glass, and are coated with a shallow layer of magnetic material typically 10–20 nm in depth, with an outer layer of carbon for protection. For reference, standard copy paper is 0.07–0.18 millimetre (70,000–180,000 nm).

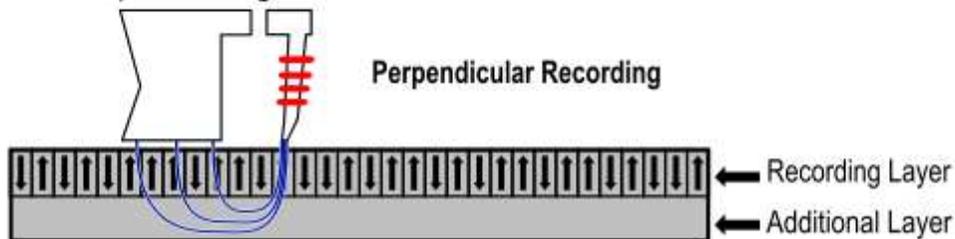


A cross section of the magnetic surface in action. In this case the binary data are encoded using frequency modulation.

"Ring" writing element



"Monopole" writing element



Perpendicular recording

The platters are spun at speeds varying from 3,000 RPM in energy-efficient portable devices, to 15,000 RPM for high performance servers. Information is written to, and read from a platter as it rotates past devices called read-and-write heads that operate very close (tens of nanometers in new drives) over the magnetic surface. The read-and-write head is used to detect and modify the magnetization of the material immediately under it. In modern drives there is one head for each magnetic platter surface on the spindle, mounted on a common arm. An actuator arm (or access arm) moves the heads on an arc (roughly

radially) across the platters as they spin, allowing each head to access almost the entire surface of the platter as it spins. The arm is moved using a voice coil actuator or in some older designs a stepper motor.

The magnetic surface of each platter is conceptually divided into many small sub-micrometer-sized magnetic regions referred to as magnetic domains. In older disk designs the regions were oriented horizontally and parallel to the disk surface, but beginning about 2005, the orientation was changed to perpendicular to allow for closer magnetic domain spacing. Due to the polycrystalline nature of the magnetic material each of these magnetic regions is composed of a few hundred magnetic grains. Magnetic grains are typically 10 nm in size and each form a single magnetic domain. Each magnetic region in total forms a magnetic dipole which generates a magnetic field.

For reliable storage of data, the recording material needs to resist self-demagnetization, which occurs when the magnetic domains repel each other. Magnetic domains written too densely together to a weakly magnetizable material will degrade over time due to physical rotation of one or more domains to cancel out these forces. The domains rotate sideways to a halfway position that weakens the readability of the domain and relieves the magnetic stresses. Older hard disks used iron(III) oxide as the magnetic material, but current disks use a cobalt-based alloy.

A write head magnetizes a region by generating a strong local magnetic field. Early HDDs used an electromagnet both to magnetize the region and to then read its magnetic field by using electromagnetic induction. Later versions of inductive heads included metal in Gap (MIG) heads and thin film heads. As data density increased, read heads using magnetoresistance (MR) came into use; the electrical resistance of the head changed according to the strength of the magnetism from the platter. Later development made use of spintronics; in these heads, the magnetoresistive effect was much greater than in earlier types, and was dubbed "giant" magnetoresistance (GMR). In today's heads, the read and write elements are separate, but in close proximity, on the head portion of an actuator arm. The read element is typically magneto-resistive while the write element is typically thin-film inductive.

The heads are kept from contacting the platter surface by the air that is extremely close to the platter; that air moves at or near the platter speed. The record and playback head are mounted on a block called a slider, and the surface next to the platter is shaped to keep it just barely out of contact. This forms a type of air bearing.

In modern drives, the small size of the magnetic regions creates the danger that their magnetic state might be lost because of thermal effects. To counter this, the platters are coated with two parallel magnetic layers, separated by a 3-atom layer of the non-magnetic element ruthenium, and the two layers are magnetized in opposite orientation, thus reinforcing each other. Another technology used to overcome thermal effects to allow greater recording densities is perpendicular recording, first shipped in 2005, and as of 2007 the technology was used in many HDDs.

Components



A hard disk drive with the disks and motor hub removed showing the copper colored stator coils surrounding a bearing at the center of the spindle motor. The orange stripe along the side of the arm is a thin printed-circuit cable. The spindle bearing is in the center. The actuator is in the upper left.

A typical hard disk drive has two electric motors; a disk motor to spin the disks and an actuator (motor) to position the read/write head assembly across the spinning disks.

The disk motor has an external rotor attached to the disks; the stator windings are fixed in place.

Opposite the actuator at the end of the head support arm is the read-write head (near center in photo); thin printed-circuit cables connect the read-write heads to amplifier electronics mounted at the pivot of the actuator. A flexible, somewhat U-shaped, ribbon cable, seen edge-on below and to the left of the actuator arm continues the connection to the controller board on the opposite side.

The head support arm is very light, but also stiff; in modern drives, acceleration at the head reaches 550 Gs.

The silver-colored structure at the upper left of the first image is the top plate of the actuator, a permanent-magnet and moving coil motor that swings the heads to the desired

position (it is shown removed in the second image). The plate supports a squat neodymium-iron-boron (NIB) high-flux magnet. Beneath this plate is the moving coil, often referred to as the *voice coil* by analogy to the coil in loudspeakers, which is attached to the actuator hub, and beneath that is a second NIB magnet, mounted on the bottom plate of the motor (some drives only have one magnet).

The voice coil itself is shaped rather like an arrowhead, and made of doubly coated copper magnet wire. The inner layer is insulation, and the outer is thermoplastic, which bonds the coil together after it is wound on a form, making it self-supporting. The portions of the coil along the two sides of the arrowhead (which point to the actuator bearing center) interact with the magnetic field, developing a tangential force that rotates the actuator. Current flowing radially outward along one side of the arrowhead and radially inward on the other produces the tangential force. If the magnetic field were uniform, each side would generate opposing forces that would cancel each other out. Therefore the surface of the magnet is half N pole, half S pole, with the radial dividing line in the middle, causing the two sides of the coil to see opposite magnetic fields and produce forces that add instead of canceling. Currents along the top and bottom of the coil produce radial forces that do not rotate the head.

Error handling

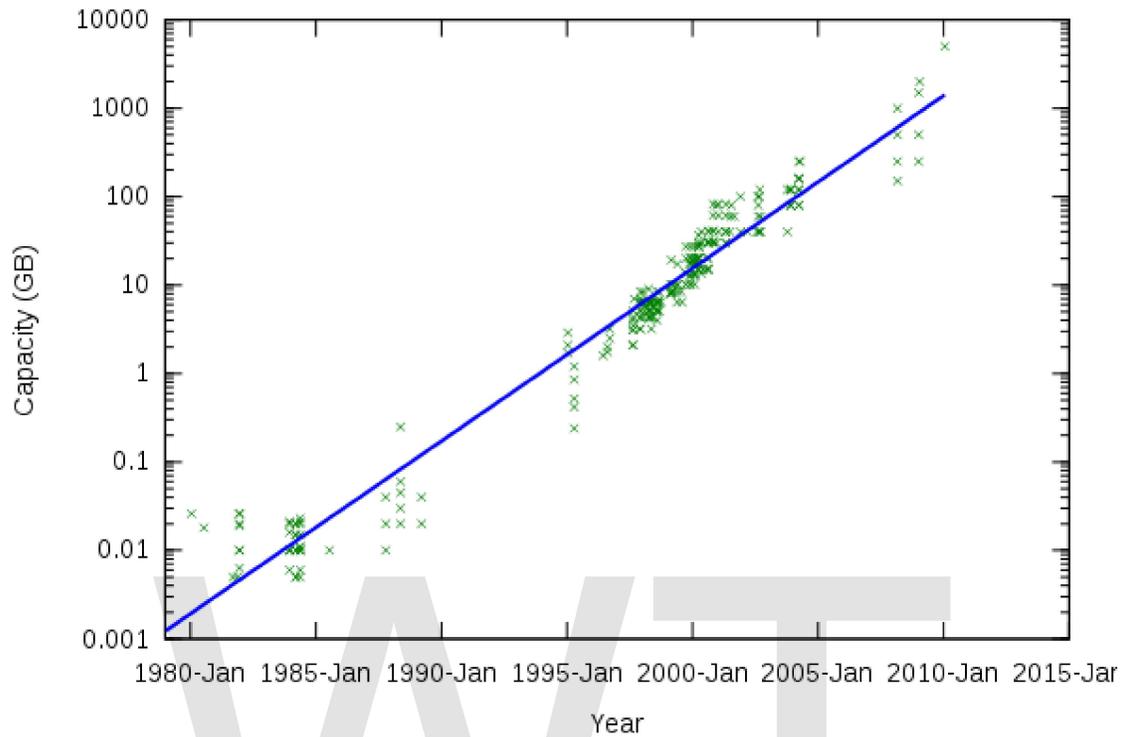
Modern drives also make extensive use of Error Correcting Codes (ECCs), particularly Reed–Solomon error correction. These techniques store extra bits for each block of data that are determined by mathematical formulas. The extra bits allow many errors to be fixed. While these extra bits take up space on the hard drive, they allow higher recording densities to be employed, resulting in much larger storage capacity for user data. In 2009, in the newest drives, low-density parity-check codes (LDPC) are supplanting Reed-Solomon. LDPC codes enable performance close to the Shannon Limit and thus allow for the highest storage density available.

