Hard Disk Drive Performance Measures

By Mark E. Donaldson

This paper is just one in a series of 11 designed to increase one's knowledge and understanding of the computer and the computer's hard disk drive. Hopefully this will enable you to improve your hard disk drive's performance. All 11 papers must be used jointly. However, they are written about varying elements of hard drive performance, and each serve well as stand alone descriptive and technical aids. The other ten papers are entitled, ***The Boot Process***, ***The Master Boot Record and Partition Table***, ***CHS Translation***, ***File System Tables***, ***File Systems***, ***Hard Disk Drive Barrier***, ***Disk Fragmentation***, ***Hard Disk Drive Configuration***, ***Hard Drive Optimization and Performance***, and ***Hard Disk Drive Technologies***.

HARD DISK SPEED

The factors that effect the speed of a hard disk drive and the measurements used to describe hard-disk performance are:

- Rotation speed
- Number of sectors per track
- Access Time
- Seek time
- Head switch time/Cylinder switch time
- Rotational latency
- Command Overhead
- Data access time
- Cache on the HD
- How data is organized on the disk
- Transfer rates
- Interface (EIDE or SCSI)

Hard-disk manufacturers measure speed in terms of access time, seek times latency, and transfer rate. These measurements, or benchmarks, often appear in hard-disk advertisements, in comparisons, and on specification sheets.

Here are some basics:

On a harddisk, data is stored in the magnetic coating of the disk. The so called head, held by an actor arm, is used to write and read data. This disk rotates with a constant turn time, measured in revolutions per minute (rpm). Data is organized on a disk in cylinders, tracks and sectors. Cylinders are concentric tracks on the surface of the disk. A track is divided into sectors. A harddisk has a head on each side of a disk. Nowadays, the actuator arm is moved by a servo-motor (not a step-motor which needs more time while swinging in after moving over the desired track). All harddisks have reserved sectors, which are used automatically by the drive logic if there is a defect in the media.

Hard Disk Drive Performance Measures

By Mark E. Donaldson

Rotational Speed

Typical harddisks have a rotation speed from 4,500 to 7,200 rpm, and a 10,000 rpm drive just hit the market. The faster the rotation, the higher the transfer rate, but also the louder and hotter the HD. You may need to cool a 7200 rpm disk with an extra fan, or its life would be much shorter. Modern HDs read all sectors of a track in one turn (Interleave 1:1). The rotation speed is constant.

Sectors Per Track

Modern harddisks use different track sizes. The outer parts of a disk have more space for sectors than the inner parts (zoned recording). Usually, HDs begin to write from the outside to the inside of a disk. Hence, data written or read at the beginning of a HD is accessed and transferred faster rate.

Access time

Access time, a term frequently used in discussions of performance, is the interval of time between the moment a drive receives a request for data, and the moment the drive begins delivering the data. The access time for a hard disk is a combination of three factors: seek time, latency, and command overhead.

Seek Time & Head Switching Time

The fastest seek time occurs when moving from one track directly to the next. The slowest seek time is the so called full-stroke between the outer and inner tracks. Some harddisks (especially SCSI drives) don't execute the seek command correctly. These drives position the head somewhere close to the desired track or leave the head where it was. The seek time everyone is interested in is the ***average seek time***, defined as the time it takes to position the drive's heads for a randomly located request. Yes, you are correct: seek time should be smaller if the disk is smaller (5 ¼ vs 3 ½).

All heads of a harddisk are carried on one actuator arm, so all heads are on the same cylinder. Head switch time measures the average time the drive takes to switch between two of the heads when reading or writing data.. ***Cylinder switch time*** is the average time it takes to move the heads to the next track when reading or writing data. All these times are measured in milliseconds (ms).

Seek time is the time it takes the read/write heads to move from their current location to the track where the desired information is located. Since the desired track could be located on the other side of the platter, or on an adjacent track, the seek time will vary for each individual seek transaction. In actuality, the average seek time for any arbitrary track is equal to the time required to seek across one-third of the tracks. Some drive manufacturers also cite **track-to-track seek time**, which is simply the amount of time to move from one track to an adjacent track. Today's hard drives have track-to-track seek times between three and five milliseconds and average seek times of less than 10 milliseconds.

Hard Disk Drive Performance Measures

By Mark E. Donaldson

Rotational Latency

Each track on a hard drive contains multiple sectors. Once the read/write heads seek the correct track, the heads remain in place and idle until the correct sector passes under them. This waiting time Is called latency. The average latency is equal to the time it takes the platter to make a one-half revolution; for example, on a platter spinning at 3600 RPM, one revolution takes 16.67 milliseconds, so the average latency is 8.3 milliseconds. The latency is Identical on those drives that spin at the same speed. Some of today's hard-drive models have platters that spin at 5500 RPM or more, thus reducing latency.

