CS 320
Ch 6 – External Memory

Figure 6.1 shows a typical read/write head on a magnetic disk system. **Read and write heads separate. Read head uses a material that changes resistance in response to a magnetic field. This allows a faster read time.**

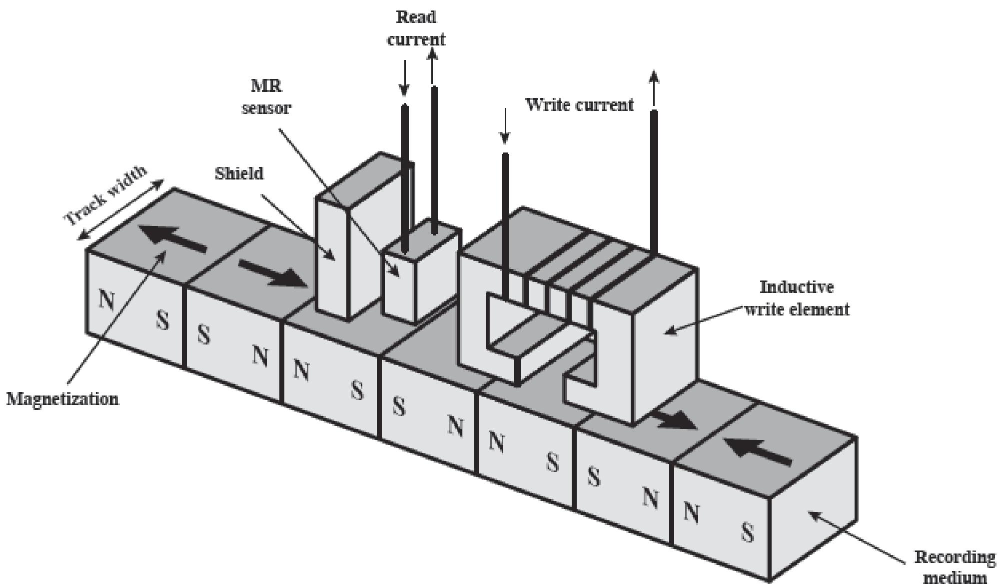

Figure 6.1 Inductive Write/Magnetoresistive Read Head

A track on a magnetic disk is a c**oncentric rings where data is stored.**

A gap is used to s**eparate tracks to prevent data bleeding from one to the other.**

A sector is a p**ie shaped wedge around a disk. Each sector has many track pieces.**

An intersector gap is used to s**eparate the sectors.**

Sectors are of **fixed size of 512 bytes per sector.**

Note that the spacing between bits along an inside track is closer together than the spacing between bits on the outside track.

Disk are typically read using *constant angular velocity*. This means **that the disk turns at a fixed rate (degrees per second is constant). Since bits on the outside are spaced further apart but are moving faster they come off the disk at the same rate as bits on the inside which are denser and moving slower.**

The disadvantage of CAV is that s**ome disk space is lost because of extra space in outer tracks.**

Multiple zone recording is more efficient. **Sectors are divided into zones and each zone has its own bit density. Thus as you move from an inner zone to an outer zone the bit density remains piecewise constant. To compensate the disk controller reads different zones at different rates with data coming off the outer tracks faster than data coming off the inner tracks.**

Sectors distinguished from one another by the disk controller by **writing the sector numbers at the beginning of each sector on each track. This is what lays out the sectors. This is called soft disk formatting.**

Figure 6.4 illustrates disk formatting. There are **30 fixed length sectors per track with 600 bytes per sector. 512 bytes are used for the data field and the remaining bytes are used for formatting and gaps.**

A) The data field consists of a synch byte, data bytes, and a crc. The **Sync byte is used to synchronize clock, the CRC is cyclic redundancy check bits.**

CRC bits

Suppose I want to transmit the message m = {1, 4, 5, 7, 4}. I add these up to get 21. I take 21 mod 8 to get 5. I then subtract 8 – 5 to get 3. I modify m to m = {1, 4, 5, 7, 4, 3}. The receiver adds this up to get 24 and 24 mod 8 is 0 so there are no errors. If any number is changed I will very likely not get 0 for a result indicating that the message is in error. This is a check sum.

