# Hard Disk Connectors and Jumpers

- Several different connectors and jumpers are used to configure the hard disk and connect it to the rest of the system.

- The number and types of connectors on the hard disk depend on the data interface it uses to connect to the system, the manufacturer of the drive, and any special features that the drive may possess.

- Instructions for setting common jumpers are usually printed right on the drive; full instructions for all jumpers will be in the product's manual, or on the manufacturer's web site.

# Some of the connectors and jumper pins

- Some of the connectors and jumper pins on a 3.5", 36 GB, 10,000 RPM SCSI Cheetah drive.

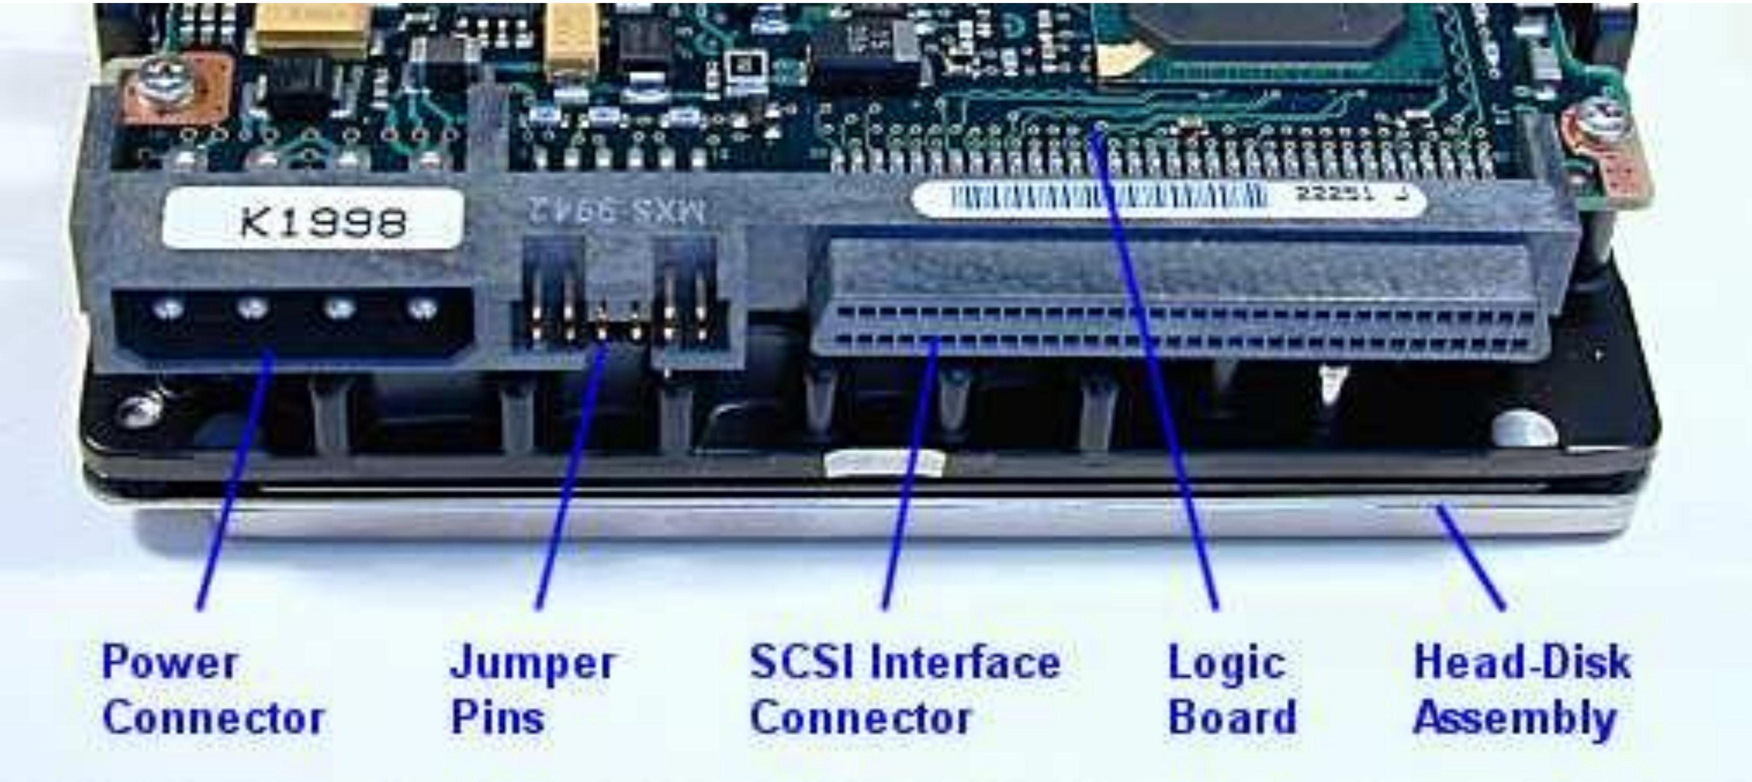

# Power Connector

- Hard disk drives use a standard, 4-pin male connector plug, that takes one of the power connectors coming from the power supply. This keyed, 4-wire plastic connector provides +5 and +12 voltage to the hard disk. See this discussion of drive power plugs in the power supply reference for more information

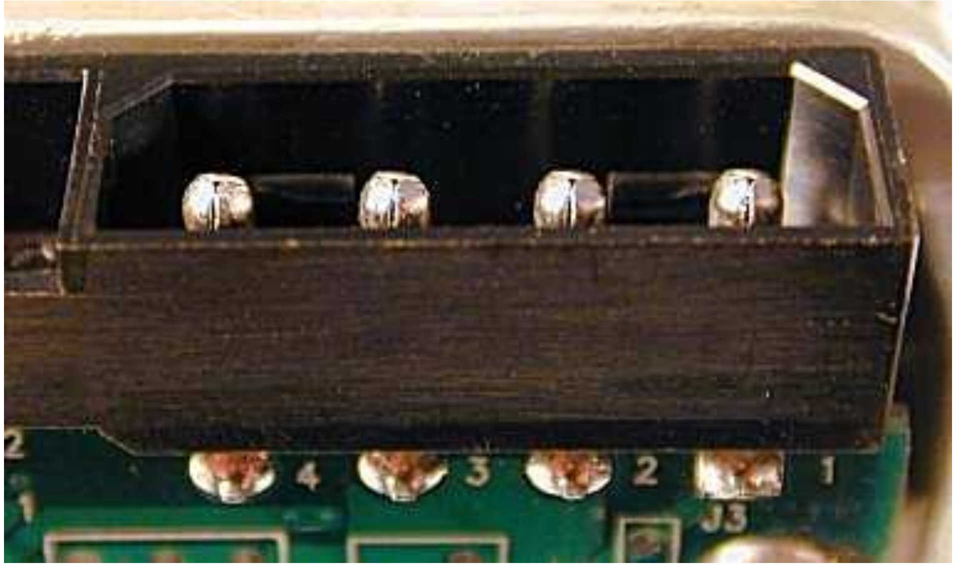


- A standard hard disk power connector. Note the "4, 3, 2, 1" pin markings on the logic board, and the square shape of solder pad for pin #1.

# Data Interface Connector

- Modern hard disk drives use one of two interfaces: IDE (ATA) and its variants, or SCSI. You can tell immediately by looking at the back of the hard disk which interface is being used by the drive:
- **IDE/ATA:** A **40-pin** rectangular connector.
- **SCSI:** A **50-pin, 68-pin, or 80-pin** D-shaped connector (the same shape used for serial and parallel port connectors). A 50-pin connector means the device is narrow SCSI; 68 pins means wide SCSI; 80 pins means wide SCSI using single connector attachment (SCA)

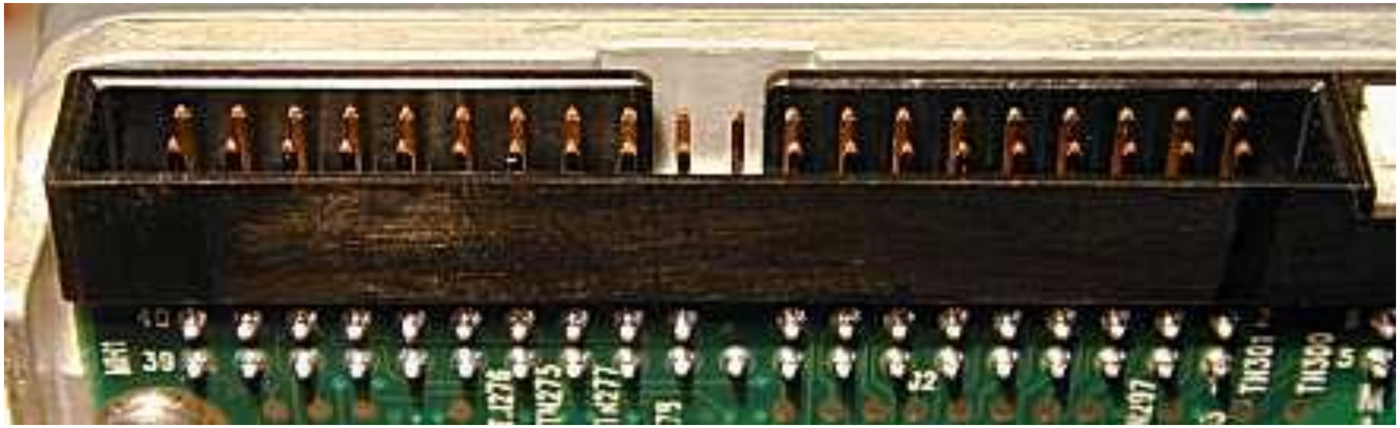


- A standard hard disk IDE/ATA data interface connector.
If you look closely you can see markings for pins #1, #2, #39 and #40.
A 50-pin SCSI interface connector looks identical except that it has 25 columns of pins instead of 20 (they look so much alike that getting the two mixed up is common).