Command Overhead

The command overhead is the time it takes the controller to process a data request. This includes determining the physical location of the data on the correct platter, directing the actuator to move the head stack assembly to the correct track, reading the data, and forwarding it to the computer. For today's hard drives, the disk overhead is relatively Insignificant.

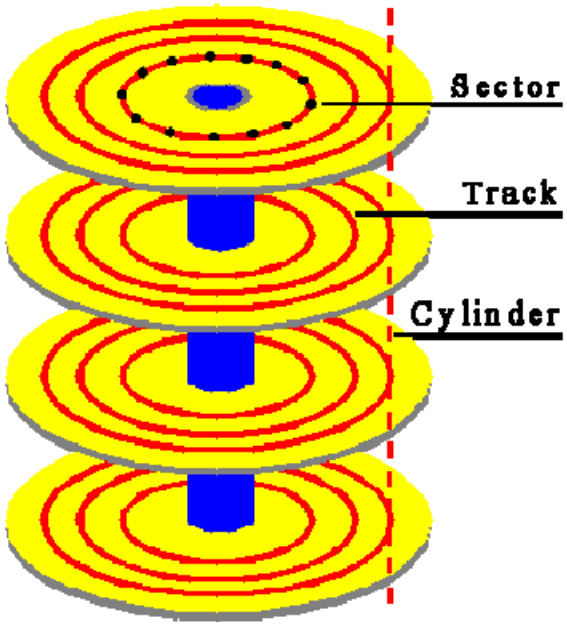

After the head is positioned over the desired track, it has to wait for the right sector. This time is called rotational latency and is measured in ms. The faster the drives spins, the shorter the rotational latency time. The average time is the time the disk needs to turn half way around, usually about 4ms (7200rpm) to 6ms (5400rpm).

Data access time is the combination of seek time, head switch time and rotational latency and is measured in ms. As you now know, the seek time only tells you about how fast the head is positioned over a wanted cylinder. Until data is read or written you will have to add the head switch time for finding the track and also the rotational latency time for finding the wanted sector.

Transfer Rates

Hard disks are also evaluated by their transfer rate, which generally refers to the rate at which data can be read from, or written to, the drive. The speed of the platters, density of the data bits, and access time affect the transfer rate. Transfer rates become particularly important when reading and writing large files. Today's drives have transfer rates ranging between five and thirty Megabits per second.

Most of today's hard drives include a small amount of RAM that is used to cache, or temporarily store, data. Some drive specifications refer to a burst transfer rate, or the speed at which data can be read from, or written to, the cache. The sustained transfer rate more accurately reflects the amount of data that can be accessed from the platters In a given amount of time.

Hard Disk Drive Performance Measures

By Mark E. Donaldson

In the pictures you can see the several ways how data can be stored physically on the harddisk. With a benchmark program that calculates the transfer rate or seek time of the whole harddisk you can see if your drive is using a 'vertical' or a 'horizontal' mapping. Depending on what kind of read/write heads and servo-motors (for positioning the actuator arm) are used it is faster to switch heads or to change tracks.

Traditional harddisks orders their capacity in **vertical mapping**. The data is read/written from one cylinder first, starting at the top track down to the bottom, before the heads are moved to the next cylinder.

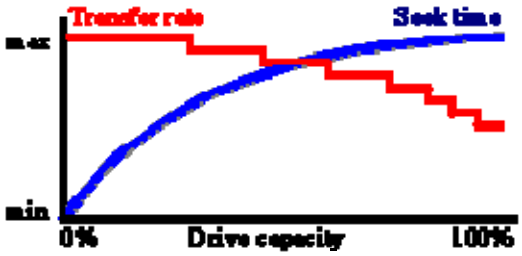

In ***horizontal mapping***, the data is read/written starting from the outer cylinder to the inner cylinder, before switching the heads to the next track.

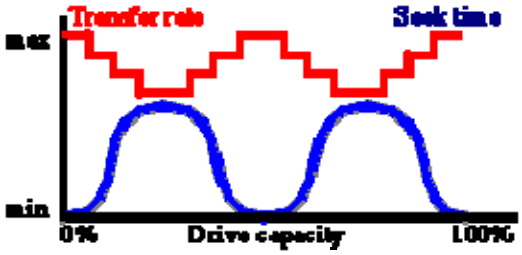

Some harddisks use a combination of vertical and horizontal mapping. As you can see in the below pictures, transfer rate is higher when data is read or written to the outer parts of a disk. The reason is that there is more space for sectors. The number of sectors varies in steps. Usually on a disk there are 10 to 20 zones (called 'notches') with a constant sector number. That's the reason why you see the steps in the transfer rates.