A CRC is similar but instead of adding the bits in the message we put them into a feedback shift register

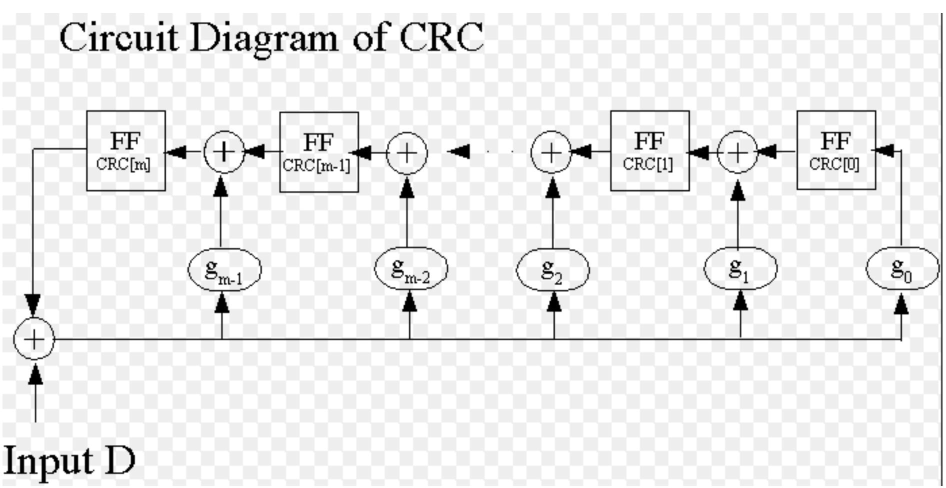

B) The ID field has the sector number in addition to the head number. Most disk systems have **Multiple heads or multiple platters.**

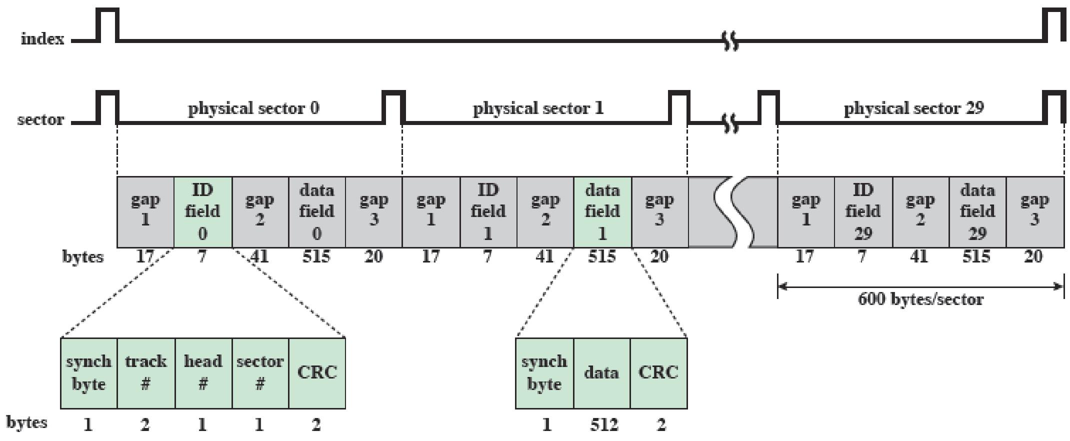

Figure 6.4 Winchester Disk Format (Seagate ST506)

Physical characteristics:

A) Head motion may be **Fixed and movable. Fixed is one per track (faster) and movable is usually one per surface.**

B) Disk portability refers to whether or not the disk is n**onremovable and removable (floppy or zip) disks.**

C) Disks may be single and double sided. **Some systems use both sides of a platter.**

D) A platter is a single disk. **Some disks have multiple spinning disks (platters)**

E) There are three types of head mechanisms. **Contact (floppy), fixed gap, and aerodynamic (Winchester).**

F) The Winchester head was invented by **IBM.**

A cylinder consists of all of t**he tracks of the same diameter on multiple platters.**

A typical hard disk turns at **7200 rpm. Some are faster but this is a common number.**

The *seek time* is the **time it takes the head to move from where it is to the selected track.**

The seek time consists of the s**tartup time and traversal time. The traversal time has a start time and a settling time of its own. This is mostly due to the mass of the head and the arm.**

The *rotational delay or rotational latency is* **the time beginning at the point where the head is at the track until the correct sector rotates around to be at the head.**

A) For a disk rotating at 7200 rpm you can calculate the average rotational latency as follows: **60sec/7200 rotations = 8.3msec per rotation. If information, on average is half way around the average rotational latency is 4.15 msec.**

The disk access time is defined as the **Seek time + rotational latency.**

The transfer time is t**he time it takes to actually read the data after the access time.**

The transfer time calculated as follows: **Trans Time = number of bytes/(rev/sec * number of bytes/track). Thus the transfer time is dependent on how fast the disk rotates since the electronics of the transfer is on the order of a thousand times faster than this.**

Access time of a disk is mostly governed by t**he time it takes to move the head and waiting for the rotational delay. Thus if data is randomly distributed rather than sequentially distributed the disk access time can be substantially slower.**

****** RAID ******

RAID is **Redundant Array of Independent Disks. This is an industry standard multidisk database design.**

RAID is seen by the computer operating system as a single disk system **Even though RAID is a set of disks.**

RAID systems are classified by the RAID level as shown in the table below.

