

Chapter 10. Hard Disk Drives

10.1 Introduction

From the early days of the computer to the present, computer storage has been classified into a primary working (RAM) memory which is usually volatile and the non-volatile secondary or backup storage. For secondary storage, paper tapes and cards were used in the early computers, giving way subsequently to magnetic tapes, drums and disks. The pace of development of the magnetic disk drive since its conception in the early 1950's has been such that it is now a standard component in all except the smallest hand held computer system.

The disk drive industry is fast-paced and competitive. With each product generation or cycle, new storage densities are achieved by employing new designs, technologies and materials. Often, these new technologies are introduced by start-up companies, and so we observe that with the supplanting of the older technologies and processes, often the older manufacturers that are slower with the introduction of these technologies become non-competitive and are supplanted also by the newer companies. Thus we have observed frequent "shake-ups" among the disk manufacturers, with many start-up companies, frequent closures, and dynamic re-structuring and mergers and acquisitions taking place.

10.1.1 Disk Drive History

The first magnetic drum was probably used in the Manchester University's Mark I built in 1948. In the early 1950's, IBM conceived and designed the first magnetic disk drive, a Direct Access Storage Device (DASD) in IBM terminology. The pace of increments in the storage density has been rapid as can be seen from the above table.

In this chapter we will focus on Winchester disk drives, currently the most important disk drive technology. These drives are characterised by having non-removable disk media in an environmentally-sealed enclosure. The read-write heads are mounted on light weight flexures and "flies" aerodynamically several μ -in away from the media surface. There is no contact between the head and the media.

There are two stories behind the name Winchester disks; one is that the disk was developed at IBM's facility at *Winchester*, New York State; the other is that the first model number was given as 3030, which is also the model number of the well-known *Winchester* Rifle popular in the Wild West.

10.2 Drive Construction

An exploded view of a typical Winchester disk is shown in Figure 0-1. The low flying height of the head over the media surface necessitates a tight control of the operating environment.

Any dust particles, finger prints and other contaminants can cause the head to “crash” on the media surface resulting in damage to the head and media and the possible loss of valuable stored information.




Figure 0-1. Exploded view of components of the Seagate ST-212 drive.

Figure 0-2 illustrates the relative dimensions of some possible contaminants compared to the flying height of the disk. A sealed enclosure is used to house the media and the head assembly. In addition, the air within the enclosure is cycled through the a fine filter as shown in Figure 0-3Figure 0-1. During manufacture, the drives are assembled at laminar flow benches in clean room environments.




Figure 0-2. Contamination vs. flying height




Figure 0-3. Airflow and filter system in the Seagate ST-212 disk.

10.2.1 Recording Media





Figure 0-4. Track / sector organisation in the recording disk media.

Each drive will have one or more disk platters, each with two magnetic surfaces. The substrate of the earlier disk used aluminium. These were more sensitive to thermal expansion. The magnetic material is applied as a thin coating on the surface. Current drives usually have a glass disk as the substrate and the magnetic material plated or sputtered. Attention is given to


the manufacturing process to ensure a very smooth surface. Figure 0-4 shows the track layout of the media.

10.2.2 Winchester Slider

The Winchester sliders are the carriers built to lift the head micro-inches above the media. Air passing under the air bearing surfaces (ABS) provide the required “lift” and their design has to take into account the weight, velocity, and skew to achieve an uniformed flying height.



IBM 50% TRI-RAIL



MERCURY 3

Figure 0-5. Modern Winchester sliders (flying heads).

10.3 Trends

We first look at some of the trends that are taking place in the industry.

10.3.1 Capacity




Figure 0-6 Forecast of the growth of disk capacity for a single disk.

Each step in the increase in capacity has been a result of increment in the areal storage density. The improvements in areal density also enables the reduction in the physical size of

the drives. When the 3.5 in disk was introduced, its capacity was 40 MB. The 2.5 in drive would not be acceptable by the market until it could also store 40 MB. In recent years, the pace of capacity increases has accelerated and 2 GB disks are commonly available by the end of 1995. From these curves, the conclusion can be drawn that the typical product life cycle is very short indeed!

For a review of the growth of Areal density, refer to the attached IEEE paper: Edward Grochowski and David A. Thompson, “*Outlook for Maintaining Areal Density Growth in Magnetic Recording.*” in IEEE transactions on Magnetics, Vol. 30, No. 6, Nov, 1994.

10.3.2 Form Factor

The introduction of the 3.5 in. floppy disk encouraged the adoption of the 3.5 in. hard disk. As lap-top and notebook computers gained in popularity, the market demanded smaller and lighter drives for these systems. Subnotebooks generally use a 2.5 in. or 1.8 in. drives and this trend is expected to continue.




Figure 0-7. The shrinking form factor - result of market pressure.

