

**DataRecover bvba**

**Burchtstraat 2  
9300 Aalst  
Belgium**

**Tel. : 0475 / 76.66.73**

**Fax : 053 / 41.75.84**




**Registratie nr. : 0826.283.018  
BTW nr. : BE 0826.283.018**

**Bank : 001-6132677-25  
Iban : BE70 0016 1326 7725  
Bic : GEBABEBB**

**Email : info@datarecover.be  
Website : www.datarecover.be**

## Construction and Operation of the Hard Disk

To many people, a hard disk is a "black box" of sorts--it is thought of as just a small device that "somehow" stores data. There is nothing wrong with this approach of course, as long as all you care about is that it stores data. If you use your hard disk as more than just a place to "keep stuff", then you want to know more about your hard disk. It is hard to really understand the factors that affect performance, reliability and interfacing without knowing how the drive works internally. Fortunately, most hard disks are basically the same on the inside. While the technology evolves, many of the basics are unchanged from the first PC hard disks in the early 1980s.



In this section we dive into the guts of the hard disk and discover what makes it tick. We look at the various key components, discuss how the hard disk is put together, and explore the various important technologies and how they work together to let you read and write data to the hard disk. My goal is to go beyond the basics, and help you really understand the design decisions and tradeoffs made by hard disk engineers, and the ways that new technologies are being employed to increase capacity and improve performance.

# Hard Disk Operational Overview

As an illustration, I'll describe here in words how the various components in the disk interoperate when they receive a request for data. Hopefully this will provide some context for the descriptions of the components that follow in later sections.

A hard disk uses round, flat disks called *platters*, coated on both sides with a special *media* material designed to store information in the form of magnetic patterns. The platters are mounted by cutting a hole in the center and stacking them onto a *spindle*. The platters rotate at high speed, driven by a special *spindle motor* connected to the spindle. Special electromagnetic read/write devices called *heads* are mounted onto *sliders* and used to either record information onto the disk or read information from it. The sliders are mounted onto *arms*, all of which are mechanically connected into a single assembly and positioned over the surface of the disk by a device called an *actuator*. A *logic board* controls the activity of the other components and communicates with the rest of the PC.


Each *surface* of each platter on the disk can hold tens of billions of individual bits of data. These are organized into larger "chunks" for convenience, and to allow for easier and faster access to information. Each platter has two heads, one on the top of the platter and one on the bottom, so a hard disk with three platters (normally) has six surfaces and six total heads. Each platter has its information recorded in concentric circles called *tracks*. Each track is further broken down into smaller pieces called *sectors*, each of which holds 512 bytes of information.

The entire hard disk must be manufactured to a high degree of precision due to the extreme miniaturization of the components, and the importance of the hard disk's role in the PC. The main part of the disk is isolated from outside air to ensure that no contaminants get onto the platters, which could cause damage to the read/write heads.

Here's an example case showing in brief what happens in the disk each time a piece of information needs to be read from it. This is a highly simplified example because it ignores factors such as disk caching, error correction, and many of the other special techniques that systems use today to increase performance and reliability. For example, sectors are not read individually on most PCs; they are grouped together into continuous chunks called *clusters*. A typical job, such as loading a file into a spreadsheet program, can involve thousands or even millions of individual disk accesses, and loading a 20 MB file 512 bytes at a time would be rather inefficient:

1. The first step in accessing the disk is to figure out where on the disk to look for the needed information. Between them, the application, operating system, system BIOS and possibly any special driver software for the disk, do the job of determining what part of the disk to read.
2. The location on the disk undergoes one or more translation steps until a final request can be made to the drive with an address expressed in terms of its *geometry*. The geometry of the drive is normally expressed in terms of the cylinder, head and sector that the system wants the drive to read. (A cylinder is equivalent to a track for addressing purposes). A request is sent to the drive over the disk drive interface giving it this address and asking for the sector to be read.
3. The hard disk's control program first checks to see if the information requested is already in the hard disk's own internal buffer (or *cache*). If it is then the controller supplies the information immediately, without needing to look on the surface of the disk itself.
4. In most cases the disk drive is already spinning. If it isn't (because power management has instructed the disk to "spin down" to save energy) then the drive's controller board will activate the spindle motor to "spin up" the drive to operating speed.
5. The controller board interprets the address it received for the read, and performs any necessary additional translation steps that take into account the particular characteristics of