# Data Interface Connector

- The connectors on hard disk drives are generally in the form of a **2xN rectangular grid of pins** (where N is 20, 25, 34 or 40 depending on the interface).

- Older ST-506 (also called MFM, RLL) and ESDI hard disks used two data connectors, one 34 pins and the other 20 pins. These connectors were often not in the form of pins but rather card edge connectors, such as those used by ISA expansion cards. Some SCSI connectors may have different shapes, especially older ones.

- While most current SCSI interface connectors are keyed to prevent incorrect insertion (because they are D-shaped), this is not always the case for other interfaces. For this reason, it is important to make sure that the cable is oriented the correct way before plugging it in. The cable has a red stripe to indicate wire #1 and the hard disk uses markers of one form or another to indicate the matching pin #1.

# IDE/ATA Configuration Jumpers

- IDE/ATA hard disks are fairly standard in terms of jumpers. There are usually only a few and they don't vary greatly from drive to drive. Here are the jumpers you will normally find:

- **Drive Select:** Since there can be two drives (master and slave) on the same IDE channel, a jumpers is normally used to tell each drive if it should function as a master or slave on the IDE channel. For a single drive on a channel, most manufacturers instruct that the drive be jumpered as master, while some manufacturers (notably Western Digital) have a separate setting for a single drive as opposed to a master on a channel with a slave. The terms "master" and "slave" are misleading since the drives really have no operational relationship.

- **Slave Present:** Some drives have an additional jumper that is used to tell a drive configured as master that there is also a slave drive on the ATA channel. This is only required for some older drives that don't support standard master/slave IDE channel signaling.

- **Cable Select:** Some configurations use a special cable to determine which drive is master and which is slave, and when this system is used a cable select jumper is normally enabled.

- **Size Restriction Jumper:** Some larger hard disk drives don't work properly in older PCs that don't have a BIOS program modern enough to recognize them. To get around this, some drives have special jumpers that, when set, will cause them to appear as a smaller size than they really are to the BIOS, for compatibility. For example, some 2.5 GB hard disks have a jumper that will cause them to appear as a 2.1 GB hard disk to a system that won't support anything over 2.1 GB. These are also sometimes called *capacity limitation jumpers* and vary from manufacturer to manufacturer.

# IDE/ATA Configuration Jumpers

- Jumper block for an IDE hard disk.

- The jumpers are labeled:
- "MA" (master)
- "SL" (slave)
- "CS" (cable select).
- Other IDE drives will have slightly different jumper configuration or placement.

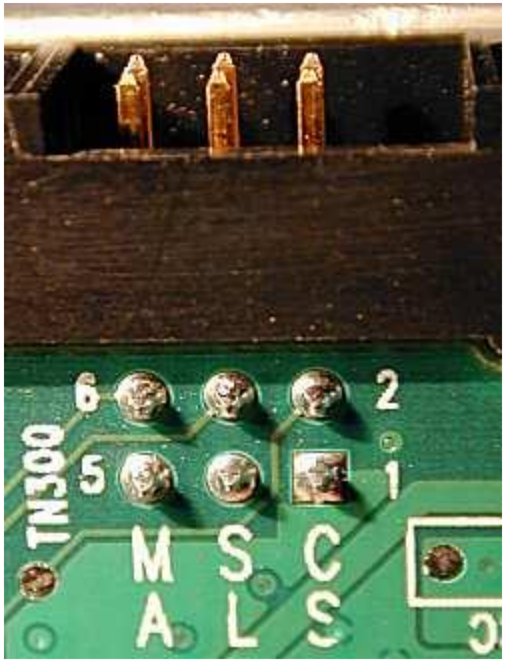

# SCSI Configuration Jumpers

- SCSI hard disks have more sophisticated controllers than their IDE/ATA cousins, and as a result typically have many more jumpers that can be set to control their operation. They also tend to vary much more from manufacturer to manufacturer, and from model to model, in the number and types of jumpers they have. Typically the following are the most common and important SCSI drive jumpers:
- **SCSI Device ID:** Every device on a SCSI bus must be uniquely identified for addressing purposes. Narrows SCSI drives will have a set of three jumpers that can be used to assign the disk an ID number from 0 to 7. Wide SCSI drives will have four jumpers to enable ID numbers from 0 to 15. Some systems don't use jumpers to configure SCSI device IDs.
- **Termination Activate:** The devices on the ends of the SCSI bus must terminate the bus for it to function properly. If the hard disk is at the end of the bus, setting this jumper will cause it to terminate the bus for proper operation. Not all drives support termination.
- **Disable Auto Start:** When present, this jumper will tell the drive not to automatically spin up when the power is applied, but instead wait for a start command over the SCSI bus. This is usually done to prevent excessive startup load on the power supply. Some manufacturers invert the sense of this jumper; they disable startup by default and provide an "Enable Auto Start" jumper.
- **Delay Auto Start:** This jumper tells the drive to start automatically, but wait a predefined number of seconds from when power is applied. It is also used to offset motor startup load on systems with many drives.
- **Stagger Spin:** An "enhanced version" of "Delay Auto Start". When a system with many hard drives has this option set for each unit, the drives stagger their startup time by multiplying a user-defined constant times their SCSI device ID. This ensures no two drives on the same SCSI channel will start up simultaneously.
- **Narrow/Wide:** Some drives have a jumper to control whether they will function in narrow or wide mode.
- **Force SE:** Allows newer Ultra2, Wide Ultra2, Ultra160, Ultra160+ or other LVD SCSI drives to be forced to use single-ended (SE) operation instead of low voltage differential (LVD).
- **Disable Parity:** Turns off parity checking on the SCSI bus, for compatibility with host adapters that don't support the featurer.

# SCSI Configuration Jumpers

- Option header block signals and functions for the Quantum Atlas 10K SCSI drive

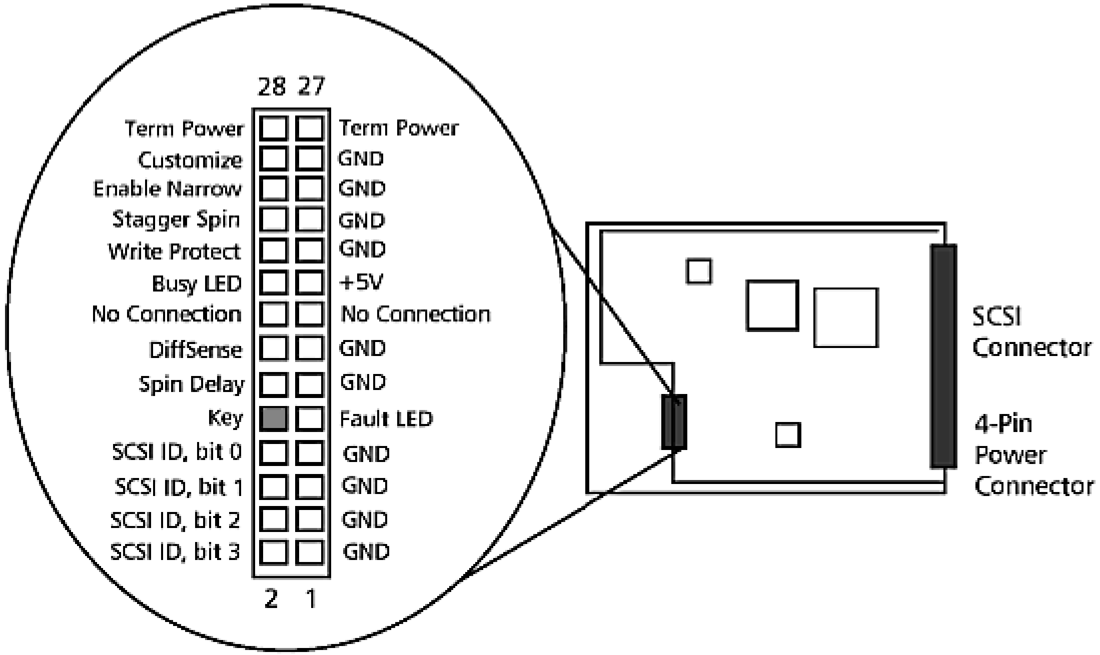


- This is list is not intended to be exhaustive; many SCSI drives have additional special features that are enabled through more jumpers. Some drives have replaced some of their jumpers with software commands sent over the SCSI interface. SCSI jumpers are often clustered together into what is called an *option block*.

# LED Connector

- Hard disks use a **light-emitting diode or LED to indicate drive activity** (if you are unfamiliar with LEDs, see the page on basic electrical components for some background information). The hard disk activity LED is a very useful indicator that generally tells the PC user at a glance when the system is active. The first PC hard disks shipped with a faceplate (or bezel) on the front. The hard disk was mounted into an external hard drive bay (in place of a floppy disk drive) and its integral LED was visible from the front of the PC, because the drive's front was actually protruding from the case, much as floppy disk drives still do.