10.3.3 Recording Density

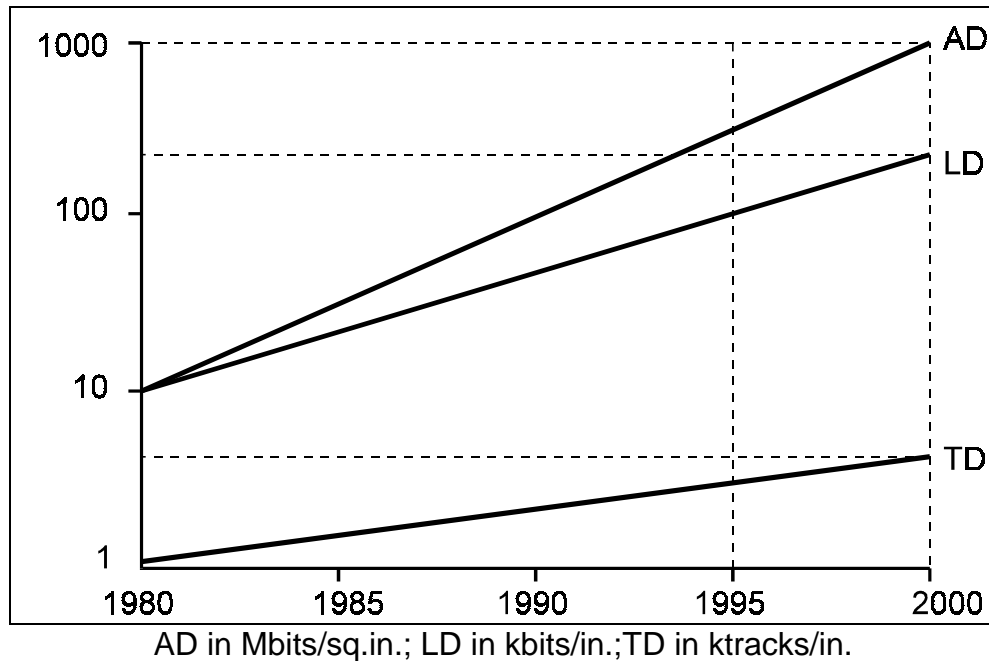


Figure 0-8. Recording Density Growth Forecast

The industry has set itself a target of achieving an areal density of 1Gbits/sq.in. by the year 2000. The curves in Figure 0-8 shows that improvements in both linear recording density and track density are required.

To achieve these targets, advances in many related technologies are required. Work is progressing in the development of sputtered thin film media to increase the remanence and coercivity and reduce the noise of the magnetic material. The finishing of the surface has to be very smooth and hard so that heads can fly lower and yet survive head crashes.

Magnetoresistive (MR) head technology and Partial Response Maximum Likelihood (PRML) read channels are the two most significant solutions now being employed to increase areal density (bits of data per square inch) and boost performance. Alone, each delivers substantial improvements in certain areas over traditional drive technologies, such as inductive heads and peak detection read channels. Together, they reduce the need for many of the capacity and performance trade-offs inherent to disk drive design, while accelerating the decrease in costs per MB.

MR heads and PRML read channels were first used together in 1990 in large-scale storage systems from IBM.

Before looking at the synergy of MR heads and PRML read channels (which will be covered in next chapter), it's important to understand the advantages each technology contributes by itself.

Presently, thin film heads are used, but the transition is being made to MIG (metal-in-gap) and MR (magneto-resistive) heads with separate read and write elements. Having separate heads allows the construction of wider write and narrower read heads to suit the differing operating needs.

10.3.4 Magnetoresistive (MR) Heads

The most economical and practical method for increasing hard disk drive capacities is to increase areal density — fit more bits of data onto the surface of the disk, as opposed to adding disks and heads to the drive. But as density increases, the bit patterns recorded on the disk necessarily grow smaller. This weakens the signal generated by traditional inductive technology read heads, making it difficult to properly identify the patterns. Several methods have been used to combat this. For example, the head can be made to fly closer to the surface of the disk or the disk can be made to spin faster to increase the strength of the signal. Turns, or coils of thin copper conductors, can also be added around the head to boost the read signal (which increases proportionally to the number of turns).

Each of these solutions, however, has its drawbacks. Flying the head closer increases the risk of crashes. Speeding the disk strengthens the signal, but also increases data frequencies; and today's inductive heads cannot perform at very high frequencies. Meanwhile, adding turns helps with the read process but hinders the write process by limiting the frequency with which current reversal can occur for write operations.

MR heads, on the other hand, employ independent read and write elements — using an inductive element (with few turns relative to inductive heads) for write operations, and an independent magnetoresistive element for read operations. The separate read element can also be made narrower to better read tightly spaced data tracks, thus sidestepping the dangers of misalignment. MR heads also produce a strong signal when reading extremely closely spaced bits, regardless of linear disk speed. This means that disks do not have to spin faster in order to accommodate increased density. But, if they do (in order to maximize data rates in high-performance drives), the write head can be optimized for high-frequency write operations without degrading readback performance.