- the drive. The hard disk's logic program then looks at the final number of the cylinder requested. The cylinder number tells the disk which track to look at on the surface of the disk. The board instructs the actuator to move the read/write heads to the appropriate track.
6. When the heads are in the correct position, the controller activates the head specified in the correct read location. The head begins reading the track looking for the sector that was asked for. It waits for the disk to rotate the correct sector number under itself, and then reads the contents of the sector.
  7. The controller board coordinates the flow of information from the hard disk into a temporary storage area (buffer). It then sends the information over the hard disk interface, usually to the system memory, satisfying the system's request for data.



## Hard Disk Platters and Media

Every hard disk contains one or more flat disks that are used to actually hold the data in the drive. These disks are called *platters* (sometimes also "disks" or "discs"). They are composed of two main substances: a *substrate* material that forms the bulk of the platter and gives it structure and rigidity, and a *magnetic media coating* which actually holds the magnetic impulses that represent the data. Hard disks get their name from the rigidity of the platters used, as compared to floppy disks and other media which use flexible "platters" (actually, they aren't usually even called platters when the material is flexible.)

The platters are "where the action is"--this is where the data itself is recorded. For this reason the quality of the platters and particularly, their media coating, is critical. The surfaces of each platter are precision machined and treated to remove any imperfections, and the hard disk itself is assembled in a *clean room* to reduce the chances of any dirt or contamination getting onto the platters.



This Barracuda hard disk has 10 platters.

The form factor of the hard disk also has a great influence on the number of platters in a drive. Even if hard disk engineers wanted to put lots of platters in a particular model, the standard PC "slimline" hard disk form factor is limited to 1 inch in height, which limits the number of platters that can be put in a single unit. Larger 1.6-inch "half height" drives are often found in servers and usually have many more platters than desktop PC drives. Of course, engineers are constantly working to reduce the amount of clearance required between platters, so they can increase the number of platters in drives of a given height.

### Tracks and Sectors

Platters are organized into specific structures to enable the organized storage and retrieval of data. Each platter is broken into *tracks*--tens of thousands of them--which are tightly-packed concentric circles. These are similar in structure to the annual rings of a tree (but *not* similar to the grooves in a vinyl record album, which form a connected spiral and not concentric rings).

A track holds too much information to be suitable as the smallest unit of storage on a disk, so each one is further broken down into *sectors*. A sector is normally the smallest individually-addressable unit of information stored on a hard disk, and normally holds 512 bytes of information. The first PC hard disks typically held 17 sectors per track. Today's hard disks can have *thousands* of sectors in a single track, and make use of zoned recording to allow more sectors on the larger outer tracks of the disk.




A platter from a 5.25" hard disk, with 20 concentric tracks drawn over the surface. This is far lower than the density of even the oldest hard disks; even if visible, the tracks on a modern hard disk would require high magnification to resolve. Each track is divided into 16 imaginary sectors. Older hard disks had the same number of sectors per track, but new ones use zoned recording with a different number of sectors per track in different zones of tracks.

## Head Actuator

The actuator is the device used to position the head arms to different tracks on the surface of the platter (actually, to different cylinders, since all head arms are moved as a synchronous unit, so each arm moves to the same track number of its respective surface). The actuator is a very important part of the hard disk, because changing from track to track is the only operation on the hard disk that requires active movement: changing heads is an electronic function, and changing sectors involves waiting for the right sector number to spin around and come under the head (passive movement). Changing tracks means the heads must be shifted, and so making sure this movement can be done quickly and accurately is of paramount importance. This is especially so because physical motion is so *slow* compared to anything electronic--typically a factor of 1,000 times slower or more.