- It was quickly realized that having the disks mounted internally to the case made more sense than using external drive bays, but the LED was still desirable. So a remote LED was mounted to the case and a wire run to a two-pin connector on the hard disk itself. This system worked fine when there was just one hard disk, but became a problem in systems that had two or three hard disks. Eventually, the case LED was made to connect to the hard disk *controller* instead, to show activity on any of the hard disks that were managed by the controller.

# Hard Disk Logic Board

- **All modern hard disks** are made with an **intelligent circuit board** integrated into the hard disk unit. **Early hard disks were "dumb**", meaning that virtually all of the control logic for controlling the hard disk itself was placed into the controller plugged into the PC; there were very little smarts on the drive itself, which had to be told specifically how to perform every action. This design meant that it was necessary for controllers to be very *generalized*; they could not be customized to the particulars of a given disk drive since they had to be able to work with any type. Older drives were similar enough, and sufficiently undemanding in terms of performance that this arrangement was acceptable. As newer drives were introduced with more features and faster speed, this approach became quite impractical, and once electronics miniaturization progressed far enough, it made sense to move most of the control functions to the drive itself.

- The most common interface for PC hard disks is called *IDE*, which in fact stands for *Integrated Drive Electronics*. This name is something of a misnomer today. When it was introduced, IDE was distinguished from the other interfaces of the day by having the integrated electronics on the drive, instead of on the controller card plugged into the system bus like older interfaces. However, the term really refers to where the control logic is and not the interface itself, and since all hard disks today use integrated electronics the name doesn't mean anything any more, despite the fact that everyone continues to use it. The other popular PC hard disk interface today, SCSI, also uses drives that have integrated controllers. The more correct name for the IDE interface is *AT Attachment* or *ATA*;

# Hard Disk Logic Board

- The logic board of a Cheetah 10,000 RPM 36 GB hard disk drive. The main interface and power connectors are on the right-hand side; auxiliary connectors on the bottom and left side. The bottom of the spindle motor protrudes through a round hole made for it in the circuit board.

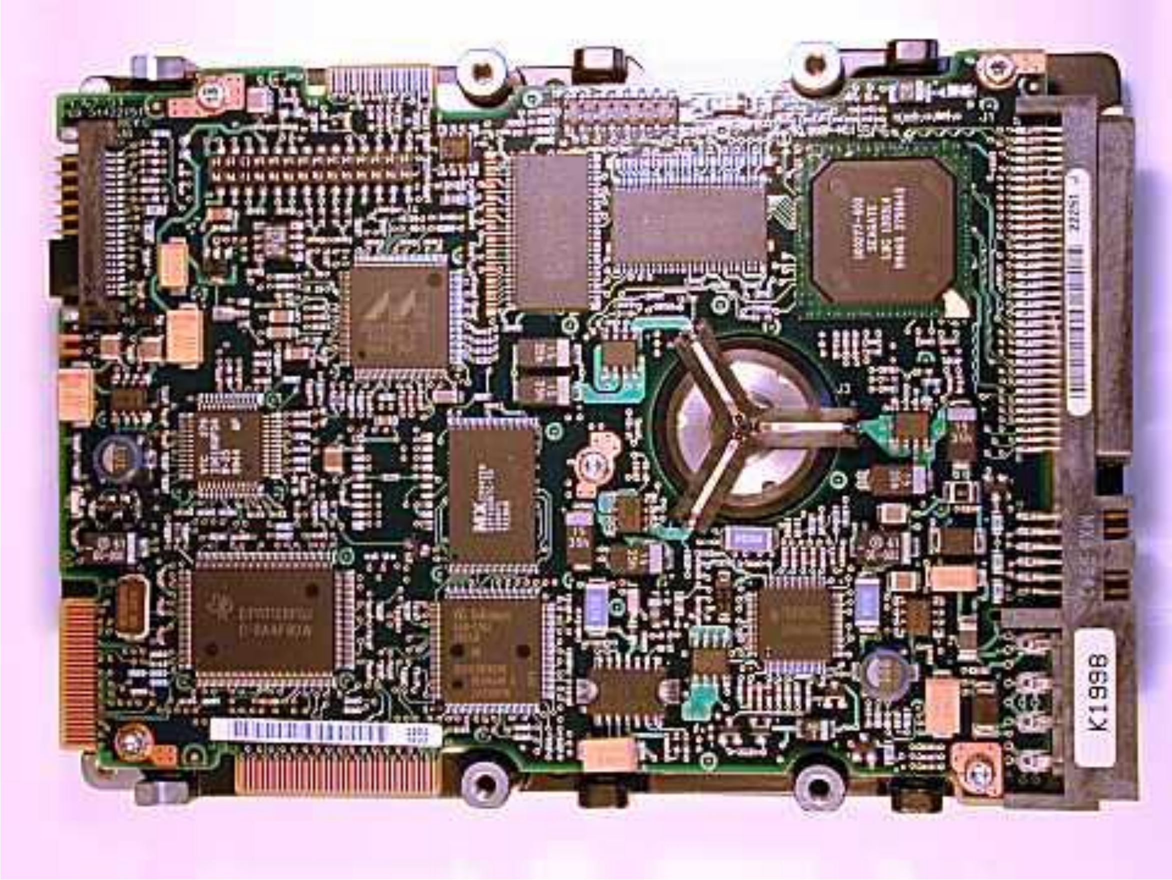

# Hard Disk Logic Board

- Today's hard disks contain logic boards that are in most ways more sophisticated than an entire early PC!
- In fact, most of them **contain more memory** and faster internal processors than an entire PC of even the mid-1980s.
- The **logic board performs several important functions**, and as hard disks become faster and more sophisticated, more functions are added to the logic board. This means the logic circuits need to be more powerful, to handle changes like:
  - geometry translation
  - advanced reliability features
  - more complicated head technologies
  - faster interfaces
  - higher bandwidth data streaming from the disk itself.

# Control Circuitry

- The drive's internal logic board contains a microprocessor and internal memory, and other structures and circuits that control what happens inside the drive. In many ways, this is **like a small embedded PC** within the hard disk itself.

- The **control circuitry** of the drive performs the **following functions** :
  - Controlling the spindle motor, including making sure the spindle runs at the correct speed.
  - Controlling the actuator's movement to various tracks.
  - Managing all read/write operations.
  - Implementing power management features.
  - Handling geometry translation.
  - Managing the internal cache and optimization features such as pre-fetch.
  - Coordinating and integrating the other functions mentioned in this section, such as the flow of information over the hard disk interface, optimizing multiple requests, converting data to and from the form the read/write heads require it, etc.
  - Implementing all advanced performance and reliability features.

# Interface Hardware

- While the drive electronics on a modern disk are moved to the disk itself, there is still a controller card or integrated controller on the motherboard, that the disk must talk to. The difference between older drives and newer ones is that older controllers actually ran the internals of the disk itself, while the newer ones just run the data interface between the disk and the rest of the system. Every hard disk's logic contains interface hardware to manage this flow of information between itself and the controller it is talking to.
- An integrated IDE/ATA controller chip from a Fujitsu hard disk logic board.


- The interface hardware itself ranges from relatively simple (slower, older IDE/ATA interfaces) to relatively complex (newer, faster IDE/ATA and SCSI interfaces). In particular, SCSI interface hard disks have considerable "smarts" on them to handle the increased sophistication and the wide variety of commands available on that interface.

# Firmware

- Since modern hard disks have internal **microprocessors**, they also have **internal "software" that runs them**. These routines are what run the control logic and make the drive work.

- Of course this isn't really software in the conventional sense, because these instructions are embedded into read-only memory.

- This code is analogous to the system BIOS: low-level, hardware-based control routines, embedded in ROM. It is usually called *firmware*, with the word "firm" intending to connote something in between "hard" and "soft".

- The functions that run the logic board's circuitry could be implemented strictly with hardware devices, as was done with early drives. However, this would be expensive and inflexible for today's sophisticated controllers, since it would make it difficult to update or adapt the logic to match changes in hard disks or the devices they interface with.

# Multiple Command Control and Reordering

- The *IDE/ATA interface* used on most PCs is generally **limited to a single** transaction outstanding on the interface at any given time. This means the hard disk can not only do just one thing at a time, it can't even keep track of what it needs to do next. It must be fed commands one at a time by the controller. In contrast, newer drives using the SCSI interface generally include the **ability to handle multiple requests**, up to a certain number. This advanced feature, sometimes called *command queuing and reordering* or *multiple command queuing*, is very useful for servers and other systems being used by multiple people (while its absence is generally not a problem for most single-user PCs equipped with IDE/ATA hardware.)