Typical hard drives attempt to "remap" the data in a physical sector that is going bad to a spare physical sector—hopefully while the errors in that bad sector are still few enough that the ECC can recover the data without loss. The S.M.A.R.T. system counts the total number of errors in the entire hard drive fixed by ECC, and the total number of remappings, in an attempt to predict hard drive failure.

Future development

Because of bit-flipping errors and other issues, perpendicular recording densities may be supplanted by other magnetic recording technologies. Toshiba is promoting bit-patterned recording (BPR), while Xyratex are developing heat-assisted magnetic recording (HAMR).

Capacity



PC hard disk drive capacity (in GB) over time. The vertical axis is logarithmic, so the fit line corresponds to exponential growth.

Capacity measurements



A disassembled and labeled 1997 hard drive. All major components were placed on a mirror, which created the symmetrical reflections.

Hard disk manufacturers quote disk capacity in multiples of SI-standard powers of 1000, where a *terabyte* is 1000 gigabytes and a *gigabyte* is 1000 megabytes. With file systems that report capacity in powers of 1024, available space appears somewhat less than advertised capacity. The discrepancy between the two methods of reporting sizes had serious financial consequences for at least one hard drive manufacturer when a class action suit argued the different methods effectively misled consumers.

Semiconductor memory chips are organized so that memory sizes are expressed in multiples of powers of two. Hard disks by contrast have no inherent binary size. Capacity is the product of the number of heads, number of tracks, number of sectors per track, and the size of each sector. Sector sizes are standardized for convenience at 256 or 512 and more recently 4096 bytes, which are powers of two. This can cause some confusion because operating systems may report the formatted capacity of a hard drive using binary prefix units which increment by powers of 1024. For example, Microsoft Windows reports disk capacity both in a decimal integer to 12 or more digits and in binary prefix units to three significant digits.

A one terabyte (1 TB) disk drive would be expected to hold around 1 trillion bytes (1,000,000,000,000) or 1000 GB; and indeed most 1 TB hard drives will contain slightly more than this number. However some operating system utilities would report this as around 931 GB or 953,674 MB. (The actual number for a formatted capacity will be somewhat smaller still, depending on the file system.) Following are the several ways of reporting one Terabyte.

SI prefixes (hard drive)	equivalent	Binary prefixes (OS)	equivalent
1 TB (Terabyte)	$1 * 1000^4 \text{ B}$	0.9095 TiB (Tebibyte)	$0.9095 * 1024^4 \text{ B}$
1000 GB (Gigabyte)	$1000 * 1000^3 \text{ B}$	931.3 GiB (Gibibyte)	$931.3 * 1024^3 \text{ B}$
1,000,000 MB (Megabyte)	$1,000,000 * 1000^2 \text{ B}$	953,674.3 MiB (Mebibyte)	$953,674.3 * 1024^2 \text{ B}$
1,000,000,000 kB (Kilobyte)	$1,000,000,000 * 1000 \text{ B}$	976,562,500 KiB (Kibibyte)	$976,562,500 * 1024 \text{ B}$
1,000,000,000,000 B (byte)	-	1,000,000,000,000 B (byte)	-

Addressing data on large drives

The capacity of an HDD can be calculated by multiplying the number of cylinders by the number of heads by the number of sectors by the number of bytes/sector (most commonly 512). Drives with the ATA interface and a capacity of eight gigabytes or more behave as if they were structured into 16383 cylinders, 16 heads, and 63 sectors, for compatibility with older operating systems. Unlike in the 1980s, the cylinder, head, sector (C/H/S) counts reported to the CPU by a modern ATA drive are no longer actual physical parameters since the reported numbers are constrained by historic operating-system interfaces and with zone bit recording the actual number of sectors varies by zone. Disks

with SCSI interface address each sector with a unique integer number; the operating system remains ignorant of their head or cylinder count.

The old C/H/S scheme has been replaced by logical block addressing. In some cases, to try to "force-fit" the C/H/S scheme to large-capacity drives, the number of heads was given as 64, although no modern drive has anywhere near 32 platters.

Not all the space on a hard drive is available for user files. The operating system file system uses some of the disk space to organize files on the disk, recording their file names and the sequence of disk areas that represent the file. Examples of data structures stored on disk to retrieve files include the MS DOS file allocation table (FAT), and UNIX inodes, as well as other operating system data structures. This file system overhead is usually less than 1% on drives larger than 100 MB.

For RAID drives, data integrity and fault-tolerance requirements also reduce the realized capacity. For example, a RAID1 drive will be about half the total capacity as a result of data mirroring. For RAID5 drives with x drives you would lose $1/x$ of your space to parity. RAID drives are multiple drives that appear to be one drive to the user, but provides some fault-tolerance.

A general rule of thumb to quickly convert the manufacturer's hard disk capacity to the standard Microsoft Windows formatted capacity is $0.93 * \text{capacity of HDD from manufacturer}$ for HDDs less than a terabyte and $0.91 * \text{capacity of HDD from manufacturer}$ for HDDs equal to or greater than 1 terabyte.

HDD Formatting

The presentation of an HDD to its host is determined by its controller. This may differ substantially from the drive's native interface particularly in mainframes or servers.

Modern HDDs, such as SAS and SATA drives, appear at their interfaces as a contiguous set of logical blocks; typically 512 bytes long but the industry is in the process of changing to 4,096 byte logical blocks.

The process of initializing these logical blocks on the physical disk platters is called *low level formatting* which is usually performed at the factory and is not normally changed in the field. *High level formatting* then writes the file system structures into selected logical blocks to make the remaining logical blocks available to the host OS and its applications.

Form factors



5¼" full height 110 MB HDD, 2½" (8.5 mm) 6495 MB HDD, US/UK pennies for comparison.



Six hard drives with 8", 5.25", 3.5", 2.5", 1.8", and 1" disks, partially disassembled to show platters and read-write heads, with a ruler showing inches.

Mainframe and minicomputer hard disks were of widely varying dimensions, typically in free standing cabinets the size of washing machines (e.g. HP 7935 and DEC RP06 Disk Drives) or designed so that dimensions enabled placement in a 19" rack (e.g. Diablo Model 31). In 1962, IBM introduced its model 1311 disk, which used 14 inch (nominal size) platters. This became a standard size for mainframe and minicomputer drives for many years, but such large platters were never used with microprocessor-based systems.

With increasing sales of microcomputers having built in floppy-disk drives (FDDs), HDDs that would fit to the FDD mountings became desirable, and this led to the evolution of the market towards drives with certain **Form factors**, initially derived from the sizes of 8-inch, 5.25-inch, and 3.5-inch floppy disk drives. Smaller sizes than 3.5 inches have emerged as popular in the marketplace and/or been decided by various industry groups.

- **8 inch:** 9.5 in × 4.624 in × 14.25 in (241.3 mm × 117.5 mm × 362 mm)
In 1979, Shugart Associates' SA1000 was the first form factor compatible HDD, having the same dimensions and a compatible interface to the 8" FDD.
- **5.25 inch:** 5.75 in × 3.25 in × 8 in (146.1 mm × 82.55 mm × 203 mm)
This smaller form factor, first used in an HDD by Seagate in 1980, was the same

size as full-height 5¼-inch-diameter (130 mm) FDD, 3.25-inches high. This is twice as high as "half height"; i.e., 1.63 in (41.4 mm). Most desktop models of drives for optical 120 mm disks (DVD, CD) use the half height 5¼" dimension, but it fell out of fashion for HDDs. The Quantum Bigfoot HDD was the last to use it in the late 1990s, with "low-profile" (≈25 mm) and "ultra-low-profile" (≈20 mm) high versions.

- **3.5 inch:** 4 in × 1 in × 5.75 in (101.6 mm × 25.4 mm × 146 mm) = 376.77344 cm³
This smaller form factor, first used in an HDD by Rodime in 1983, was the same size as the "half height" 3½" FDD, i.e., 1.63 inches high. Today it has been largely superseded by 1-inch high "slimline" or "low-profile" versions of this form factor which is used by most desktop HDDs.