Figure 0-9 A conceptual cutaway view of a magnetoresistive read head.

10.3.5 Lighter and yet stronger flexures

Lighter and yet stronger flexures and heads are required to enable the lower flying height, a higher head-media relative velocity, and yet maintain the robustness and stability to cope with the working environment found in portable and notebook computers.




Figure 0-10 Evolution of sliders

The high track density requirements are to be achieved with improvements in the VCM (voice-coil motor) design, servo control electronics, and design of the embedded or dedicated servo information written on the disks.

10.3.6 Performance metrics

A number of parameters are used to measure and characterise the performance of the disk drive.

Table 1 Definitions of some disk performance parameters.

Access time:	i. Average latency time + time for a random seek, or ii. Average latency time + time taken for 1/3 of full stroke seek (Estimate).
Seek time	Time to move the Read/Write Head from current position to the desired track location
Random seek time	Time to move from any random track to another random track at any time.
Full stroke:	Head movement from track 0 to track N-1 or vice versa, where N is the total number of tracks
Average latency:	Time taken for 1/2 of a disk revolution

Like the floppy disks, stepper motors were initially used to position the head over the required track. A closed-loop system is used as the track spacing decreased. Voice-coil motors (VCM) have been used in the large disk drives with removable disk packs. VCMs, which are more expensive, were introduced into Winchester disks in response to the demand for faster access times.

The access time in current products have almost equal contribution from the latency time and the seek time. Spindle speed is being increased to around 7,200 rpm to lower the average latency. VCM motor design using stronger permanent magnets, DSP-based servo positioning circuits and lighter flexures and heads are directions of development taken to reduce seek times. Basic servo theory indicates that more powerful motors will improve response but this has to be balanced with the design to minimise power consumption.

Table 2. Improvements in average access times

Access time (ms)	Year	Comments
85	1980	
65	1984	↑ stepper (1/3 stroke)
28	1987	stepper (random)
20	1989	↓ voice coil (random)
15	1991	
<10	1995	Increase rotational speed

10.3.7 Power Consumption & Management

Table 3. Improvements in power consumption of disk drives.

Power	Year	Comments
25W	1980	
20W	1984	
15W	1987	
10W	1989	↑ 12 V and 5 V operation
2W	1991	5 V only
1.3W	1992	3.3 V

10.3.7.1 Power Management

Multiple power down modes using CMOS IC's:

Table 4. Power-down modes of current disk drives.

Normal	Spindle running
Idle	Spindle running but actuator circuit off
Standby	Spindle off and actuator parked
Sleep	All switches off and processor waiting for interrupt

More efficient spindle motors and actuators .

10.3.8 Head Media Interface Issues

10.3.8.1 Flying Height

Over the years the flying height has been lowered consistently, with the requirement that there should be no head media contact during operation that can result in damage.

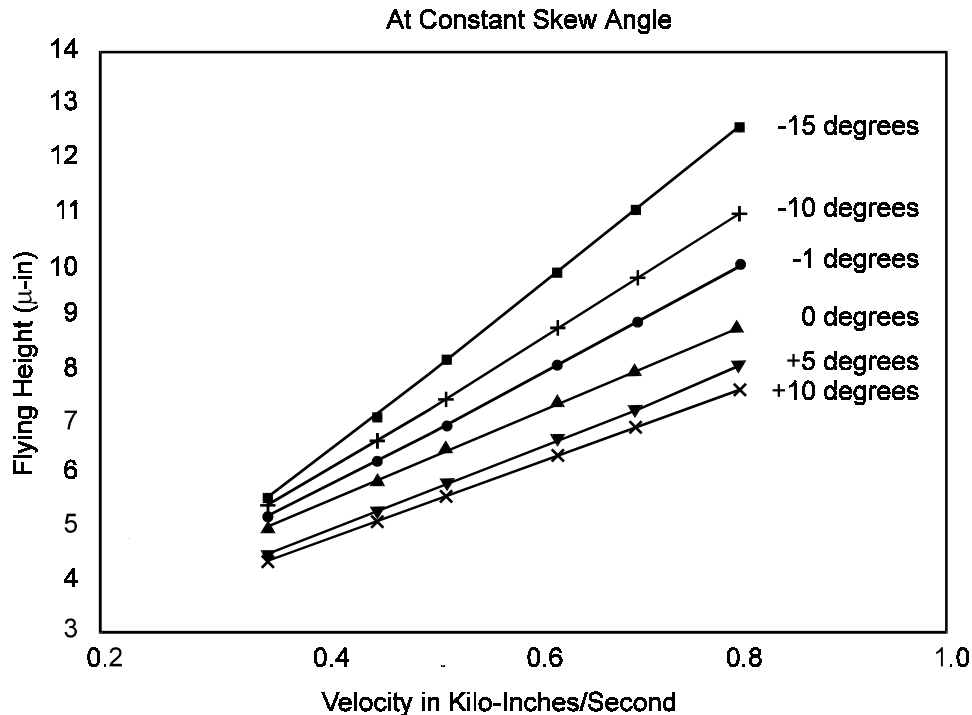


Figure 0-11. Flying height vs. head velocity for different skew angles.