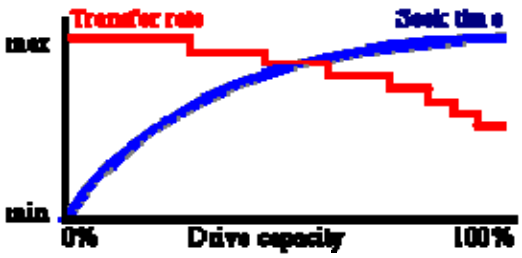

Some harddisk use the combination of 'vertical' and 'horizontal' mapping. The 'horizontal' mapping is used in the zones, the 'vertical' mapping between the zones. However, transfer rate and seek time look the same to 'vertical' mapping. If you are going to buy a HD you have may need to know what kind of mapping the drive does. If you need constant transfer rates (for video, audio) you should get a drive which doesn't do the horizontal mapping. However, drives with horizontal mapping are not very common.

Cache

All modern HDs have their own cache varying in size and organization. The cache is normally used for writing and reading. On SCSI HDs you may have to enable write caching, because often it is disabled by default. This varies from drive to drive. You will have to check the cache status with a program like ftp://ftp.seagate.com/techsuppt/seagate_utils/aspiid15.zip. You may be surprised that it is not the cache size that is important, but the organization of the cache itself (write / read cache or look ahead cache.

With most EIDE drives, the PCs system memory is also used for storing the HDs firmware (e.g. software or BIOS). When the drive powers up, it reads the firmware from special sectors. By doing this, manufacturers save money by eliminating the need for ROM chips, but also give you the ability to easily update your drives BIOS if it is necessary (Like for the WD drives which had problems with some motherboard BIOS' resulting in head crashes!).

Drive Interface (EIDE or SCSI)

Currently there are 2 different interfaces: EIDE and SCSI. You will find an EIDE controllers integrated with the motherboard and that EIDE harddisks are much cheaper than SCSI drives. For SCSI you need an extra controller, because there aren't a lot of motherboards with integrated SCSI controllers. Together with the higher price of a SCSI disk a SCSI system is more expensive than EIDE.

The EIDE interface has a primary and a secondary channel that will connect to two devices each, for a total of four. That could be a harddisk, CD-ROM or disk changers. Lately there have been tape backups with EIDE connectors, but you need special backup software. Scanners and CD-writers aren't available with EIDE interface, only with SCSI.. You can connect up to 7 devices to a SCSI bus or 15 devices to a Wide SCSI. In a standard environment, the performance of single harddisk won't improve much from the SCSI interface. Rather, the power of SCSI is that several devices can use the bus at the same time, not using the bus while they don't need it. So, we see the best benefit from SCSI when several devices are all used on the same bus. The SCSI interface comes in several types: 8-bit (50 wire data cable) or 16-bit (68 wire data cable, Wide SCSI). The clock can be 5 MHz (SCSI 1), 10 MHz (Fast SCSI), 20 MHz (Fast-20 or Ultra SCSI) or 40 MHz (Ultra-2 SCSI).

On one EIDE channel, the 2 devices have to take turns controlling the bus. If there is a harddisk and a CD-ROM on the same channel, the harddisk has to wait until a request to the CD-ROM has finished. Because CD-ROM's are relatively slow, there is a degradation of performance. That's why everybody tells you to connect the CD-ROM to the secondary channel and your

harddisk to the primary. The primary and secondary channels work more or less independently of one another (it's a matter of the EIDE controller chip). The theoretical transfer rate of EIDE is up to 16.6 Mbytes/s in PIO mode 4 or DMA mode 2.

Organization of Data On Disks

You now know, a harddisk has cylinders, heads and sectors. If you look in your BIOS you will find these 3 values listed for each harddisk in your computer. You learned that a harddisk doesn't have a fixed sector size as they had in earlier days. Today, these values are only used for compatibility with DOS, as they have nothing to do with the physical geometry of the drive. The harddisk controller calculates these values into a **logical block address (LBA)** and then this LBA value is converted into the real cylinder, head and sector values. Modern BIOS are able to use LBA, so limitations like the 504 MB barrier are now gone. Cylinder, heads and sectors are still used in DOS environments. SCSI drives have always used LBA to access data on the harddisk. Modern operating systems access data via LBA directly without using the BIOS.