- **2.5 inch:** 2.75 in × 0.275–0.59 in × 3.945 in (69.85 mm × 7–15 mm × 100 mm) = 48.895–104.775 cm³

This smaller form factor was introduced by PrairieTek in 1988; there is no corresponding FDD. It is widely used today for hard-disk drives in mobile devices (laptops, music players, etc.) and as of 2008 replacing 3.5 inch enterprise-class drives. It is also used in the Playstation 3 and Xbox 360 video game consoles. Today, the dominant height of this form factor is 9.5 mm for laptop drives (usually having two platters inside), but higher capacity drives have a height of 12.5 mm (usually having three platters). Enterprise-class drives can have a height up to 15 mm. Seagate has released a wafer-thin 7mm drive aimed at entry level laptops and high end netbooks in December 2009.

- **1.8 inch:** 54 mm × 8 mm × 71 mm = 30.672 cm³
This form factor, originally introduced by Integral Peripherals in 1993, has evolved into the ATA-7 LIF with dimensions as stated. It was increasingly used in digital audio players and subnotebooks, but is rarely used today. An original variant exists for 2–5GB sized HDDs that fit directly into a PC card expansion slot. These became popular for their use in iPods and other HDD based MP3 players.

- **1 inch:** 42.8 mm × 5 mm × 36.4 mm
This form factor was introduced in 1999 as IBM's Microdrive to fit inside a CF Type II slot. Samsung calls the same form factor "**1.3 inch**" drive in its product literature.

- **0.85 inch:** 24 mm × 5 mm × 32 mm
Toshiba announced this form factor in January 2004 for use in mobile phones and similar applications, including SD/MMC slot compatible HDDs optimized for video storage on 4G handsets. Toshiba currently sells a 4 GB (MK4001MTD) and 8 GB (MK8003MTD) version and holds the Guinness World Record for the smallest hard disk drive.

3.5-inch and 2.5-inch hard disks currently dominate the market.

By 2009 all manufacturers had discontinued the development of new products for the 1.3-inch, 1-inch and 0.85-inch form factors due to falling prices of flash memory, which is slightly more stable and resistant to damage from impact and/or dropping.

The inch-based nickname of all these form factors usually do not indicate any actual product dimension (which are specified in millimeters for more recent form factors), but just roughly indicate a size relative to disk diameters, in the interest of historic continuity.

Current hard disk form factors

Form factor	Width	Height	Largest capacity	Platters (Max)
3.5"	102 mm	25.4 mm	3 TB (2010)	5
2.5"	69.9 mm	7–15 mm	1.5 TB (2010)	4
1.8"	54 mm	8 mm	320 GB (2009)	3

Obsolete hard disk form factors

Form factor	Width	Largest capacity	Platters (Max)
5.25" FH	146 mm	47 GB (1998)	14
5.25" HH	146 mm	19.3 GB (1998)	4
1.3"	43 mm	40 GB (2007)	1
1" (CFII/ZIF/IDE-Flex)	42 mm	20 GB (2006)	1
0.85"	24 mm	8 GB (2004)	1

Performance characteristics

Access Time

The factors that limit the time to access the data on a hard disk drive (Access time) are mostly related to the mechanical nature of the rotating disks and moving heads. Seek time is a measure of how long it takes the head assembly to travel to the track of the disk that contains data. Latency is rotational delay incurred because the desired disk sector may not be directly under the head when data transfer is requested. These two delays are on the order of milliseconds each. The bit rate or data transfer rate once the head is in the right position creates delay which is a function of the number of blocks transferred; typically relatively small, but can be quite long with the transfer of large contiguous files. Delay may also occur if the drive disks are stopped to save energy.

An HDD's **Average Access Time** is its average Seek time which technically is the time to do all possible seeks divided by the number of all possible seeks, but in practice is determined by statistical methods or simply approximated as the time of a seek over one-third of the number of tracks

Defragmentation is a procedure used to minimize delay in retrieving data by moving related items to physically proximate areas on the disk. Some computer operating systems perform defragmentation automatically. Although automatic defragmentation is

intended to reduce access delays, the procedure can slow response when performed while the computer is in use.

Access time can be improved by increasing rotational speed, thus reducing latency and/or by decreasing seek time. Increasing areal density increases throughput by increasing data rate and by increasing the amount of data under a set of heads, thereby potentially reducing seek activity for a given amount of data. Based on historic trends, analysts predict a future growth in HDD areal density (and therefore capacity) of about 40% per year. Access times have not kept up with throughput increases, which themselves have not kept up with growth in storage capacity.

Seek time

Average Seek time ranges from 3 ms for high-end server drives, to 15 ms for mobile drives, with the most common mobile drives at about 12 ms and the most common desktop type typically being around 9 ms. The first HDD had an average seek time of about 600 ms and by the middle 1970s HDDs were available with seek times of about 25 ms. Some early PC drives used a stepper motor to move the heads, and as a result had seek times as slow as 80–120 ms, but this was quickly improved by voice coil type actuation in the 1980s, reducing seek times to around 20 ms. Seek time has continued to improve slowly over time.

Latency

Latency is the delay for the rotation of the disk to bring the required disk sector under the read-write mechanism. It depends on rotational speed of a disk, measured in revolutions per minute (RPM). Average rotational delay is shown in the table below, based on the empirical relation that the average latency in milliseconds for such a drive is one-half the rotational period:

Spindle [rpm]	Average latency [ms]
4200	7.14
5400	5.56
7200	4.17
10000	3
15000	2

Data transfer rate

As of 2010, a typical 7200 rpm desktop hard drive has a sustained "disk-to-buffer" data transfer rate up to 1030 Mbits/sec. This rate depends on the track location, so it will be higher for data on the outer tracks (where there are more data sectors) and lower toward the inner tracks (where there are fewer data sectors); and is generally somewhat higher for 10,000 rpm drives. A current widely used standard for the "buffer-to-computer" interface is 3.0 Gbit/s SATA, which can send about 300 megabyte/s from the buffer to

the computer, and thus is still comfortably ahead of today's disk-to-buffer transfer rates. Data transfer rate (read/write) can be measured by writing a large file to disk using special file generator tools, then reading back the file. Transfer rate can be influenced by file system fragmentation and the layout of the files.

HDD data transfer rate depends upon the rotational speed of the platters and the data recording density. Because heat and vibration limit rotational speed, advancing density becomes the main method to improve sequential transfer rates. Areal density advances by increasing both the number of tracks across the disk and the number of sectors per track, the later will increase the data transfer rate (for a given RPM). Since data transfer rate performance only tracks one of the two components of areal density, its performance improves at lower rate,

Power consumption

Power consumption has become increasingly important, not only in mobile devices such as laptops but also in server and desktop markets. Increasing data center machine density has led to problems delivering sufficient power to devices (especially for spin up), and getting rid of the waste heat subsequently produced, as well as environmental and electrical cost concerns. Heat dissipation directly tied to power consumption, and as drive age, disk failure rates increase at higher drive temperatures. Similar issues exist for large companies with thousands of desktop PCs. Smaller form factor drives often use less power than larger drives. One interesting development in this area is actively controlling the seek speed so that the head arrives at its destination only just in time to read the sector, rather than arriving as quickly as possible and then having to wait for the sector to come around (i.e. the rotational latency). Many of the hard drive companies are now producing Green Drives that require much less power and cooling. Many of these Green Drives spin slower (<5,400 rpm compared to 7,200, 10,000 or 15,000 rpm) and also generate less waste heat. Power consumption can also be reduced by parking the drive heads when the disk is not in use reducing friction, adjusting spin speeds according to transfer rates, and disabling internal components when not in use.

Also in systems where there might be multiple hard disk drives, there are various ways of controlling when the hard drives spin up since the highest current is drawn at that time.

- On SCSI hard disk drives, the SCSI controller can directly control spin up and spin down of the drives.
- On Parallel ATA (aka PATA) and Serial ATA (SATA) hard disk drives, some support power-up in standby or PUIS. The hard disk drive will not spin up until the controller or system BIOS issues a specific command to do so. This limits the power draw or consumption upon power on.
- Some SATA II hard disk drives support staggered spin-up, allowing the computer to spin up the drives in sequence to reduce load on the power supply when booting.

Power management

Most hard disk drives today support some form of power management which uses a number of specific power modes that save energy by reducing performance. When implemented an HDD will change between a full power mode to one or more power saving modes as a function of drive usage. Recovery from the deepest mode, typically called Sleep, may take as long as several seconds.

Audible noise

Measured in dBA, audible noise is significant for certain applications, such as DVRs, digital audio recording and quiet computers. Low noise disks typically use fluid bearings, slower rotational speeds (usually 5,400 rpm) and reduce the seek speed under load (AAM) to reduce audible clicks and crunching sounds. Drives in smaller form factors (e.g. 2.5 inch) are often quieter than larger drives.

