Disk Geometry: How Data is Organized on a Hard Disk Drive

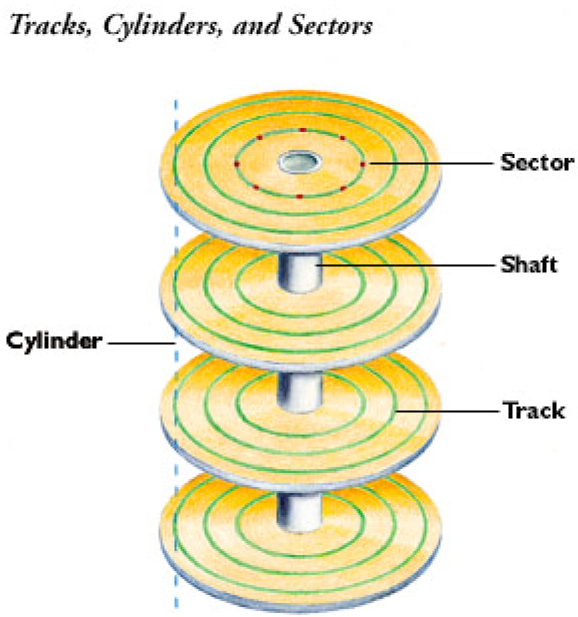

The surface of the drive platter is organized with coordinates, much like a map. Data is stored in concentric tracks on the surfaces of each platter. (A platter has two sides, and thus, two data recording surfaces.) A typical disk drive can have more than 2,000 tracks per inch (TPI) on its recording surface. A cylinder describes the group of all tracks located at a given head position across all platters. To allow for easier access to data, each track is divided into individually addressable sectors.

The process of organizing the disk surface into tracks and sectors is called formatting, and almost all hard disk drives today come preformatted by the manufacturer. The process of formatting a hard drive applies addressing data to the platter's surface. In almost all systems, including PCs and Macintoshes, sectors typically contain 512 bytes of user data plus addressing information used by the drive electronics (although some proprietary systems use other sector lengths). The disk drive controller, which resides on the drive's PCB, uses the formatting information and addresses - much like a tourist uses a city map - to guide data into and out of a specific location on the hard drive. Without formatting instructions, neither the controller nor the operating system would know where to store data or how to retrieve it.

In earlier hard drive designs, the number of sectors per track was fixed and, because the outer tracks on a platter have a larger circumference than the inner tracks, space on the outer tracks was wasted. The number of sectors that would fit on the innermost track constrained the number of sectors per track for the entire platter. However, many of today's advanced drives use a formatting technique called Multiple Zone Recording to pack more data onto the surface of the disk. Multiple Zone Recording allows the number of sectors per track to be adjusted so more sectors are stored on the larger, outer tracks. By dividing the outer tracks into more sectors, data can be packed uniformly throughout the surface of a platter, disk surface is used more efficiently, and higher capacities can be achieved with fewer platters. The number of sectors per track on a typical 3.5-inch disk ranges from 60 to 120 under a Multiple Zone Recording scheme. Not only is effective storage capacity increased by as much as 25 percent with Multiple Zone Recording, but the disk-to-buffer transfer rate also is boosted. With more bytes per track, data in the outer zones is read at a faster rate. Quantum Corporation is a pioneer in Multiple Zone Recording, and was the first manufacturer to implement Multiple Zone Recording on 2.5-inch disk drive products.

Read/Write Heads:Skimming the Surface

Read/write heads are the single most costly component of a hard disk drive, and their characteristics have a great impact on drive design and performance. Despite their expense, the head's basic design and objective are relatively simple: a head is a piece of magnetic material, formed almost in the shape of a "C" with a small opening or gap. A coil of wire is wound around this core to construct an electromagnet. In writing to the disk, current flowing through the coil creates a magnetic field across the gap that magnetizes the disk coating layer under the head. In reading from the disk, the read/write head senses an electronic current pulse through the coil when the gap passes over a flux reversal on the disk.

As technology increases, bits are packed more densely, and the space required to store a bit shrinks. At the same time, the tiny size of the stored data bit causes the signal produced by the head when reading the bit to become weaker and harder to read. As a result, the fundamental challenge in packing bits closer together is finding a way to fly the heads closer to the media to increase the amplitude of the signal. The hard disk drive industry has made great strides on this front. In 1973, flying heights averaged 17 microinches. Today's heads fly at just three microinches, with 2- to 2.5-microinch flying heights expected soon. And, in the not too distant future, read/write heads might even make contact with the media, enabling data to be packed even more densely on the platter surface but offering the additional challenge of eliminating added wear on the disk media and read/write heads. (For more information, see "Alternative Disk Media and Contact Recording" in chapter 4.)