- If the hard drive's logic circuitry receive multiple commands to read or write from the disk, it must process them and figure out where on the disk the data is for each request. Some requests may be filled from the internal cache; these would generally be filled immediately. For the remainder, the controller must decide in which order to perform them. Since random reads or writes on even the fastest hard disks take thousands of times longer than computing operations, this determination is very important to overall performance. There are probably dozens of different specific algorithms that could be used to decide which commands get fulfilled first. However, they generally fall into these three categories:

# Multiple Command Control and Reordering

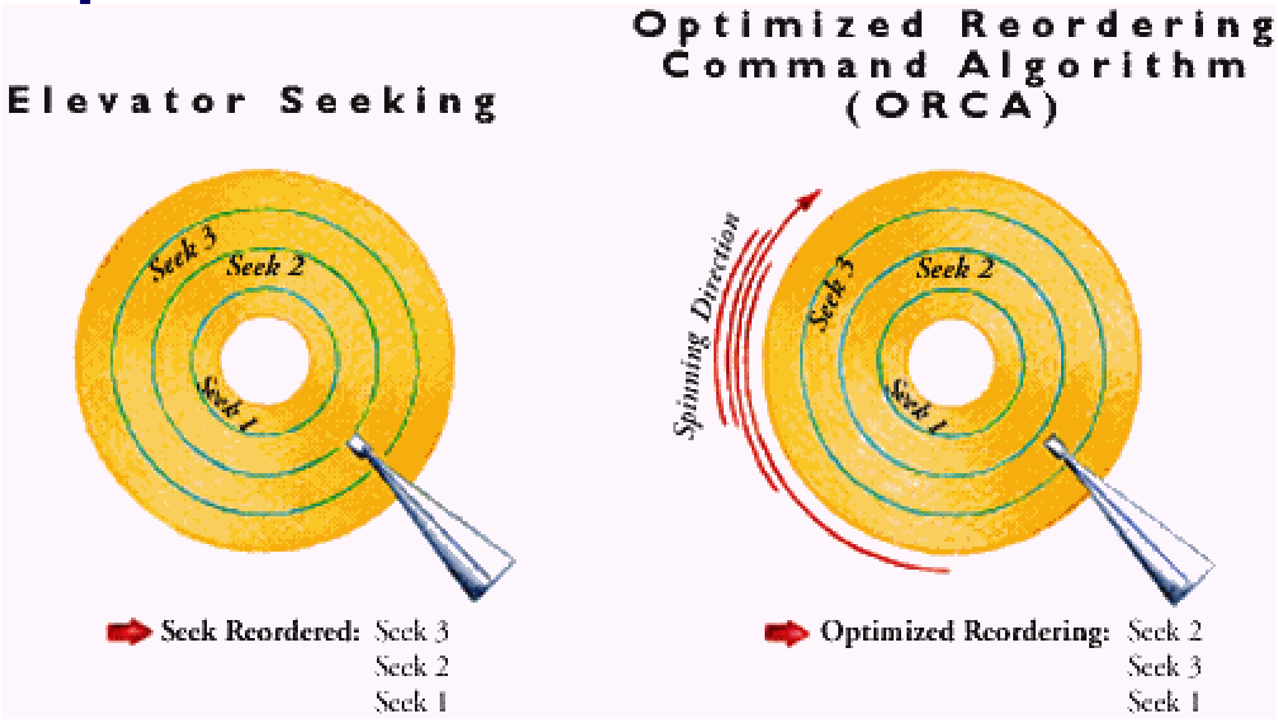


- Quantum's version of command optimization is called "ORCA", as you can see above. The diagram shows graphically what the idea is with this technology.
  Note when comparing the two figures above, that on the one at right, "Seek 3" has been "rotated" counter-clockwise a bit. That's significant: if you applied "ORCA" to the figure at left, with "Seek 2" and "Seek 3" so close to each other, it would probably not be possible to do them in the order "2, 3, 1" because there wouldn't be enough time to change tracks before "Seek 3" rotated past the head. If you tried, you'd end up doing worse than if you just did "2, 1, 3". This is all part of what the algorithm has to figure out as part of its job.

# Hard Disk Cache and Cache Circuitry

- **All modern hard disks contain an integrated *cache***, also often called a *buffer*. The purpose of this cache is not dissimilar to other caches used in the PC, even though it is not normally thought of as part of the regular PC cache hierarchy. The function of cache is to act as a buffer between a relatively fast device and a relatively slow one. For hard disks, the cache is used to hold the results of recent reads from the disk, and also to "pre-fetch" information that is likely to be requested in the near future, for example, the sector or sectors immediately after the one just requested.

- The use of cache improves performance of any hard disk, by reducing the number of physical accesses to the disk on repeated reads and allowing data to stream from the disk uninterrupted when the bus is busy.

- Most modern hard disks have between 512 KiB and 2 MiB of internal cache memory, although some high-performance SCSI drives have as much as 16 MiB, more than many whole PCs have!

# Cache Circuitry and Operation

- The reason that the hard disk's cache is important is due to the sheer difference in the speeds of the hard disk and the hard disk interface.

- Finding a piece of data on the hard disk involves random positioning, and incurs a penalty of *milliseconds* as the hard disk actuator is moved and the disk rotates around on the spindle. In today's PCs, a millisecond is an *eternity*. On a typical IDE/ATA hard disk, transferring a 4,096-byte block of **data from the disk's internal cache** is over **100 times faster than actually finding it and reading it from the platters.**

- *That* is why hard disks have internal buffers. :^) If a seek isn't required (say, for reading a long string of consecutive sectors from the disk) the difference in speed isn't nearly as great, but the buffer is still much faster.

# Cache Circuitry and Operation

- General concepts behind the operation of an internal hard disk cache

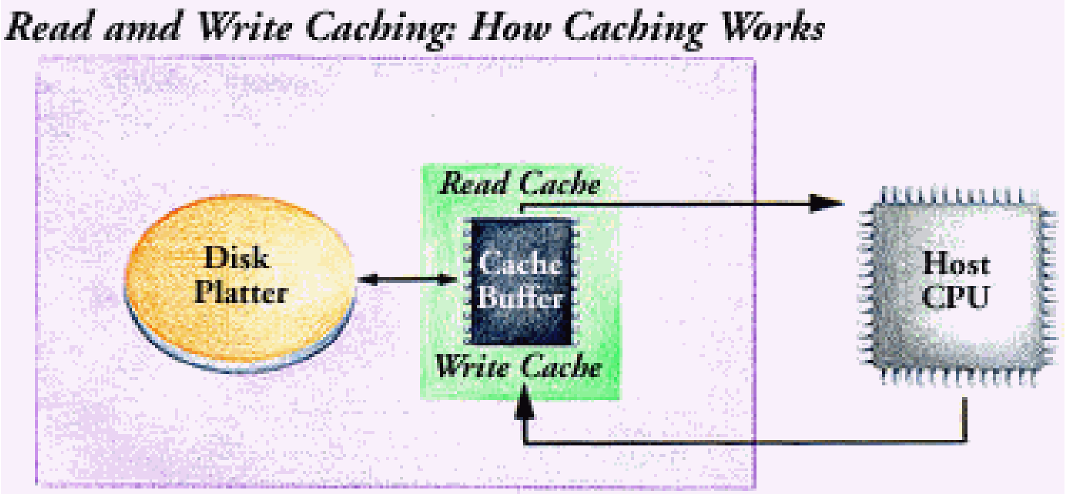


- The basic principle behind the operation of a simple cache is straightforward. Reading data from the hard disk is generally done in blocks of various sizes, not just one 512-byte sector at a time. The cache is broken into "segments", or pieces, each of which can contain one block of data. When a request is made for data from the hard disk, the cache circuitry is first queried to see if the data is present in any of the segments of the cache. If it is present, it is supplied to the logic board without access to the hard disk's platters being necessary. If the data is not in the cache, it is read from the hard disk, supplied to the controller, and then placed into the cache in the event that it gets asked for again. Since the cache is limited in size, there are only so many pieces of data that can be held before the segments must be recycled. Typically the oldest piece of data is replaced with the newest one. This is called *circular*, *first-in, first-out* (*FIFO*) or *wrap-around* caching.

# Cache Circuitry and Operation

- In an effort to improve performance, most hard disk manufacturers today have **implemented enhancements to their cache management circuitry**, particularly on high-end SCSI drives:

- **Adaptive Segmentation:** Conventional caches are chopped into a number of **equal-sized segments**. Since requests can be made for data blocks of different sizes, this can lead to some of the cache's storage in some segments being "left over" and hence wasted (in exactly the same way that slack results in waste in the FAT file system). Many **newer drives dynamically resize the segments based on how much space is required for each access**, to ensure greater utilization. It can also change the number of segments. This is more complex to handle than fixed-size segments, and it can result in waste itself if the space isn't managed properly.

- **Pre-Fetch:** The drive's cache logic, based on analyzing access and usage patterns of the drive, attempts to load into part of the cache data that has not been requested yet but that it *anticipates* will be requested soon. Usually, this means loading additional data beyond that which was just read from the disk, since it is statistically more likely to be requested next. When done correctly, this will improve performance to some degree.

- **User Control:** High-end drives have implemented a set of commands that allows the user detailed control of the drive cache's operation. This includes letting the user enable or disable caching, set the size of segments, turn on or off adaptive segmentation and pre-fetch, and so on.

# Cache Circuitry and Operation

- While obviously improving performance, the limitations of the internal buffer should be fairly obvious. For starters, it helps very little if you are doing a lot of random accesses to data in different parts of the disk, because if the disk has not loaded a piece of data recently in the past, it won't be in the cache. The buffer is also of little help if you are reading a large amount of data from the disk, because normally it is pretty small: if copying a 10 MiB file for example, on a typical disk with a 512 kiB buffer, at *most* 5% of the file could be in the buffer: the rest must be read from the disk itself.

- Due to these limitations, the cache doesn't have as much of an impact on overall system performance as you might think. How much it helps depends on its size to some extent, but at least as much on the intelligence of its circuitry; just like the hard disk's logic overall. And just like the logic overall, it's hard to determine in many cases exactly what the cache logic on a given drive is like.

# Cache Size

- In the last couple of years, hard disk manufacturers have dramatically increased the size of the hard disk buffers in their products. Even as recently as the late 1990s, 256 to 512 kiB was common on consumer drives, and it was not unusual to find only 512 kiB buffers on even some SCSI units (though many had from 1 MiB to 4 MiB). Today, 2 MiB buffers are common on retail IDE/ATA drives, and some SCSI drives are now available with a whopping 16 MiB!