Shock resistance

Shock resistance is especially important for mobile devices. Some laptops now include active hard drive protection that parks the disk heads if the machine is dropped, hopefully before impact, to offer the greatest possible chance of survival in such an event. Maximum shock tolerance to date is 350 g for operating and 1000 g for non-operating.

Access and interfaces

Hard disk drives are accessed over one of a number of bus types, including parallel ATA (P-ATA, also called IDE or EIDE), Serial ATA (SATA), SCSI, Serial Attached SCSI (SAS), and Fibre Channel. Bridge circuitry is sometimes used to connect hard disk drives to buses that they cannot communicate with natively, such as IEEE 1394, USB and SCSI.

For the ST-506 interface, the data encoding scheme as written to the disk surface was also important. The first ST-506 disks used Modified Frequency Modulation (MFM) encoding, and transferred data at a rate of 5 megabits per second. Later controllers using 2,7 RLL (or just "RLL") encoding caused 50% more data to appear under the heads compared to one rotation of an MFM drive, increasing data storage and data transfer rate by 50%, to 7.5 megabits per second.

Many ST-506 interface disk drives were only specified by the manufacturer to run at the 1/3 lower MFM data transfer rate compared to RLL, while other drive models (usually more expensive versions of the same drive) were specified to run at the higher RLL data transfer rate. In some cases, a drive had sufficient margin to allow the MFM specified model to run at the denser/faster RLL data transfer rate (not recommended nor guaranteed by manufacturers). Also, any RLL-certified drive could run on any MFM controller, but with 1/3 less data capacity and as much as 1/3 less data transfer rate compared to its RLL specifications.

Enhanced Small Disk Interface (ESDI) also supported multiple data rates (ESDI disks always used 2,7 RLL, but at 10, 15 or 20 megabits per second), but this was usually negotiated automatically by the disk drive and controller; most of the time, however, 15 or 20 megabit ESDI disk drives were not downward compatible (i.e. a 15 or 20 megabit disk drive would not run on a 10 megabit controller). ESDI disk drives typically also had jumpers to set the number of sectors per track and (in some cases) sector size.

Modern hard drives present a consistent interface to the rest of the computer, no matter what data encoding scheme is used internally. Typically a DSP in the electronics inside the hard drive takes the raw analog voltages from the read head and uses PRML and Reed–Solomon error correction to decode the sector boundaries and sector data, then sends that data out the standard interface. That DSP also watches the error rate detected by error detection and correction, and performs bad sector remapping, data collection for Self-Monitoring, Analysis, and Reporting Technology, and other internal tasks.

SCSI originally had just one signaling frequency of 5 MHz for a maximum data rate of 5 megabytes/second over 8 parallel conductors, but later this was increased dramatically. The SCSI bus speed had no bearing on the disk's internal speed because of buffering between the SCSI bus and the disk drive's internal data bus; however, many early disk drives had very small buffers, and thus had to be reformatted to a different interleave (just like ST-506 disks) when used on slow computers, such as early Commodore Amiga, IBM PC compatibles and Apple Macintoshes.

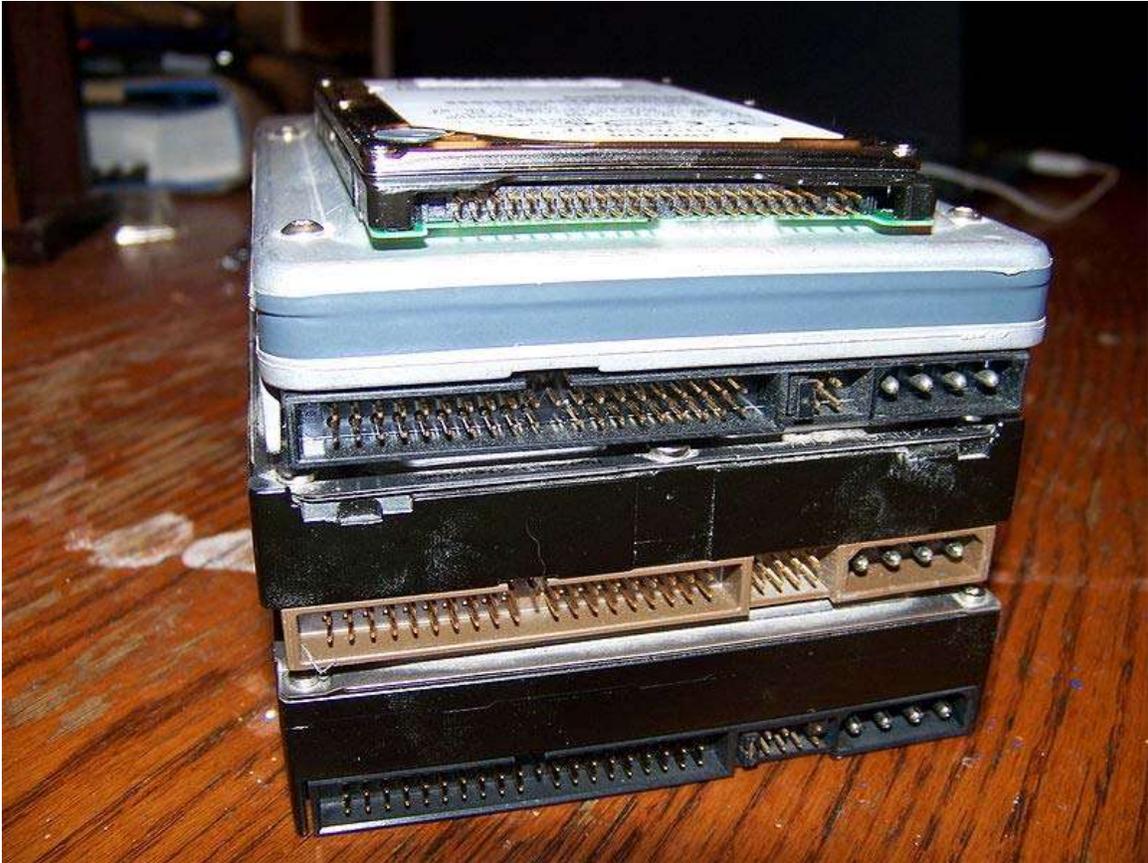
ATA disks have typically had no problems with interleave or data rate, due to their controller design, but many early models were incompatible with each other and could not run with two devices on the same physical cable in a master/slave setup. This was mostly remedied by the mid-1990s, when ATA's specification was standardized and the details began to be cleaned up, but still causes problems occasionally (especially with CD-ROM and DVD-ROM disks, and when mixing Ultra DMA and non-UDMA devices).

Serial ATA does away with master/slave setups entirely, placing each disk on its own channel (with its own set of I/O ports) instead.

FireWire/IEEE 1394 and USB(1.0/2.0) HDDs are external units containing generally ATA or SCSI disks with ports on the back allowing very simple and effective expansion and mobility. Most FireWire/IEEE 1394 models are able to daisy-chain in order to continue adding peripherals without requiring additional ports on the computer itself. USB however, is a point to point network and does not allow for daisy-chaining. USB hubs are used to increase the number of available ports and are used for devices that do not require charging since the current supplied by hubs is typically lower than what's available from the built-in USB ports.

Disk interface families used in personal computers

Notable families of disk interfaces include:



Several Parallel ATA hard disk drives

- Historical **bit serial interfaces** connect a hard disk drive (HDD) to a hard disk controller (HDC) with two cables, one for control and one for data. (Each drive also has an additional cable for power, usually connecting it directly to the power supply unit). The HDC provided significant functions such as serial/parallel conversion, data separation, and track formatting, and required matching to the drive (after formatting) in order to assure reliability. Each control cable could serve two or more drives, while a dedicated (and smaller) data cable served each drive.
 - ST506 used MFM (Modified Frequency Modulation) for the data encoding method.
 - ST412 was available in either MFM or RLL (Run Length Limited) encoding variants.
 - Enhanced Small Disk Interface (ESDI) was an industry standard interface similar to ST412 supporting higher data rates between the processor and the disk drive.
- Modern **bit serial interfaces** connect a hard disk drive to a host bus interface adapter (today typically integrated into the "south bridge") with one data/control cable. (As for historical *bit serial interfaces* above, each drive also has an additional power cable, usually direct to the power supply unit.)

- Fibre Channel (FC), is a successor to parallel SCSI interface on enterprise market. It is a serial protocol. In disk drives usually the Fibre Channel Arbitrated Loop (FC-AL) connection topology is used. FC has much broader usage than mere disk interfaces, and it is the cornerstone of storage area networks (SANs). Recently other protocols for this field, like iSCSI and ATA over Ethernet have been developed as well. Confusingly, drives usually use *copper* twisted-pair cables for Fibre Channel, not fibre optics. The latter are traditionally reserved for larger devices, such as servers or disk array controllers.
- Serial ATA (SATA). The SATA data cable has one data pair for differential transmission of data to the device, and one pair for differential receiving from the device, just like EIA-422. That requires that data be transmitted serially. Similar differential signaling system is used in RS485, LocalTalk, USB, Firewire, and differential SCSI.
- Serial Attached SCSI (SAS). The SAS is a new generation serial communication protocol for devices designed to allow for much higher speed data transfers and is compatible with SATA. SAS uses a mechanically identical data and power connector to standard 3.5-inch SATA1/SATA2 HDDs, and many server-oriented SAS RAID controllers are also capable of addressing SATA hard drives. SAS uses serial communication instead of the parallel method found in traditional SCSI devices but still uses SCSI commands.