//

Areal Density: The Measure of Disk Capacity

The overall capacity of a drive depends on how densely information (i.e., bits) can be recorded on the disk media. The bits per square inch (BPSI) number is called the areal density. Areal density is calculated by taking the number of bits per inch (BPI) that can be written to and read from each track, and multiplying that number by the number of tracks per inch (TPI) that can be packed onto the disk. BPI depends on the read/write head, recording media, disk rpm, and the speed at which the electronics can accept bits. TPI depends on the read/write head, recording media, the mechanical accuracy with which the head can be positioned on its actuator arm, and the ability of the disk to spin in a perfect circle. Increasing areal density can come by increasing either or both of these factors.

//

Storing and Retrieving Data: The User's View

Let's take a look at what happens when you retrieve data from a hard disk drive. When you issue a command to open an existing file, the application program you're running prompts you to enter the name of the file to open. It then passes the file name to the operating system, which determines where the file is located on the disk drive - the head number, cylinder, and sector identification. The operating system transfers this information to the disk controller, which drives an actuator motor connected to the actuator arm to position the heads over the right track. As the disk rotates, the appropriate head reads the address of each sector on the track. When the desired sector appears under the read/write head, the entire contents of the sector containing the necessary data are read into a special, ultra-fast memory, called cache, on the drive's PCB. Then, the disk drive interface chip sends the necessary information to the computer's main memory.

(For more information on hard drive cache memory, see "Improving Performance with a Cache Buffer" in chapter 5. For more on how hard drives and the host computers in which they reside communicate, see "Making the Connection: Buses and Hard Disk Drive Interfaces" later in this chapter.) Storing data on a hard drive is a similar process to retrieving data, only reversed. The host computer operating system is responsible for remembering the addresses for each file on the disk and which sectors are available for new data. If the file you want to store is large - for example, a 10 MB CAD/CAM drawing - the operating system instructs the controller where to begin writing information to the disk. The controller moves the read/ write heads to the appropriate track and writing begins. When the first track is full, the heads write to the same track on successive platter surfaces. If still more track capacity is required to store all the data, the head moves to the next available track with sufficient contiguous space and writes the data there.

Although an extraordinary amount of care and effort goes into making the platters for hard disk drives, it is not economically feasible to manufacture 100 percent defect-free media. Therefore, all modern hard drives have a defect management strategy built into the disk controller to provide defect-free operation in the field. Defect management involves setting aside some spare sectors on each disk surface to replace a limited number of defective sectors. At the end of the manufacturing process, the entire disk surface is scanned for defects and the disk controller stores a map of their locations. When the operating system requests that information be written to one of the bad sectors, the disk controller transparently maps it to one of the spares. The disk controller continuously updates the defect map, to map out any new bad sectors that might occur.

Tricky Business: Head Positioning

Given the tolerances involved in drive electronics and the speed at which drives operate, the ability to locate data and move the heads accurately to read or write data is an astounding accomplishment. Since the tracks on a platter are located about 300 microinches apart and the heads are flying three microinches above the surface of the disk, accurately positioning the read/write heads is comparable to flying a jet plane one foot off the ground while maintaining a course of flight directly over the center divider of a freeway.

A number of variables work against the process of accurate head positioning. These include temperature variations, which cause minute expansions and contractions of the disk platter, as well as shock and vibration. To counter the effects of these variables and ensure precise head alignment and positioning, hard drives incorporate an electromechanical technique called servo positioning, which provides feedback to the drive electronics to control the position of the head.

There are two primary servo head positioning techniques. Most hard drive manufacturers have abandoned the older method, dedicated servo, for the newer method, embedded servo technology.

- Dedicated servo uses positioning information residing on a single dedicated platter surface accessed by a single dedicated head - a mechanism to which the other heads are slaved. Dedicated servos require one entire side of a platter, a significant percentage of the total disk space - particularly in a hard drive with few platters.
- Embedded servo is emerging as the leading servo implementation for the next several generations of hard disk drive products. Embedded servo overcomes the limitations of dedicated servos by interspersing the servo information with the

data in the form of prerecorded servo burst wedges special data patterns recorded on each of the tracks. When the heads arrive at the intended track location, they read the servo bursts and send back the information to the drive electronics. Then, the drive electronics adjust the position of the actuator motor, which positions the heads so that they receive the maximum signal from the bursts. (The maximum signal only occurs when the head is exactly over the center of the track.) They provide the most accurate, error-free, and cost-effective head positioning technique for small form factor drives.