- I believe there are two main reasons for this dramatic increase in buffer sizes. The first is that memory prices have dropped precipitously over the last few years. With the cost of memory only about $1 per MiB today, it doesn't cost much to increase the amount the manufacturers put into their drives. The second is related to marketing: hard disk purchasers have a perception that doubling or quadrupling the size of the buffer will have a great impact on the performance of the hardware.

# Cache Size

- The **cache chip** from a **Seagate Barracuda** hard disk logic board. This chip is the entire cache: it's a 4 Mib chip, which is 512 kiB, the size of the cache on this drive. Some caches use more than one chip, especially the larger ones.

# Cache Size

- The size of the disk's cache *is* important to its overall impact in improving the performance of the system, for the same reason that adding system memory will improve system performance, and why increasing the system cache will improve performance as well. However, the attention that the size of the hard disk buffer is getting today is largely unwarranted. It has become yet another "magic number" of the hardware world that is tossed around too loosely and overemphasized by salespeople. In fact, *a benchmarking comparison done by StorageReview.com* **showed very little performance difference between 512 kiB and 1 MiB buffer** versions of the same Maxtor hard drive. See this section for more on this performance metric.
- So, where does this leave us? Basically, with the realization that the size of the buffer is important only to an extent, and that only large differences (4 MiB vs. 512 kiB) are likely to have a significant impact on performance. Also remember that the size of the **drive's internal buffer will be small on most systems compared to the amount of system memory** set aside by the operating system for its **disk cache**. These two caches, the one inside the drive and the one the operating system uses to avoid having to deal with the drive at all, perform a similar function, and really work together to improve performance.

# Write Caching

- Caching reads from the hard disk and caching writes to the hard disk are similar in some ways, but very different in others. They are the same in their overall objective: to decouple the fast PC from the slow mechanics of the hard disk. The key difference is that a write involves a *change* to the hard disk, while a read does not.
- With no write caching, every write to the hard disk involves a performance hit while the system waits for the hard disk to access the correct location on the hard disk and write the data. As mentioned in the general discussion of the cache circuitry and operation, this takes at least 10 milliseconds on most drives, which is a long time in the computer world and really slows down performance as the system waits for the hard disk. This mode of operation is called ***write-through* caching**. (The contents of the area written actually are put into the cache in case it needs to be *read* again later, but the write to the disk always occurs at the same time.)
- When write caching is enabled, when the system sends a write to the hard disk, the logic circuit records the write in its much faster cache, and then immediately sends back an acknowledgement to the operating system saying, in essence, "all done!" The rest of the system can then proceed on its merry way without having to sit around waiting for the actuator to position and the disk to spin, and so on. This is called ***write-back* caching**, because the data is stored in the cache and only "written back" to the platters later on.

# Write Caching

- Due to this risk, in some situations write caching is not used at all. This is especially true for applications where high data integrity is critical. Due to the improvement in performance that write caching offers, however, it is increasingly being used despite the risk, and the risk is being mitigated through the use of additional technology. The most common technique is simply ensuring that the power does not go off! In high-end server environments, with their uninterruptible power supplies and even redundant power supplies, having unfilled cached writes is much less of a concern. For added peace of mind, better drives that employ write caching have a "write flush" feature that tells the drive to immediately write to disk any pending writes in its cache. This is a command that would commonly be sent before the UPS batteries ran out if a power interruption was detected by the system, or just before the system was to be shut down for any other reason.

# Hard Disk Form Factor

- Most hard disks are designed to be installed on the inside of the PC, and are produced in one of a dozen or so standard sizes and shapes. These standards are called hard disk *form factors* and refer primarily to its external dimensions. The reason for standardizing on form factors is compatibility. Without these standards, hard disks would have to be custom-made to fit different PCs. By agreeing on standards shapes and sizes for hard disks--as well as standard interfaces of course--it is possible for any of the thousands of PC makers to purchase units from any hard disk manufacturer and know that there won't be problems with fit or form during installation.

- Over the life of the PC there have only been a few different hard disk form factors. Since changing a form factor standard requires coordination from the makers of other components (such as the makers of system cases) there is resistance in the industry to change the standard sizes unless there is a compelling reason to do so. (For example, when laptop PCs became popular new, smaller drives were created to save space and power, important goals in the world of mobile computing.)

# Hard Disk Form Factor

- Form factors are generally described by a single metric. For example, the most common form factors today are "3.5-inch" and "2.5-inch".
- These numbers generally refer to the *width* of the drive, but they can be both vague and misleading (nice, huh? :^) ) They usually were chosen for historical reasons and in typically were based on either the platter size of drives that use the form factor, or the width of drives using that form factor. Obviously a single number cannot represent both, and in some cases, it represents neither! For example, 3.5" hard disks are generally 4" wide and use 3.74" platters. :^) (The name in this case comes from the fact that the drives fit in the same space as a 3.5" floppy disk drive!) Much more about the relationship between form factors and platters can be found in the discussion of platter size in the media section. You will also find there a detailed description of the trend towards smaller platters in modern hard disks.

# Hard Disk Form Factor

- The five most popular internal form factors for PC hard disks. Clockwise from the left: 5.25", 3.5", 2.5", PC Card and CompactFlash

# Hard Disk Form Factor

- In this section I examine the major form factors that have been used for internal hard drives in PCs. This includes details on the dimensions of the form factor, especially the different heights associated with each. (Most form factors are actually a *family* of form factors, with different drives varying in the height dimension). In addition to the standard internal drive form factors, I briefly discuss external drives and also removable drive trays, which are sort of a "hybrid" of internal and external designs.
- **Note:** There may be drives available in form factors other than those listed here; I believe I have them all but as always, could be wrong. :^)

# 5.25" Form Factor

- The 5.25" form factor is the oldest in the PC world. Used for the first hard disks on the original IBM PC/XT back in the early 1980s, this form factor has been used for most of the PC's life span, but is now obsolete.
- The basis of the form factor is the 5.25" drive bay used in the first PCs for 5.25" floppy disk drives (themselves obsolete today). These bays still exist today in modern PCs, but are now used primarily for CD-ROM/DVD drives and similar devices, not hard disks. The 5.25" form factor was replaced by the 3.5" form factor for two main reason: first, 5.25" drives are *big* and take up a lot of space; second, 3.5" drives offer better performance;
- The use of 5.25" drives continued as late as the mid-1990s for high-end drives used in servers and other applications where the large size of the platters in these drive was needed to allow drives with high capacities to be created. They mostly disappeared from consumer PC many years prior to that.

# 5.25" Form Factor

- A 3.5" form factor hard disk piggybacked on a 5.25" form factor hard disk to contrast their dimensions. The 5.25" drive here is a Quantum Bigfoot and is the same height as a regular 3.5" low profile drive.

# 5.25" Form Factor

- 5.25" drives generally use **5.12"** platters and have a width of 5.75" and depth of 8.0".
- For many years they were found in only **two different height profiles**: *full-height*, meaning the same height as the floppy drive on the original PC and the bay it used (3.25"); and *half-height*, which is of course half that number.

- In the 1990s, Quantum launched a new line of 5.25" drives named the ***Bigfoot* family**, which reintroduced 5.25" drives to the consumer marketplace. These were sold as "economy" drives and due to the larger platter size, offered a lot of capacity--but due to slower spindle speeds and a "value line" design, not much performance. They were popular with many PC manufacturers but eventually were phased out. These drives used what had up to that point been non-standard heights for 5.25" drives, typically 1" high or less. Quantum calls these drives ***low-profile*** or ***ultra-low-profile***.

# 5.25" Form Factor

- Here are the statistics and applications of the different profiles used in the 5.25" form factor:

| Form Factor | Width (in) | Depth (in) | Height (in) | Application |
|---|---|---|---|---|
| **5.25" Full-Height** | 5.75 | 8.0 | 3.25 | All drives in early 1980s; Large capacity drives with many platters as late as the mid-1990s |
| **5.25" Half-Height** | 5.75 | 8.0 | 1.63 | Early 1980s through early 1990s |
| **5.25" Low-Profile** | 5.75 | 8.0 | 1.0 | Quantum Bigfoot, mid-to-late 1990s |
| **5.25" Ultra-Low-Profile** | 5.75 | 8.0 | 0.75 - 0.80 | Quantum Bigfoot, mid-to-late 1990s |

- Interestingly, despite the general trend to smaller drives, the drives that continued to use the 5.25" form factor through the late 1980s and early 1990s were more often found as full-height devices than half-height ones. This may be due to the fact that their niche became applications where a lot of storage was needed, so the ability to fit many more platters in that nice, roomy 3.25" high package was attractive.