Inner view of a 1998 Seagate hard disk drive which used Parallel ATA interface

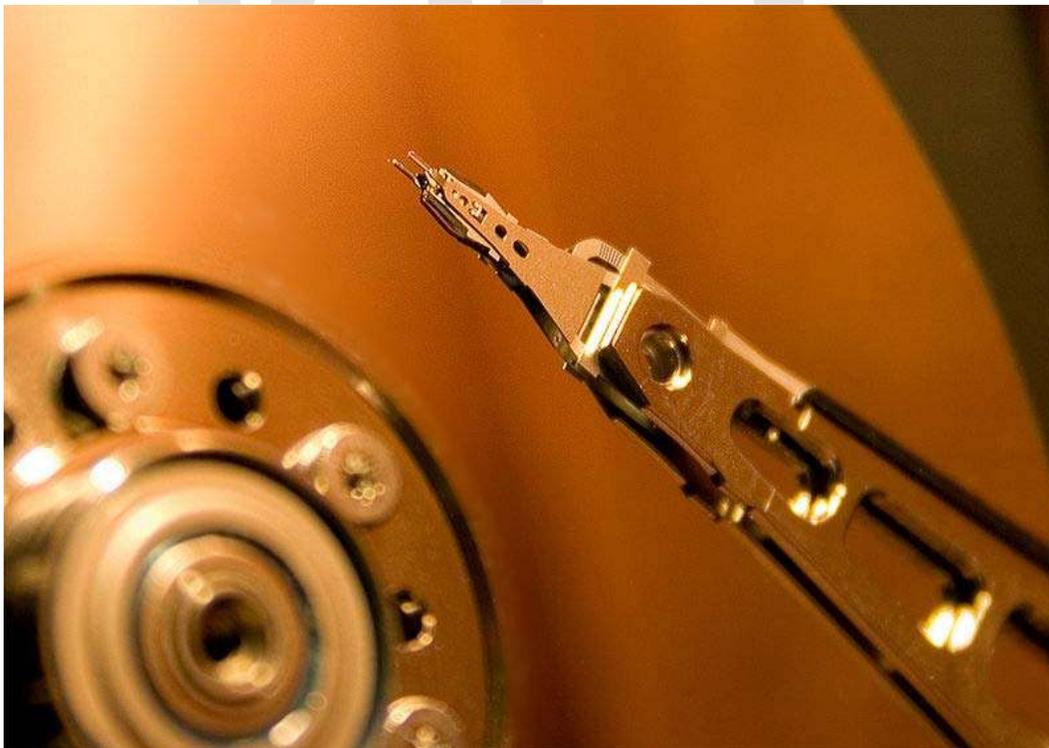
- **Word serial interfaces** connect a hard disk drive to a host bus adapter (today typically integrated into the "south bridge") with one cable for combined data/control. (As for all *bit serial interfaces* above, each drive also has an additional power cable, usually direct to the power supply unit.) The earliest versions of these interfaces typically had a 8 bit parallel data transfer to/from the

drive, but 16-bit versions became much more common, and there are 32 bit versions. Modern variants have serial data transfer. The word nature of data transfer makes the design of a host bus adapter significantly simpler than that of the precursor HDD controller.

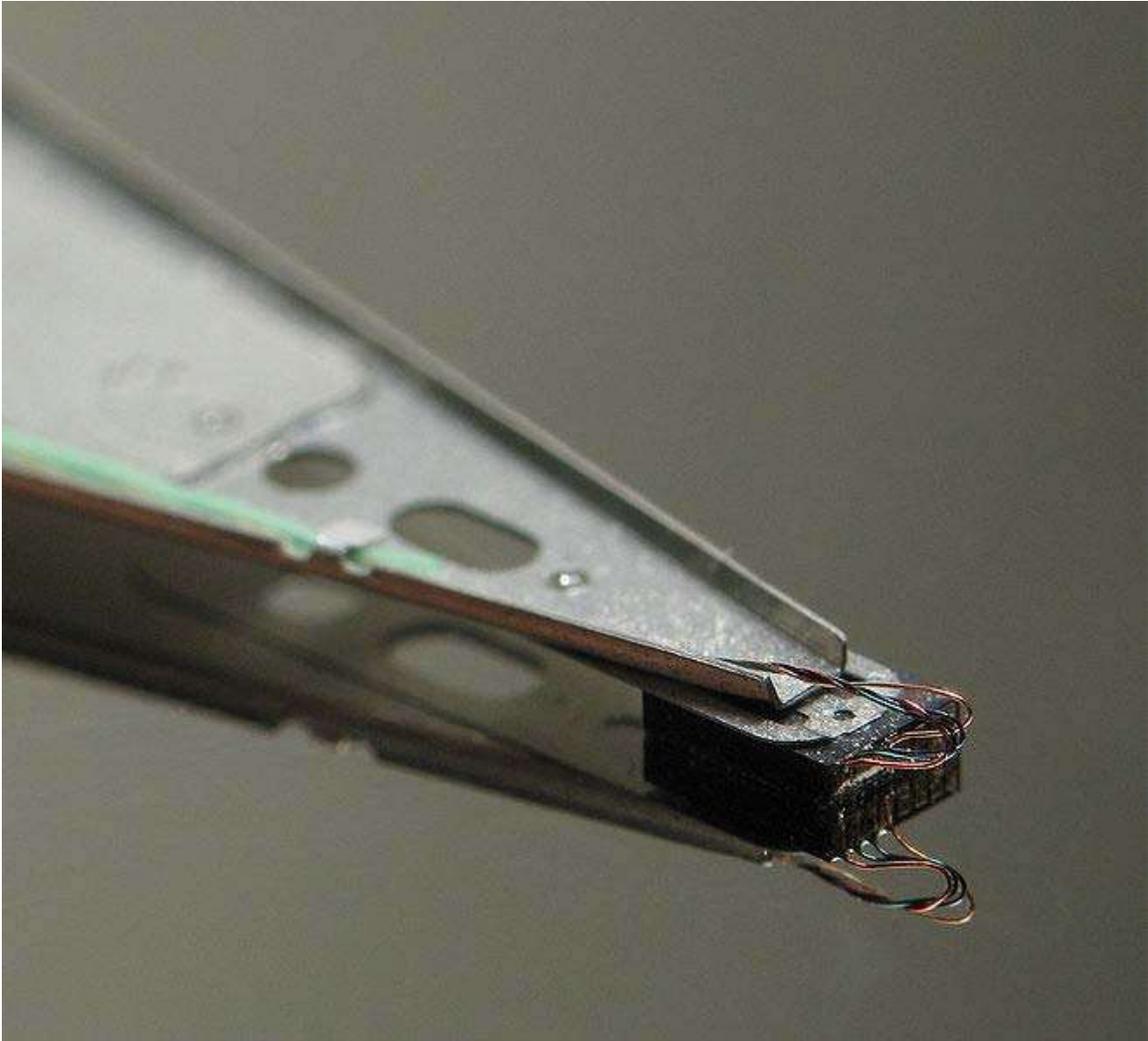
- Integrated Drive Electronics (IDE), later renamed to ATA, with the alias P-ATA ("parallel ATA") retroactively added upon introduction of the new variant Serial ATA. The original name reflected the innovative integration of HDD controller with HDD itself, which was not found in earlier disks. Moving the HDD controller from the interface card to the disk drive helped to standardize interfaces, and to reduce the cost and complexity. The 40-pin IDE/ATA connection transfers 16 bits of data at a time on the data cable. The data cable was originally 40-conductor, but later higher speed requirements for data transfer to and from the hard drive led to an "ultra DMA" mode, known as UDMA. Progressively swifter versions of this standard ultimately added the requirement for a 80-conductor variant of the same cable, where half of the conductors provides grounding necessary for enhanced high-speed signal quality by reducing cross talk. The interface for 80-conductor only has 39 pins, the missing pin acting as a key to prevent incorrect insertion of the connector to an incompatible socket, a common cause of disk and controller damage.
- EIDE was an unofficial update (by Western Digital) to the original IDE standard, with the key improvement being the use of direct memory access (DMA) to transfer data between the disk and the computer without the involvement of the CPU, an improvement later adopted by the official ATA standards. By directly transferring data between memory and disk, DMA eliminates the need for the CPU to copy byte per byte, therefore allowing it to process other tasks while the data transfer occurs.
- Small Computer System Interface (SCSI), originally named SASI for Shugart Associates System Interface, was an early competitor of ESDI. SCSI disks were standard on servers, workstations, Commodore Amiga, and Apple Macintosh computers through the mid-1990s, by which time most models had been transitioned to IDE (and later, SATA) family disks. Only in 2005 did the capacity of SCSI disks fall behind IDE disk technology, though the highest-performance disks are still available in SCSI and Fibre Channel only. The range limitations of the data cable allows for external SCSI devices. Originally SCSI data cables used single ended (common mode) data transmission, but server class SCSI could use differential transmission, either low voltage differential (LVD) or high voltage differential (HVD). ("Low" and "High" voltages for differential SCSI are relative to SCSI standards and do not meet the meaning of low voltage and high voltage as used in general electrical engineering contexts, as apply e.g. to statutory electrical codes; both LVD and HVD use low voltage signals (3.3 V and 5 V respectively) in general terminology.)