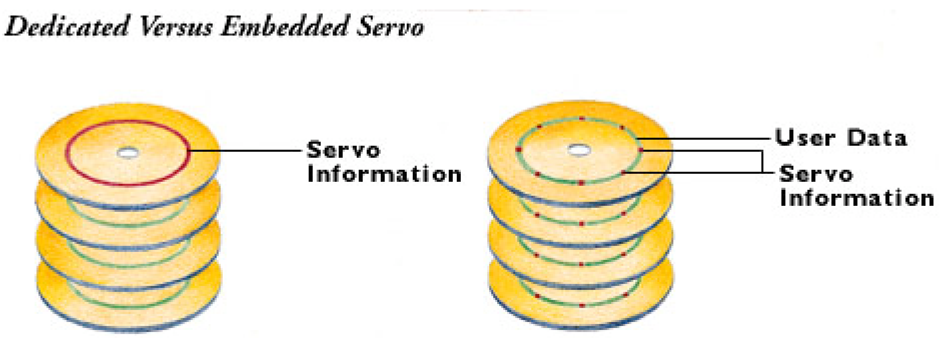

Embedded servo technology for hard disk drive head positioning has been difficult to implement in drives which use Multiple Zone Recording techniques. The varying number of sectors on tracks in different zones vastly complicates the task of reading interspersed servo data. Quantum was the first hard drive manufacturer to effectively address this problem by developing specialized servo feedback and controller ASICs that efficiently handle the task of separating the servo information from the user data.

Making the Connection: Buses and Hard Disk Drive Interfaces

By definition, an interface is anything that allows two separate or dissimilar systems to work together or communicate. One of the most critical but least appreciated parts of a hard disk drive is the interface that allows the drive to communicate with the host system. In peripherals such as hard disk drives, this communication is achieved through adapter boards or chips on the motherboard that plug into the host system's bus. So, to understand drive interfaces, you must first understand a little about the main information corridor inside a computer: the computer bus.
Introducing the Computer Bus Acting much like a human central nervous system, the bus serves as a common conduit for carrying signals to and from various computer components: CPU, video controller, keyboard, storage, and other peripherals. Many types of bus architectures have been offered by manufacturers - both "open" and "closed." Open architectures, most notably the Industry Standard Architecture (ISA) of the early IBM and MS-DOS compatible PCs readily enables outside vendors to design, build, and market general and special purpose add-on cards. Today's main personal computer bus architectures include:

- ISA, a low-cost, relatively unsophisticated bus, available in either 8- or 16-bit versions

- Extended ISA (EISA), a much more sophisticated and powerful 32-bit superset of ISA designed for 386- and 486-based PCs
- The local bus, originally developed for graphics-intensive applications, used now as a high speed disk attachment for IBM-compatible computers
- MicroChannel Architecture (MCA), a proprietary 32-bit bus developed by IBM for its PS/2 line of personal computers
- NuBus, Apple Computer's proprietary 32-bit bus featured in its Macintosh computers
- The Personal Computer Memory Card International Association (PCMCIA) standard 16-bit interface

Efficient use of the bus is a key factor to increasing computer system performance. The speed of data transfer along the bus, which is the product of the bus width and cycle time, has the greatest effect on overall system performance. (The bus cycle time is proportional to the number of words transferred per second. The width determines the width of the transfers.) The bus, which was once 4- or 8-bits wide, typically transferred data at rates of only up to 1 MB/s. Now, the bus is commonly 16- or 32-bits wide and transfers data at speeds of 10, 20, and up to 132 MB/s in Peripheral Components Interface (PCI). The next logical development for the bus is a 64-bit wide interface, which will allow drives and other peripherals to reach even higher data transfer speeds.

(While many of the bus architectures in the list above are commonplace today, two of them, the local bus and PCMCIA, are newer bus technologies for connecting storage devices to the CPU and represent significant technological advancements. Therefore, they are discussed in detail in chapter 4, "Recent Technological Developments").

Understanding Drive Interfaces

A drive interface is required to provide communication between the computer bus and the hard drive. Among the information specified by drive interfaces is how fast the disk and controller should talk to one another, what kinds of commands they can pass back and forth, the location of data and control lines along the connecting cable, and what level of voltage they should use for data transfer. Thus, a drive interface is a standardized combination of connector configuration signal levels and functions, commands, and data transfer protocols.

As with all other high technology developments, drive interface technology has evolved significantly in the past two decades. There are two interfaces commonly in use in today's personal computers. Most IBM-compatible PCs use the Intelligent Disk Electronics (IDE) interface. Also called the AT or ATA (Advanced Technology Attachment) interface, for the PC/AT I/O bus it was designed to work with, IDE can transfer data at a rate of up to 4 MB/s. When connected to a bridge to a faster bus (e.g., local bus), AT can go up to 13 MB/s. The second most common interface is the Small Computer System Interface, or SCSI (pronounced "scuzzy"). This higher performance interface, which is used on Apple Macintosh computers as well as high- performance workstations and servers, can transfer data at a rate of up to 10 MB/s in 8-bit mode and up to 20 MB/s in 16-bit mode.

Quantum was the first company to introduce hard drive products with embedded interface controllers. In the late 1980s, usually drives and interface controllers were sold separately. Today, most hard drives feature an embedded disk drive controller, which includes an interface controller designed onto a single chip and incorporated into the single circuit board that contains the drive electronics. By the early 1990s - in less than a 5-year time span - drive manufacturers have reduced the disk drive controller and drive electronics from two full circuit cards to just five integrated circuit chips that would easily fit in the palm of your hand.