# 3.5" Form Factor

- The 3.5" form factor is *the* standard in the PC world today, and has been for about the last decade. Drives of this size are found almost exclusively now in modern desktop PCs, and even in servers and larger machines. The only major market where other form factors hold sway over 3.5" is that for laptops and other portable devices, where the reduced size of 2.5" and smaller form factors is important.
- Like the 5.25" form factor before it, the 3.5" form factor is named not for any dimension of the drives themselves, but rather for the fact that they were designed to fit into the same drive bay as 3.5" floppy disk drives. 3.5" form factor drives traditionally have used 3.74" platters with an overall drive width of 4.0" and depth of about 5.75". In recent years, 3.5" form factor drives with platters smaller than 3.74"--in some cases much smaller--have appeared on the market. Most 10,000 RPM spindle speed drives reduce the size of the platters to 3", and the new 15,000 RPM Seagate drive has platters just 2.5" in diameter. The shrinking media size is done for performance reasons, but the 3.5" form factor is maintained for compatibility (these high-end drives are designed to go into expensive servers, not laptops!) For this reason, it is no longer the case that you can tell the size of a drive's platters by its form factor. See the discussion of platter size for more details.
- 3.5" form factor drives come in two general profiles: the larger is the so-called *half-height* drive, which is 1.63" in height. This name is kind of funny, since it is "half" of a height that never existed for 3.5" form factor drives. The name was derived from the fact that these drives are the same height as half-height 5.25" form factor drives, which are half the height of full-height 3.25" high drives in that form factor. Half-height 3.5" form factor drives are still used today, but only in servers and other high-end platforms. The standard for 3.5" is 1" height, which is commonly called *slimline* or *low-profile*, but just as commonly given no name at all and assumed as the default. The reason for the smaller size being the standard is that 1" is the height of a standard 3.5" floppy disk drive and 3.5" drive bay. In addition, there are some drives that are reduced in size from the 1" standard, using for example 0.75" height. Here are the standard profiles for 3.5" form factor drives:

# 3.5" Form Factor

- Here are the standard profiles for 3.5" form factor drives:

| Form Factor | Width (in) | Depth (in) | Height (in) | Application |
|---|---|---|---|---|
| 3.5" Half-Height | 4.0 | 5.75 | 1.63 | High-end, high-capacity drives |
| 3.5" Low-Profile | 4.0 | 5.75 | 1.0 | Industry standard, most common form factor for PC hard disks |

- It is likely that the 3.5" form factor will continue to be the industry standard in PCs for years to come, due to the enormous installed base of systems that use this size, and no real compelling reason to change the form factor for the typical desktop machine.

# 2.5" Form Factor

- 2.5" form factor drives are the standard today for notebook computers (although not all notebooks use them, most do). Since the notebook market continues to grow by leaps and bounds, sales of 2.5" form factor drives have been increasing, on a percentage basis, faster than probably any other segment of the hard disk market overall. While older laptops originally used 3.5" drives, the move to 2.5" was done for several reasons that are very important to mobile PC users (you can also find related information on the reduction of platter sizes here):

- **Size Reduction:** Smaller drives take up less space and allow for laptops to be reduced in size. This trend began with the first 2.5" drives and continues with the continuous reduction in the heights of 2.5" drives (see below) and also the creation of still-smaller form factors.

- **Power Reduction:** Smaller drives use less power, important for PCs that run on batteries.

- **Enhanced Rigidity:** Smaller drives use smaller platters, which are less susceptible to damage as a result of shock, always a concern for a drive that will be moved around (often while operating!)

# 2.5" Form Factor

- An 8.4 GB, **2.5"** form factor IBM hard disk. Note the single connector in the front, which is mated to a matching connector in the laptop's hard disk bay. This allows the drive to be easily replaced at a later time.
- The connector on the hard disk itself just uses straight pins like a 3.5" hard disk form factor drive; the drive is mounted into a carrier here, and the thin circuit board you can see in the front "adapts" the regular pin connector into the single Centronics-style connector my notebook uses

# 2.5" Form Factor

- Unlike its larger, older siblings, the 2.5" form factor actually *is* named for the platter size of drives that use it (finally! :^) ) The width of a 2.5" drive is 2.75", and depth is 3.94". These drives originally came in just one height (0.75" or 19 mm). Since for any storage technology level there is a tradeoff between size and capacity, over time several different heights were created in this form factor as standards for mobile PC users with different requirements. They are usually specified in metric (mm) and to my knowledge have no fancy names:

| Form Factor | Width (in) | Depth (in) | Height (in) | Application |
|---|---|---|---|---|
| **2.5" 19 mm Height** | 2.75 | 3.94 | 0.75 | Highest-capacity 2.5" drives, used in full-featured laptop systems |
| **2.5" 17 mm Height** | 2.75 | 3.94 | 0.67 | Mid-range capacity drives used in some laptop systems |
| **2.5" 12.5 mm Height** | 2.75 | 3.94 | 0.49 | Low-capacity drives used in small laptops (subnotebooks) |
| **2.5" 9.5 mm Height** | 2.75 | 3.94 | 0.37 | Lowest-capacity drives used in very small laptops (mini-subnotebooks) |

- 2.5" drives are pretty much entrenched as the standard for laptop machines. They are also used occasionally in industrial applications, where the smaller size and increased ruggedness of portable drives is important.

# PC Card (PCMCIA) Form Factor

- One of the most popular interfaces used in the notebook PC world is the ***PC Card* interface, also sometimes called *PCMCIA*** after the group that created it in the late 1980s. This interface standard was created to allow for easier expansion of notebook systems, which at that time had very few options for adding hardware at all.

- Despite the relatively small size of cards that adhere to the PC Card standard, **hard disk engineers have managed to create hard disks to fit**. There are actually **3 PC Card form factor sizes**, defined by the PCMCIA in 1995. The width and depth of these devices is exactly the same as that of a credit card, which I am sure is not a coincidence! They are all 2.13" wide and 3.37" deep. The three sizes differ only in their height. Type I devices are 3.3 mm thick; Type II devices are 5.0 mm thick, and Type III devices are 10.5 mm thick. Originally, the intention was for solid state devices like memory, modems and the like to use the Type I and Type II devices, while the Type III devices were for hard disks. Due to the extreme height limits of PC Cards, it is difficult to make hard disks that will fit into the allowed space, and most PC Card hard disks are Type III. However, advances in miniaturization have allowed some companies to now make hard disks that actually fit into the Type II PC Card form factor as well. Since most laptops can only accept *either* two Type I/II cards or a single Type III, this is a significant advantage.

# PC Card (PCMCIA) Form Factor

- Here's a summary table of the different sizes:

| Form Factor | Width (in) | Depth (in) | Height (in/mm) | Application |
|---|---|---|---|---|
| **PC Card Type I** | 2.13 | 3.37 | 0.13 / 3.3 | Not used for hard disks (yet?) |
| **PC Card Type II** | 2.13 | 3.37 | 0.20 / 5.0 | Smaller-capacity expansion hard disks for laptops and consumer electronics |
| **PC Card Type III** | 2.13 | 3.37 | 0.41 / 10.5 | Higher-capacity expansion hard disks for laptops |

- The 2.13" width of this form factor puts a hard limit on the platter size of these drives--even 2.5" platters are too large. Most PC Card drives today use 1.8" platters. Interestingly, the first hard drive to use the PC Card form factor was probably the Hewlett Packard *Kittyhawk* drive, with much smaller 1.3" platters. This drive is a good example of a technology being "ahead of its time". It was actually introduced way back in 1992, very early on for such miniaturized technology. Unfortunately, at the time the market may not have been big enough to provide HP with sufficient revenues to keep making it. The Kittyhawk was used in early hand-helds and other small consumer electronic devices (even printers!) for a while, but was eventually discontinued, and HP is no longer making hard disk drives of any sort. If this technology had been introduced five years later, it may have been a runaway success; certainly IBM is having great success with its slightly-smaller Microdrive.

# PC Card (PCMCIA) Form Factor

- A CompactFlash card (left) and a PC Card (right).The quarter is included for size context. **Neither of these is a hard disk** (though the SanDisk is a solid-state flash card "hard disk") but are the same size and shape as hard drives of their respective form factor.

# PC Card (PCMCIA) Form Factor

- **Note:** Many companies also make "**solid state drives**" using the PC Card form factor. These perform the same function as hard disk drives, but are not hard disks at all: they are actually flash memory, a type of ROM. I discuss this in more detail in the section on the CompactFlash form factor.

- Interestingly, with a couple of exceptions, most of the smaller PC Card hard disks are not made by the bigger, well-known hard disk companies, but rather smaller niche companies. I am not sure what the reason is for this. I suspect that there just may not be enough profit potential there for the big names to bother with this market, which is small compared to the market for mainstream PC drives.

# CompactFlash Form Factor

- In much the same way that the need for expansion capabilities in laptops led to the creation of the PCMCIA and PC Card devices, a consortium of electronics and computer industry companies in 1995 formed the ***CompactFlash Association*** to promote a new form factor called, of course, ***CompactFlash*** (abbreviated *CF* or *CF*+). This form factor is similar to the PC Card form factor, but amazingly enough, even smaller. CF cards are intended to be used not in laptop PCs **but smaller electronic devices such as hand-held computers**, digital cameras and communications devices (including cellular phones).
- Unlike the PC Card standard, which is used for a wide variety of devices, CompactFlash is primarily designed around permanent storage. The "flash" in "CompactFlash" is from the primary technology used in these cards: flash memory.
- Flash memory is really electrically-erasable read-only memory, typically used in regular PCs only for holding the motherboard's BIOS code. The word "flash" refers to the ability to write and erase these ROMs electrically. Much the way you can "flash" your motherboard BIOS to update it, these flash memory storage cards have controllers in them that do this as part of their normal operation. Unlike regular memory, flash memory is of course non-volatile and retained when the power is removed.

# CompactFlash Form Factor

- The intention of the CompactFlash form factor was to allow consumer electronic devices to use these CompactFlash cards for their equivalent of a hard disk.
- Since the flash memory is not volatile, it does perform the same general function as a hard disk. Like PCMCIA devices, variants of the form factor were developed, differing only in thickness; the thicker cards provide more space to pack in additional flash memory chips for greater capacity.