Acronym or abbreviation	Meaning	Description
SASI	Shugart Associates System Interface	Historical predecessor to SCSI.
SCSI	Small Computer System Interface	Bus oriented that handles concurrent operations.
SAS	Serial Attached SCSI	Improvement of SCSI, uses serial communication instead of parallel.
ST-506	Seagate Technology	Historical Seagate interface.
ST-412	Seagate Technology	Historical Seagate interface (minor improvement over ST-506).
ESDI	Enhanced Small Disk Interface	Historical; backwards compatible with ST-412/506, but faster and more integrated.
ATA (PATA)	Advanced Technology Attachment	Successor to ST-412/506/ESDI by integrating the disk controller completely onto the device. Incapable of concurrent operations.
SATA	Serial ATA	Modification of ATA, uses serial communication instead of parallel.

Integrity



An IBM HDD head resting on a disk platter. Since the drive is not in operation, the head is simply pressed against the disk by the suspension.



Close-up of a hard disk head resting on a disk platter. A reflection of the head and its suspension is visible on the mirror-like disk.

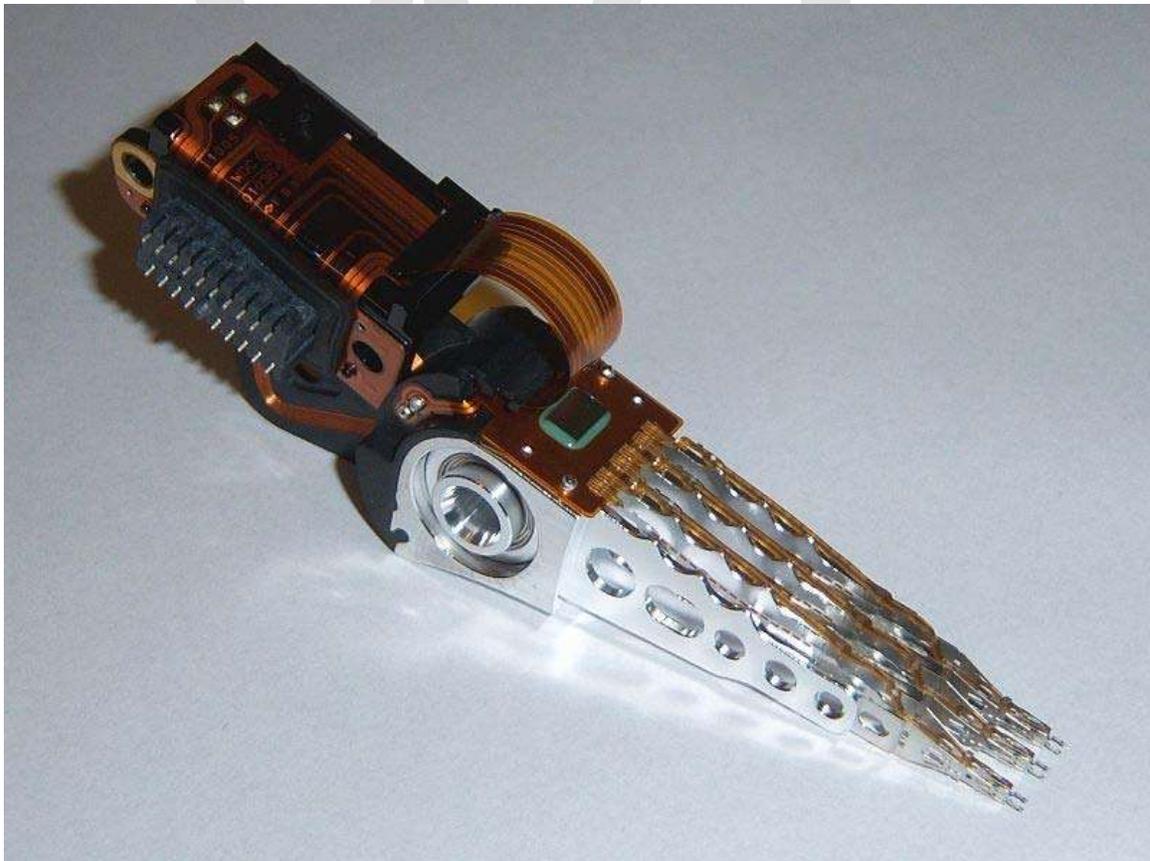
Due to the extremely close spacing between the heads and the disk surface, hard disk drives are vulnerable to being damaged by a head crash—a failure of the disk in which the head scrapes across the platter surface, often grinding away the thin magnetic film and causing data loss. Head crashes can be caused by electronic failure, a sudden power failure, physical shock, contamination of the drive's internal enclosure, wear and tear, corrosion, or poorly manufactured platters and heads.

The HDD's spindle system relies on air pressure inside the disk enclosure to support the heads at their proper *flying height* while the disk rotates. Hard disk drives require a certain range of air pressures in order to operate properly. The connection to the external environment and pressure occurs through a small hole in the enclosure (about 0.5 mm in breadth), usually with a filter on the inside (the *breather filter*). If the air pressure is too low, then there is not enough lift for the flying head, so the head gets too close to the disk, and there is a risk of head crashes and data loss. Specially manufactured sealed and

pressurized disks are needed for reliable high-altitude operation, above about 3,000 m (10,000 feet). Modern disks include temperature sensors and adjust their operation to the operating environment. Breather holes can be seen on all disk drives—they usually have a sticker next to them, warning the user not to cover the holes. The air inside the operating drive is constantly moving too, being swept in motion by friction with the spinning platters. This air passes through an internal recirculation (or "recirc") filter to remove any leftover contaminants from manufacture, any particles or chemicals that may have somehow entered the enclosure, and any particles or outgassing generated internally in normal operation. Very high humidity for extended periods can corrode the heads and platters.

For giant magnetoresistive (GMR) heads in particular, a minor head crash from contamination (that does not remove the magnetic surface of the disk) still results in the head temporarily overheating, due to friction with the disk surface, and can render the data unreadable for a short period until the head temperature stabilizes (so called "thermal asperity", a problem which can partially be dealt with by proper electronic filtering of the read signal).