Table 6.3 RAID Levels

Category	**Level**	**Description**	**Disks Required**	**Data Availability**	**Large I/O Data Transfer Capacity**	**Small I/O Request Rate**
Striping	0	Nonredundant	N	Lower than single disk	Very high	Very high for both read and write
Mirroring	1	Mirrored	$2N$	Higher than RAID 2, 3, 4, or 5; lower than RAID 6	Higher than single disk for read; similar to single disk for write	Up to twice that of a single disk for read; similar to single disk for write
Parallel access	2	Redundant via Hamming code	$N + m$	Much higher than single disk; comparable to RAID 3, 4, or 5	Highest of all listed alternatives	Approximately twice that of a single disk
	3	Bit-interleaved parity	$N + 1$	Much higher than single disk; comparable to RAID 2, 4, or 5	Highest of all listed alternatives	Approximately twice that of a single disk
Independent access	4	Block-interleaved parity	$N + 1$	Much higher than single disk; comparable to RAID 2, 3, or 5	Similar to RAID 0 for read; significantly lower than single disk for write	Similar to RAID 0 for read; significantly lower than single disk for write
	5	Block-interleaved distributed parity	$N + 1$	Much higher than single disk; comparable to RAID 2, 3, or 4	Similar to RAID 0 for read; lower than single disk for write	Similar to RAID 0 for read; generally lower than single disk for write
	6	Block-interleaved dual distributed parity	$N + 2$	Highest of all listed alternatives	Similar to RAID 0 for read; lower than RAID 5 for write	Similar to RAID 0 for read; significantly lower than RAID 5 for write

N = number of data disks; m proportional to log N

In *striping* **the data is divided into strips and the strips are distributed round robin across the multiple disks. A strip is a collection of stripes.**

The advantage of striping is that **for slow systems multiple disks can be accessed simultaneously or overlapped to get strips in parallel. This is non-redundant and is faster but provides little protection against data loss. Raid 0.**

Mirroring is the s**imple duplication of data on multiple disks. This redundancy provides some guarantee against loss of data. Raid 1 uses striping and mirroring.**

Raid 2 uses redundancy via the Hamming code. **Error correcting Hamming codes are used with each data strip. Codes are stored on separate disks. This is expensive and not used.**

Bit-interleaved parity is u**sed in Raid 3. Has a single disk to hold all parity bits. If a disk fails the data can be reconstructed. Cheaper than Raid 2. Parity is used on a small set of bits.**

Block-interleaved parity means that p**arity is computed across blocks with parity bits stored on interleaved strips all on one disk. Raid 4 uses this. Raid 5 is similar but has parity distributed across disks instead of stored on a single disk.**

Raid 6 u**ses two different check schemes and makes a reliable system even if two disk systems fail simultaneously.**

****** Solid State Disks ******
Most Solid State Disks use NAND flash.

The advantages of SSDs over HDDs include b**etter performance, less susceptible to physical abuse, longer lifespan, lower power, quieter. About 10 times faster than HDDs**

In terms of cost the **SSD is about $0.50/GByte where and HDD is about $0.15/GByte**

The main problems with SSDs include: **Cost remains higher than HHDs, SSDs can be written about 100,000 times and then becomes unusable, As the disk gets full it can be fragmented, finally NAND flash is written in blocks so a whole block must be processed even if only one byte is needed.**

***** Optical memory *****
Audio CD's and CD-ROM's are Optical memory. The use a l**aser reads pits in a mirror like surface. Pits are ¼ wavelength deep so that the total path is exactly ½**

wavelength thus cancelling the incoming wave and leaving a dark area. Both use *constant linear velocity* in a continuous spiral **track.**

The advantages of CD-ROM's over hard disks are: **CD-ROMS are much more easily duplicated. Optical disk are removable, sturdy, and portable.**

The disadvantage of CD-ROM's over hard disks are a**ccess time is slow ≈ 0.5 seconds and disks are read only.**

CD-R disks are **CD Recordable. These are write once read many CD's. They use a technology similar to CD-ROMs but use a dye layer that is activated by a high intensity laser.**

CD-Rewritable disks **use a phase change mechanism in which a laser can change the "magic" material from crystalline to amorphous. These can be rewritten > 500,000 times but not forever. Cheaper than hard drives although slower. Make good back up media.**

A DVD is a **Digital Versatile Disk. Will eventually replace tape and CD ROMs. Can be double layered and double sided. Can hold up to 17 Gytes of data. Has a higher density than CD-ROMs because of a higher frequency laser and a shorter wavelength.**

Despite DVD's tape systems remain in use because they **have very high capacities (>100 Gbytes) and very cheap.**

There are three differences between DVD and CDs that give DVDs more capacity.**1. Bits are packed more closely because of the laser with a shorter wavelength. 2.DVD has two layers and can read both from the same side. 3. DVD can be two-sided giving two layers on both sides.**