# CompactFlash Form Factor

- Here are the dimensions of the two types of CompactFlash cards:

| Form Factor | Width (in) | Depth (in) | Height (in/mm) | Application |
|---|---|---|---|---|
| **CF+ Type I** | 1.69 | 1.42 | 0.13 / 3.3 | Smaller-capacity flash cards for digital cameras, hand-held computers and consumer electronics; not used for hard disks (yet) |
| **CF+ Type II** | 1.69 | 1.42 | 0.20 / 5.0 | Larger-capacity flash cards and hard disks for digital cameras, hand-held computers and consumer electronics |

- As you can see, the CF form factors are very small: so small that they were probably never designed with the thought that anyone would make a true hard disk using them. The engineers at IBM however had a different idea! In 1999, while the makers of regular flash memory cards were struggling to reach 64 MB capacity, IBM introduced the Microdrive, a true hard disk that fits into the small confines of the CF form factor. The original Microdrive was available in either 170 MB or 340 MB capacities, which is pretty impressive considering that the drive uses a single 1" platter... even more impressive is the new Microdrive released in 2000 with a whopping 1 GB capacity! The Microdrive uses the CF Type II format, which is 5.0 mm thick. No current hard disks are made for the even smaller Type I size (only 3.3 mm thick) but I have heard rumors that at least one company is working on it, and I'd be surprised if IBM themselves didn't have something brewing in this regard also. Pretty cool.

# CompactFlash Form Factor

- IBM's amazing Microdrive.

- **Smaller drives generally have less performance than full-sized ones and the Microdrive is no exception**; it certainly doesn't compete with the newest 2.5" or 3.5" form factor drives for performance (though in many ways it is superior to flash memory chips). For its size it is certainly no slouch, though: it has a 3,600 RPM spindle speed, and very decent maximum areal density of as much as 15.2 Gbits/in2. (The original generation actually used a faster 4,500 RPM spindle; this was probably lowered to reduce power consumption and heat, both very serious issues for this form factor... however, the first Microdrives also had only one-third the areal density of the 1 GB model.) The extremely small size of the drive allows it to spin up to full speed in only half a second, a fraction of the time required by most large drives. This lets the Microdrive power down often when idle to save power, an essential feature for devices that use small batteries.
- **Tip:** If you want to use the Microdrive in a laptop, an inexpensive adapter is available to convert it to the PC Card type II form factor.
- As areal density continues to increase these drives will only continue to increase in capacity. Over the next few years they will battle for sales with flash memory cards; this market is likely to grow as the number of hand-held PCs, digital cameras and other electronic gizmos needing lots of storage continues to increase dramatically.

# Form Factor Comparison

- For ease of comparison, the summary table below lists all of the standard internal hard disk form factors with their dimensions, typical [platter sizes]() found in hard disks that use the form factor, and common applications:

| Form Factor | Profile | Platter Size (in) | Width (in) | Depth (in) | Height (in) | Application |
|---|---|---|---|---|---|---|
| 5.25" | Full-Height | 5.12 | 5.75 | 8.0 | 3.25 | All drives in early 1980s; Large capacity drives with many platters as late as the mid-1990s |
| | Half-Height | 5.12 | 5.75 | 8.0 | 1.63 | Early 1980s through early 1990s |
| | Low-Profile | 5.12 | 5.75 | 8.0 | 1.0 | Quantum Bigfoot, mid-to-late 1990s |
| | Ultra-Low-Profile | 5.12 | 5.75 | 8.0 | 0.75 - 0.80 | Quantum Bigfoot, mid-to-late 1990s |
| 3.5" | Half-Height | 2.5, 3.0, 3.74 | 4.0 | 5.75 | 1.63 | High-end, high-capacity drives |
| | Low-Profile | 2.5, 3.0, 3.74 | 4.0 | 5.75 | 1.0 | Industry standard, most common form factor for PC hard disks |
| 2.5" | 19 mm Height | 2.5 | 2.75 | 3.94 | 0.75 | Highest-capacity 2.5" drives, used in full-featured laptop systems |
| | 17 mm Height | 2.5 | 2.75 | 3.94 | 0.67 | Mid-range capacity drives used in some laptop systems |
| | 12.5 mm Height | 2.5 | 2.75 | 3.94 | 0.49 | Low-capacity drives used in small laptops (subnotebooks) |
| | 9.5 mm Height | 2.5 | 2.75 | 3.94 | 0.37 | Lowest-capacity drives used in very small laptops (mini-subnotebooks) |
| PC Card | Type I | 1.8 | 2.13 | 3.37 | 0.13 | Not used for hard disks (yet?) |
| | Type II | 1.8 | 2.13 | 3.37 | 0.20 | Smaller-capacity expansion hard disks for laptops and consumer electronics |
| | Type III | 1.8 | 2.13 | 3.37 | 0.41 | Higher-capacity expansion hard disks for laptops |
| Compact Flash | Type I | 1.0 | 1.69 | 1.42 | 0.13 | Smaller-capacity flash cards for digital cameras, hand-held computers and consumer electronics; not used for hard disks (yet) |
| | Type II | 1.0 | 1.69 | 1.42 | 0.20 | Larger-capacity flash cards and hard disks for digital cameras, hand-held computers and consumer electronics |

# External Hard Disks

- Some hard disks are available as external drives, especially ones using the SCSI interface. These really don't differ much from internal drives, except that they include an **additional outer plastic shell,** a **power supply** to run the disk, and of course, a larger price tag. :^)

- They do offer some advantages over internal drives: more expandability, **easier installation**, usually **better cooling**, and also interoperability with other systems that use SCSI. Since they are external, they don't have to be made in standardized form factors.

- At one time, these SCSI external drives were more popular than they are today. The higher cost of these drives due to the duplication of support hardware, and the extra space required to put them "somewhere" are probably the main reasons why they are less common today than they once were.

# External Hard Disks

- The IBM Travelstar E, an external hard disk using a PC Card interface card

- External drives have found a new market role of sorts as expansion and backup devices for portable computers. Many varieties are now available using either the parallel port or the PC card interface. In the latter design, the hard disk is in an external enclosure, with an interface cable that runs to a PC card. The card connects to the laptop through a PC card slot. The fact that the hard disk is not constrained by the physical limits of the small PC card slot means it can be made much larger than the small drives available in the PC card form factor, while retaining the portability advantages of the PC card interface. Some companies also make just the enclosure itself, with the PC card interface; you supply your own standard internal hard disk. This can save money by letting you "recycle" an older internal drive, and gives you the flexibility of swapping in new drives as better technology becomes available.

# Removable Hard Disk Trays

- An interesting compromise between internal and external hard disks is the *removable hard disk drive tray*. A tray is installed into a standard PC case drive bay that allows regular internal hard disks to be placed into it. You then can swap the internal hard disk with another one without opening up the case, allowing you to use hard disks as a removable storage medium. In a way, the concept is similar to the way a removable car stereo is designed. These trays are also commonly called *mobile racks*, *drive caddies* and several other names.

# Removable Hard Disk Trays

- Shown above is the *Kingwin KF-21-IPF mobile rack system* that I use on my work desktop PC for backups and file archiving.
- The drive fits into the removable tray (bottom) which fits into the stationary docking station (top).
- The stationary portion is installed into a standard 5.25" drive bay and connected to a regular IDE/ATA cable.
- On the right-hand side you can see the lock that secures the tray in place, as well as indicator lights for power and drive activity.