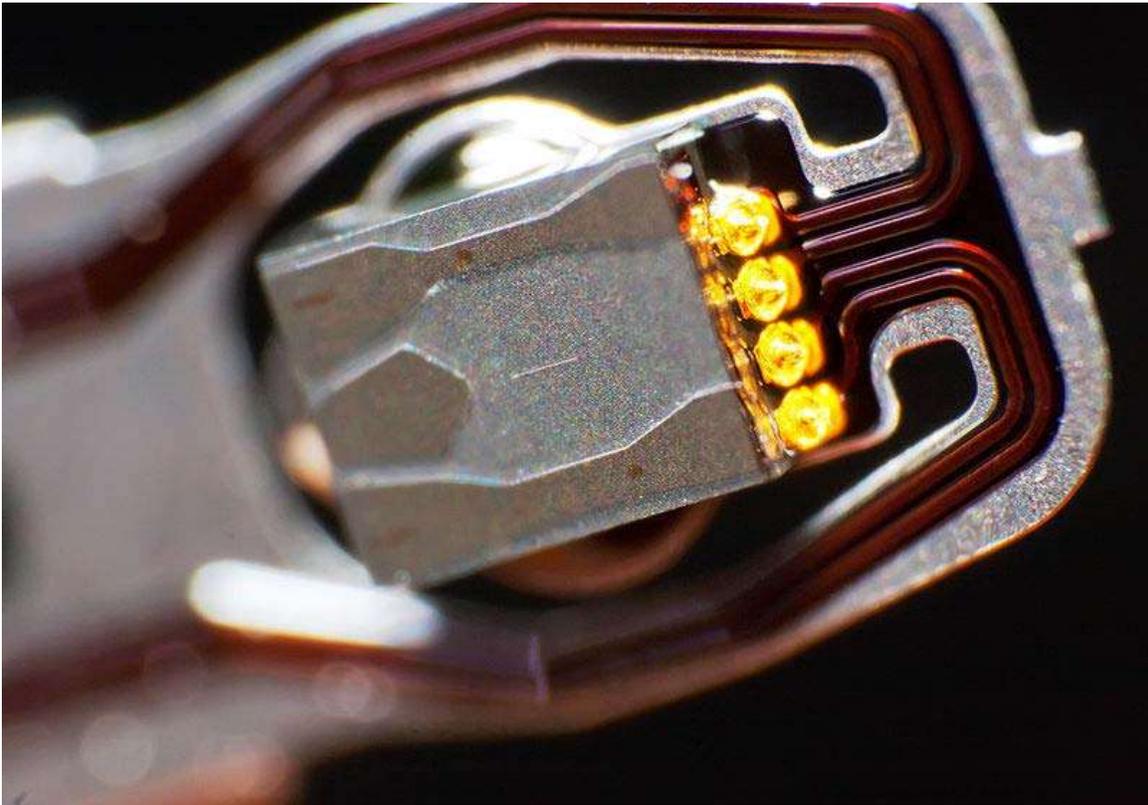
Actuation of moving arm



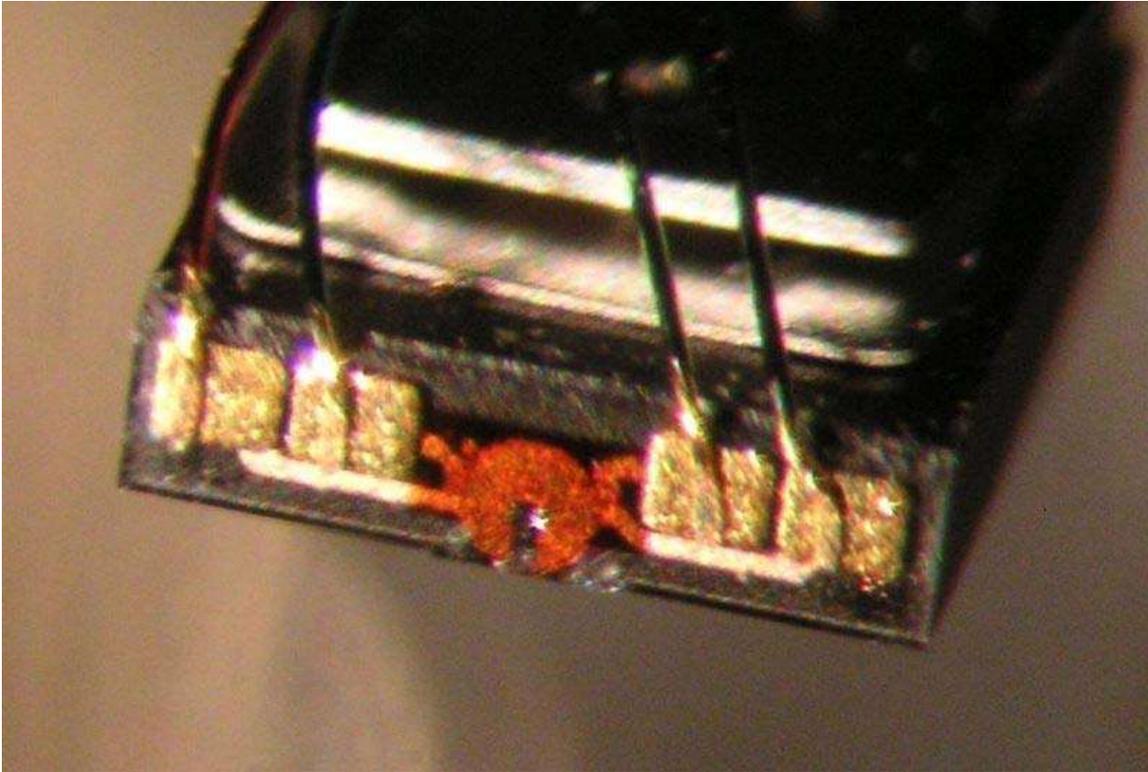
Head stack with actuator coil on the left side (partly hidden by the controller interface) and read/write heads on the right side

The hard drive's electronics control the movement of the actuator and the rotation of the disk, and perform reads and writes on demand from the disk controller. Feedback of the drive electronics is accomplished by means of special segments of the disk dedicated to servo feedback. These are either complete concentric circles (in the case of dedicated servo technology), or segments interspersed with real data (in the case of embedded servo technology). The servo feedback optimizes the signal to noise ratio of the GMR sensors by adjusting the voice-coil of the actuated arm. The spinning of the disk also uses a servo motor. Modern disk firmware is capable of scheduling reads and writes efficiently on the platter surfaces and remapping sectors of the media which have failed.

Landing zones and load/unload technology



A read/write head from a circa-1998 Fujitsu 3.5-inch hard disk. The area pictured is approximately 2.0 mm x 3.0mm.



Microphotograph of an older generation hard disk head and slider (1990s). The size of the front face (which is the "trailing face" of the slider) is about $0.3 \text{ mm} \times 1.0 \text{ mm}$. It is the location of the actual head (magnetic sensors). The non-visible bottom face of the slider is about $1.0 \text{ mm} \times 1.25 \text{ mm}$ (so-called "nano" size) and faces the platter. It contains the lithographically micro-machined air bearing surface (ABS) that allows the slider to fly in a highly controlled fashion. One functional part of the head is the round, orange structure visible in the middle—the lithographically defined copper coil of the write transducer. Also note the electric connections by wires bonded to gold-plated pads.

Modern HDDs prevent power interruptions or other malfunctions from landing its heads in the data zone by **parking** the heads either in a **landing zone** or by unloading (i.e., **load/unload**) the heads. Some early PC HDDs did not park the heads automatically and they would land on data. In some other early units the user manually parked the heads by running a program to park the HDD's heads.

A **landing zone** is an area of the platter usually near its inner diameter (ID), where no data are stored. This area is called the Contact Start/Stop (CSS) zone. Disks are designed such that either a spring or, more recently, rotational inertia in the platters is used to park the heads in the case of unexpected power loss. In this case, the spindle motor temporarily acts as a generator, providing power to the actuator.

Spring tension from the head mounting constantly pushes the heads towards the platter. While the disk is spinning, the heads are supported by an air bearing and experience no physical contact or wear. In CSS drives the sliders carrying the head sensors (often also

just called *heads*) are designed to survive a number of landings and takeoffs from the media surface, though wear and tear on these microscopic components eventually takes its toll. Most manufacturers design the sliders to survive 50,000 contact cycles before the chance of damage on startup rises above 50%. However, the decay rate is not linear: when a disk is younger and has had fewer start-stop cycles, it has a better chance of surviving the next startup than an older, higher-mileage disk (as the head literally drags along the disk's surface until the air bearing is established). For example, the Seagate Barracuda 7200.10 series of desktop hard disks are rated to 50,000 start-stop cycles, in other words no failures attributed to the head-platter interface were seen before at least 50,000 start-stop cycles during testing.

Around 1995 IBM pioneered a technology where a landing zone on the disk is made by a precision laser process (*Laser Zone Texture* = LZT) producing an array of smooth nanometer-scale "bumps" in a landing zone, thus vastly improving stiction and wear performance. This technology is still largely in use today (2008), predominantly in desktop and enterprise (3.5 inch) drives. In general, CSS technology can be prone to increased stiction (the tendency for the heads to stick to the platter surface), e.g. as a consequence of increased humidity. Excessive stiction can cause physical damage to the platter and slider or spindle motor.

Load/Unload technology relies on the heads being lifted off the platters into a safe location, thus eliminating the risks of wear and stiction altogether. The first HDD RAMAC and most early disk drives used complex mechanisms to load and unload the heads. Modern HDDs use ramp loading, first introduced by Memorex in 1967, to load/unload onto plastic "ramps" near the outer disk edge.

All HDDs today still use one of these two technologies listed above. Each has a list of advantages and drawbacks in terms of loss of storage area on the disk, relative difficulty of mechanical tolerance control, non-operating shock robustness, cost of implementation, etc.

Addressing shock robustness, IBM also created a technology for their ThinkPad line of laptop computers called the Active Protection System. When a sudden, sharp movement is detected by the built-in accelerometer in the Thinkpad, internal hard disk heads automatically unload themselves to reduce the risk of any potential data loss or scratch defects. Apple later also utilized this technology in their PowerBook, iBook, MacBook Pro, and MacBook line, known as the Sudden Motion Sensor. Sony, HP with their HP 3D DriveGuard and Toshiba have released similar technology in their notebook computers.

This accelerometer-based shock sensor has also been used for building cheap earthquake sensor networks.