Head actuators come in two general varieties:

- **Stepper Motors:** Originally, hard disk drives used a *stepper motor* to control the movement of the heads over the surface of the platters. A regular motor turns in a rotary fashion continuously; it can stop at any point in its rotation as it spins around, kind of like the second hand on a wind-up wristwatch. A stepper motor can only stop at predefined "steps" as it turns around, much the way the second hand turns on an electronic, quartz wristwatch. A hard drive using a stepper motor for an actuator attaches the arms to the motor, and each time the motor steps one position clockwise or counterclockwise, the arms move in or out one position. Each position defines a track on the surface of the disk. Stepper motors are also commonly used for both turning the spindle and positioning the head on floppy disk drives. If you have a floppy drive, find one of its motors and turn it slowly with your hand; you will feel the discrete step-wise nature of its motion.



A stepper motor actuator. The motor moves in steps, which you can feel if you move the motor shaft by hand. The shaft has two thin strips of metal wrapped around it, which are connected to a pivot that is rigidly attached to the actuator arms. As the motor shaft turns, one half of this "split band" coils onto the shaft and the other half uncoils. When the motor turns in the opposite direction the process reverses. As this occurs the pivot moves and in doing so, moves the actuator arms and the hard disk heads.

- Voice Coils:** The actuator in a modern hard disk uses a device called a *voice coil* to move the head arms in and out over the surface of the platters, and a closed-loop feedback system called a *servo system* to dynamically position the heads directly over the data tracks. The voice coil works using electromagnetic attraction and repulsion. A coil is wrapped around a metal protrusion on the end of the set of head arms. This is mounted within an assembly containing a strong permanent magnet. When current is fed to the coil, an electromagnetic field is generated that causes the heads to move in one direction or the other based on attraction or repulsion relative to the permanent magnet. By controlling the current, the heads can be told to move in or out much more precisely than using a stepper motor. The name "voice coil" comes from the resemblance of this technology to that used to drive audio speakers, which are also basically electromagnets. All PC hard disk voice coil actuators are *rotary*, meaning that the actuator changes position by rotating on an axis.



A partially-disassembled voice coil actuator. The magnet assembly has been unscrewed from its mounting and pulled to the left to expose the coil. The magnet assembly consists of two metal plates (top one easily visible above, and part of the bottom one visible.) The magnet itself is mounted on the underside of the top plate, and spacers used between the plates to create the gap for the coil assembly. Being non-ferrous the coil moves freely between the plates, rotating the actuator on its axis as its magnetic polarity is changed. (Incidentally, the magnet is strong enough that after removing the spacers between the plates, the bottom plate got "stuck" on the magnet and required considerable effort to remove!)

## Head Crashes

Since the read/write heads of a hard disk are floating on a microscopic layer of air above the disk platters themselves, it is possible that the heads can make contact with the media on the hard disk under certain circumstances. Normally, the heads only contact the surface when the drive is either starting up or stopping. Considering that a modern hard disk is turning over 100 times a second, this is not a good thing. :^)

If the heads contact the surface of the disk while it is at operational speed, the result can be loss of data, damage to the heads, damage to the surface of the disk, or all three. This is usually called a *head crash*, two of the most frightening words to any computer user. :^) The most common causes of head crashes are contamination getting stuck in the thin gap between the head and the disk, and shock applied to the hard disk while it is in operation.

Despite the lower floating height of modern hard disks, they are in many ways less susceptible to head crashes than older devices. The reason is the superior design of hard disk enclosures to eliminate contamination, more rigid internal structures and special mounting techniques designed to eliminate vibration and shock. The platters themselves usually have a protective layer on their surface that can tolerate a certain amount of abuse before it becomes a problem. Taking precautions to avoid head crashes, especially not abusing the drive physically, is obviously still common sense. Be especially careful with portable computers; I try to never move the unit while the hard disk is active.