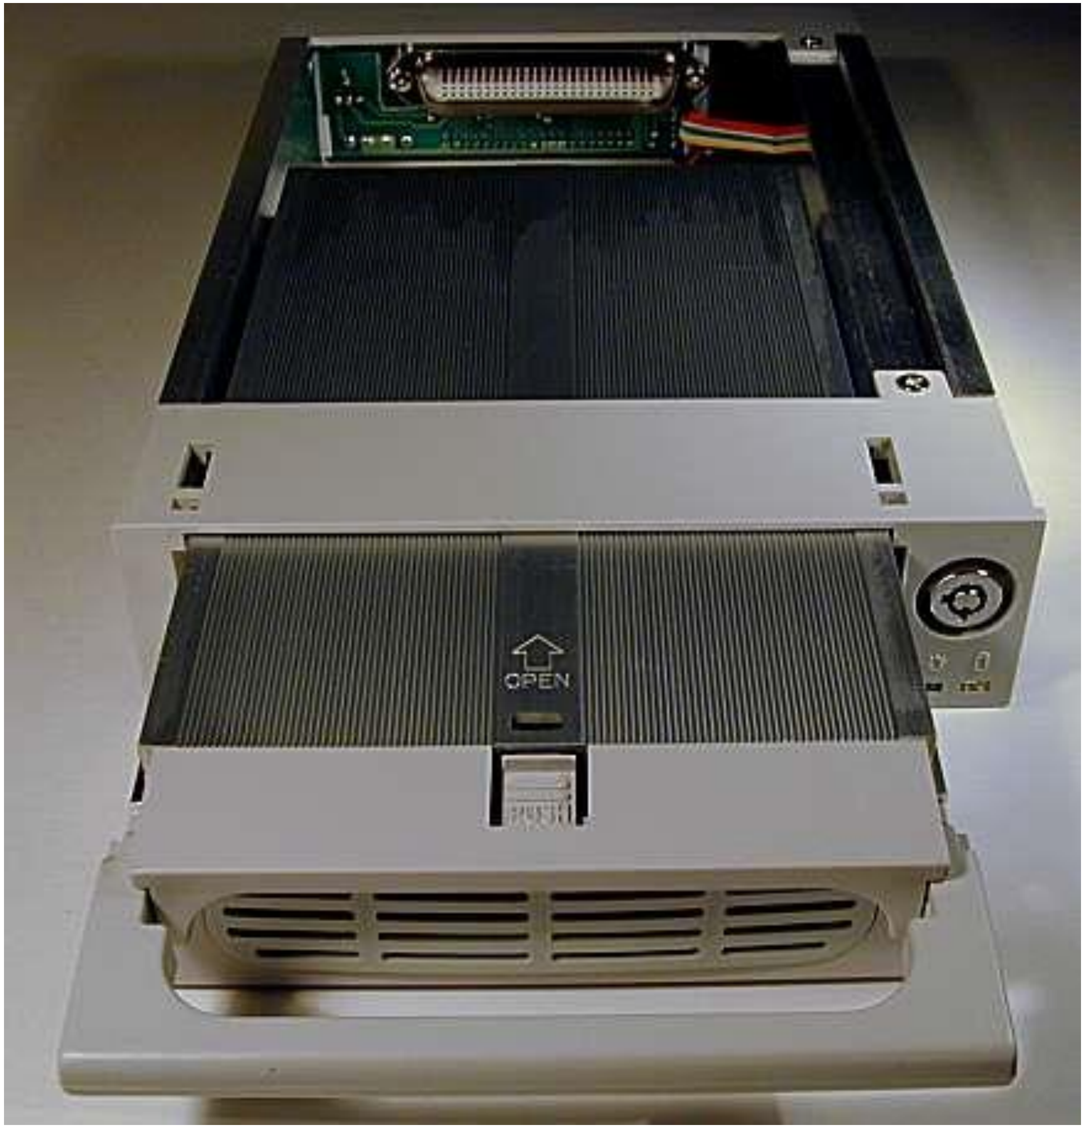

# Air Circulation and Air Filtration

- Closeup shot of the breather holes in the top of a hard disk case. Part of the breather filter can be seen just under the holes.

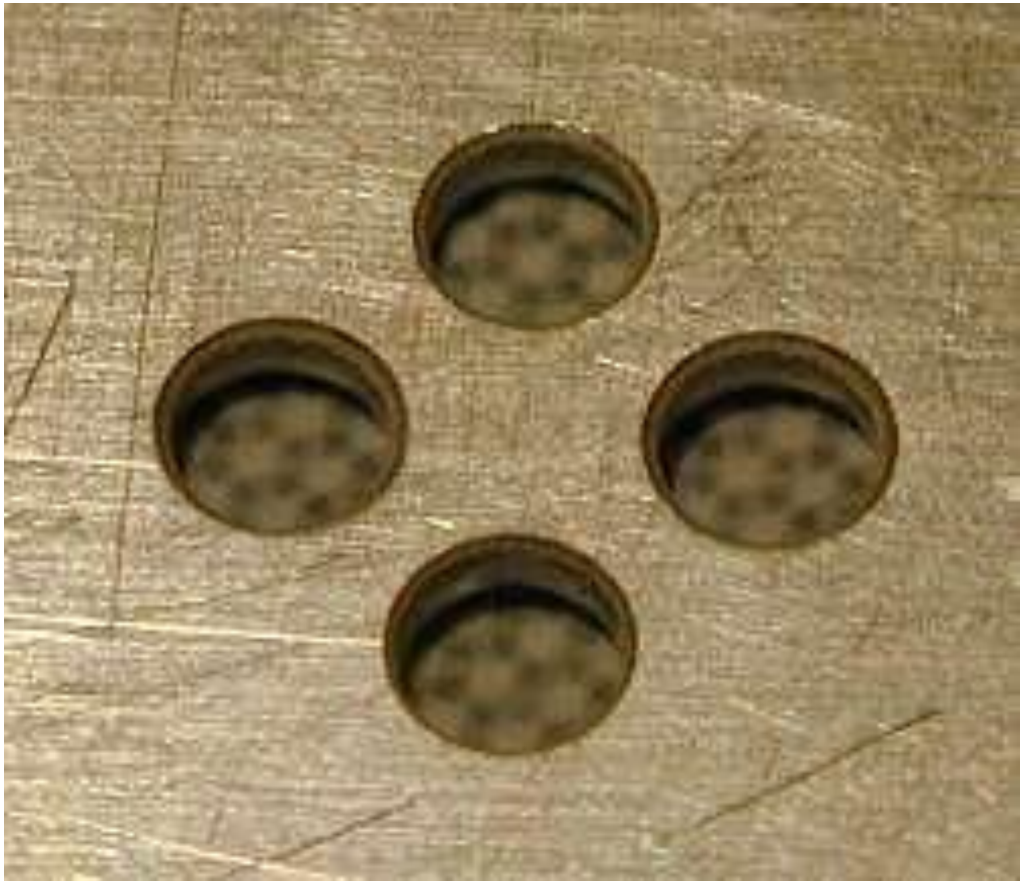

# Temperature Limits and Drive Cooling

- If the system's cooling isn't sufficient for a hot-running hard disk, you have some options. Various manufacturers make devices specifically for cooling hard disks. They usually come in two basic forms:
- **Drive Coolers:** These are essentially fan and heat sink combos, similar to those used for CPUs, but designed specially to be used with hard disks. They are attached to the drive using thermal conductive tape, and blow air directly onto the drive case to cool it.

- A drive cooler mounted on top of a standard 3.5“ form factor hard disk. The power plug is visible at right.

# Temperature Limits and Drive Cooling

- **Bay Coolers:** These devices are similar in design to the drive bay adapters available for mounting a 3.5" form factor drive into a 5.25" drive bay, except that they add cooling for the drive. The disk is mounted into the cooler, which contains one or more integrated fans. The cooler then is mounted into one of the larger 5.25" drive bays found in most PCs.

- A 5.25" bay cooler with a 3.5" form factor hard disk installed in it. This view is from the inside of the case; the external faceplate is at rear.*Original image* © *PC Power & Cooling, Inc.*

# Retail and OEM Packaging

- Most hard disk drive models are sold as **two different packages**: *OEM drives* and *retail drives* (sometimes called *retail kits*). Retail drives are of course drives that are distributed to retail stores and online dealers for sale to the general public. OEM drives are those sold to system manufacturers in large quantity, for inclusion in PCs built for resale: "OEM" stands for "original equipment manufacturer" and refers to a company that makes PCs (or other equipment).

- Retail packaged hard disk drives normally include the following (though this varies from manufacturer to manufacturer and model to model):
- **Hard Disk Drive:** The hard disk drive itself, in an anti-static bag or similar package.
- **Installation Instructions:** Instructions on how to configure and install the hard disk.
- **Drivers and/or Overlay Software:** A floppy disk or CD-ROM containing any necessary drivers or utilities, and usually, a copy of that manufacturer's version of drive overlay software for working around BIOS capacity problems in older systems.
- **Mounting Hardware:** A set of appropriately-sized screws for mounting the drive into the system case.
- **Interface Cable:** A cable of the correct type for the drive's interface.
- **Warranty Card:** A card describing the warranty provided on the drive, usually three or five years in length.
- **Pretty Box:** A very colorful box that looks nifty and holds all of the above. :^)
- In contrast, OEM drives typically contain the following:
- **Hard Disk Drive:** The hard disk drive itself, in an anti-static bag or similar package.
- **Jumpers:** One ore more jumpers needed for configuring the drive.

# Retail and OEM Packaging

- And that's it (heck, even the jumpers aren't a sure thing. :^) ) The reason that OEM packaging is so "plain" is that most OEMs do not need the additional support materials and packaging required for a proper retail package--they are just going to put the drives into PCs, not resell them to end users. If you are SuperPeeceez Inc. and make 10,000 systems a month, you just want 10,000 drives--not 10,000 fancy boxes, 10,000 warranty cards, 10,000 driver disks, and so on. So the drive manufacturers just ship OEMs lots of bare drives, with the minimum required to protect them, and that's all. By skipping all the extras, the OEM is able to buy the drives at a lower price from the manufacturer (and the fact that they are buying them by the thousand certainly helps with price as well!)


- Contents of a Western Digital hard disk retail box kit.

# Hard Disk Handling

- Hard disks are very delicate and sensitive instruments. While certainly it is true that newer (and especially *smaller*) drives are less fragile than their older counterparts, it is also true that all hard disk drives have to be properly handled to avoid damage. In most cases, handling of drives is something that happens very little anyway: you get the drive, you install it and you leave it there. So for most people handling isn't too much of an issue. For those that handle drives a great deal however, proper handling technique is essential.
- The first thing you will notice when you deal with hard disks (and most other computer components) is that they are always transported in an *anti-static bag*. This is of course to prevent the damage that can occur to the hard disk's circuits as a result of *electrostatic discharge* or *ESD*. This is such a common problem when dealing with components that I put a general warning about it right in the Site Guide.


- A "new" (well, actually a few years old, but unused) hard disk in its original anti-static bag.

# Hard Disk Handling

- Seagate has actually come up with a neat improvement on the standard anti-static bag (which has been around for years and years). They call it the ***SeaShell*** (ha ha, good one guys) and instead of being a thin plastic bag, it's a solid plastic clam-shell case that not only provides ESD protection for the drive, but physically protects it against shock as well. Instead of being a smoky gray, the SeaShell is clear plastic, so you can examine the drive without having to open up the package. And these little cases are also both recyclable and easily reusable. I hope other hard drive makers start using similar devices.

- A Seagate "SeaShell", containing a Cheetah SCSI drive.*Original image* © *Seagate Technology*