Disk failures and their metrics

Most major hard disk and motherboard vendors now support S.M.A.R.T. (Self-Monitoring, Analysis, and Reporting Technology), which measures drive characteristics

such as operating temperature, spin-up time, data error rates, etc. Certain trends and sudden changes in these parameters are thought to be associated with increased likelihood of drive failure and data loss.

However, not all failures are predictable. Normal use eventually can lead to a breakdown in the inherently fragile device, which makes it essential for the user to periodically back up the data onto a separate storage device. Failure to do so can lead to the loss of data. While it may sometimes be possible to recover lost information, it is normally an extremely costly procedure, and it is not possible to guarantee success. A 2007 study published by Google suggested very little correlation between failure rates and either high temperature or activity level; however, the correlation between manufacturer/model and failure rate was relatively strong. Statistics in this matter is kept highly secret by most entities. Google did not publish the manufacturer's names along with their respective failure rates, though they have since revealed that they use Hitachi Deskstar drives in some of their servers. While several S.M.A.R.T. parameters have an impact on failure probability, a large fraction of failed drives do not produce predictive S.M.A.R.T. parameters. S.M.A.R.T. parameters alone may not be useful for predicting individual drive failures.

A common misconception is that a colder hard drive will last longer than a hotter hard drive. The Google study seems to imply the reverse—"lower temperatures are associated with higher failure rates". Hard drives with S.M.A.R.T.-reported average temperatures below 27 °C (80.6 °F) had higher failure rates than hard drives with the highest reported average temperature of 50 °C (122 °F), failure rates at least twice as high as the optimum S.M.A.R.T.-reported temperature range of 36 °C (96.8 °F) to 47 °C (116.6 °F).

SCSI, SAS, and FC drives are typically more expensive and are traditionally used in servers and disk arrays, whereas inexpensive ATA and SATA drives evolved in the home computer market and were perceived to be less reliable. This distinction is now becoming blurred.

The mean time between failures (MTBF) of SATA drives is usually about 600,000 hours (some drives such as Western Digital Raptor have rated 1.4 million hours MTBF), while SCSI drives are rated for upwards of 1.5 million hours. However, independent research indicates that MTBF is not a reliable estimate of a drive's longevity. MTBF is conducted in laboratory environments in test chambers and is an important metric to determine the quality of a disk drive before it enters high volume production. Once the drive product is in production, the more valid metric is annualized failure rate (AFR). AFR is the percentage of real-world drive failures after shipping.

SAS drives are comparable to SCSI drives, with high MTBF and high reliability.

Enterprise S-ATA drives designed and produced for enterprise markets, unlike standard S-ATA drives, have reliability comparable to other enterprise class drives.

Typically enterprise drives (all enterprise drives, including SCSI, SAS, enterprise SATA, and FC) experience between 0.70%–0.78% annual failure rates from the total installed drives.

Eventually all mechanical hard disk drives fail, so to mitigate loss of data, some form of redundancy is needed, such as RAID or a regular backup system.

External removable drives

External removable hard disk drives connect to the computer using a USB cable or other means. External drives are used for:

- Backup of files and information
- Data recovery
- Disk cloning
- Running virtual machines
- Scratch disk for video editing applications and video recording.



A 6 GB Seagate Pocket hard drive with USB cable extended next to a 2 GB CompactFlash card.

Larger models often include full-sized 3.5" PATA or SATA desktop hard drives. Features such as biometric security or multiple interfaces generally increase cost.

Market segments

- As of July 2010, the highest capacity consumer HDDs are 3 TB.
- **"Desktop HDDs"** typically store between 120 GB and 2 TB and rotate at 5,400 to 10,000 rpm, and have a media transfer rate of 0.5 Gbit/s or higher. (1 GB = 10⁹ bytes; 1 Gbit/s = 10⁹ bit/s)
- **Enterprise HDDs** are typically used with multiple-user computers running enterprise software. Examples are
 - transaction processing databases;
 - internet infrastructure (email, webserver, e-commerce);
 - scientific computing software;
 - nearline storage management software.

The fastest enterprise HDDs spin at 10,000 or 15,000 rpm, and can achieve sequential media transfer speeds above 1.6 Gbit/s. and a sustained transfer rate up to 1 Gbit/s. Drives running at 10,000 or 15,000 rpm use smaller platters to mitigate increased power requirements (as they have less air drag) and therefore generally have lower capacity than the highest capacity desktop drives. Enterprise drives commonly operate continuously ("24/7") in demanding environments while delivering the highest possible performance without sacrificing reliability. Maximum capacity is not the primary goal, and as a result the drives are often offered in capacities that are relatively low in relation to their cost.

- **Mobile HDDs** or laptop HDDs, smaller than their desktop and enterprise counterparts, tend to be slower and have lower capacity. A typical mobile HDD spins at either 4200 rpm, 5200 rpm, 5400 rpm, or 7200 rpm, with 5400 rpm being the most prominent. 7200 rpm drives tend to be more expensive and have smaller capacities, while 4200 rpm models usually have very high storage capacities. Because of smaller platter(s), mobile HDDs generally have lower capacity than their greater desktop counterparts.

The exponential increases in disk space and data access speeds of HDDs have enabled the commercial viability of consumer products that require large storage capacities, such as digital video recorders and digital audio players. In addition, the availability of vast amounts of cheap storage has made viable a variety of web-based services with extraordinary capacity requirements, such as free-of-charge web search, web archiving, and video sharing (Google, Internet Archive, YouTube, etc.).

Sales

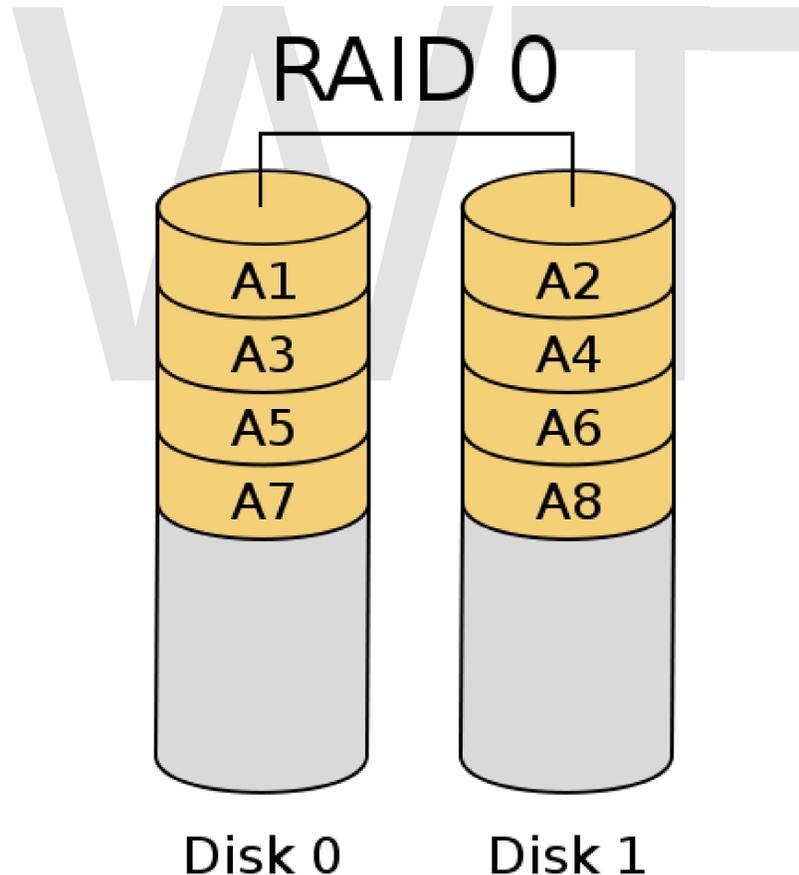
Worldwide revenue from shipments of HDDs is expected to reach \$27.7 billion in 2010, up 18.4% from \$23.4 billion in 2009 corresponding to a 2010 unit shipment forecast of 674.6 million compared to 549.5 million units in 2009.

Icons

Hard drives are traditionally symbolized as either a stylized stack of platters (in orthographic projection) or, more abstractly, as a cylinder. This is particularly found in schematic diagrams or on indicator lights, as on laptops, to indicate hard drive access. In most modern operating systems, hard drives are instead represented by an illustration or photograph of a hard drive enclosure. These are illustrated below.



Today, hard drives are symbolized by a picture of the enclosure.



Schematically, hard drives may be represented by cylinders or stacks of platters, as in this RAID diagram.



The cylinder schematic derives from hard drives internally being a stack of platters, as in these 1970s vintage disk pack (cover removed).

Manufacturers

More than 200 companies have manufactured hard disk drives. Today most drives are made by Seagate, Western Digital, Hitachi, Samsung, and Toshiba (though Toshiba does not manufacture 3.5 inch drives).