Some of the factors that determine flying height are the relative media velocity. This varies with the track radius. The gram load (weight) of the flexure, and the design and width of the air bearing surface (ABS). As the head moves from the outermost track to in the inmost track, apart from the decrease in media velocity, the skew angle, which measurements the alignment of the head to the track also changes, Figure 0-11 shows the variation in flying height due to these factors.

10.3.8.2 Contact Start Stop (CSS)

The head media combination has to withstand 40,000 starts on a single track. This is especially important for lap-tops and notebook computers with advanced power management where the disk may be shut down during periods of inactivity.

Previously CSS was only 10,000 times.

10.3.8.3 Stiction

If two objects with very flat and smooth surfaces are placed in contact, stiction, which is the bonds set up between the atoms or molecules, will hold the pieces together. When the head rests on the smooth media surface, the stiction force holding the head to the media is not allowed to exceed a pre-determined value. This force should not cause the head to break loose nor cause damage to head or media.

Stiction depends on the smoothness of the media and the presence of any contamination. The problem can be noticed when the head is park on one track (especially ID) for too long. Sometimes the starting torque of the spindle motor is too low.

Normally a breaking force of < 3gm is allowed.

10.3.9 Head Disk Assembly (HDA) Parameters

A number of parameters are used to characterise the performance of the HDA, the head disk combination.

10.3.9.1 Read Amplitude

- i. DC erase the test tracks, normally Track 0 and Track last.
- ii. Using standard write current, record data on whole of test tracks.
- iii. Measure TAA (track average amplitude) read for 1 revolution at 1F and 2F, in MFM; 1F = 1.25 MHz, 2F = 2.5 MHz.

10.3.9.2 Resolution R

The resolution of the HDA is defined as:

$$R = \frac{TAA(2F)}{TAA(1F)} \times 100\%$$

where $70\% < R < 90\%$ is acceptable range for HDA.

10.3.9.3 Media Signal-to-Noise Ratio

Media noise is the total noise less the contribution of noise attributable to the systems' electronics circuits.

$$N_{\text{media}}^2 = N_{\text{total}}^2 - N_{\text{elect}}^2$$

$$S/N \text{ ratio} = 20 \log \frac{2F(\text{signal})}{N_{\text{media}}}$$

A media signal-to-noise figure of > 35 dB at 2F is considered acceptable.

10.3.9.4 Overwrite Modulation

Overwrite modulation is a measurement of the residual 1F signal left after the same track has been overwritten for one revolution at with a 2F signal without DC erase.

$$\text{Overwrite (dB)} = 20 \log \frac{\text{Residual 1F TAA}}{\text{Initial 1F TAA}}$$

The nominal acceptable value for overwrite modulation is -30 dB.

10.3.9.5 Optimal Write Current

Digital Recording works in the saturation region of the B-H Curve. What is the optimal write current to use?




Figure 0-12. B-H curve of a hard magnet. ($H_c > 100$ Oe)

The flying height will vary from OD ($\sim 12 \mu\text{in}$) to ID. ($\sim 8 \mu\text{in}$). As the flying height changes, the optimal write current will vary. Instead of having a constantly varying current, the disk is divided into two or more zones. e.g. ST225 (20 MB 5.25" has 4 zones. These zones should not be confused with the zones in ZCAV (zoned constant angular velocity) recording formats.

The optimal write current for 5.25" drives are:

$$\text{OD: } I_w = 52 - 54 \text{ mA}$$

$$\text{ID: } I_w = 38 - 40 \text{ mA}$$

10.3.9.6 Parameter For Evaluating HDA

In the design selection of the HDA, compromises and trade-offs have to be made between TAA, Resolution, Overwrite modulation, etc. The "bottom line" is to select a combination of head, media and R/W electronics such that we are able to reliably detect the encoded data inside the "timing window".

10.3.9.7 Timing Window

To be decoded correctly, the read pulse (or the corresponding transition) has to occur while the timing window is open. Bit shifts and jitter great enough to move the transition outside the window may cause errors.

10.3.10 Interleave

System hardware and software may not be fast enough to read or write a series of sectors consecutively. This is especially in the case of earlier disk drives. For example, after writing 1 sector, the data for the next sector may not be in the buffer yet, and so cannot be written into the adjacent sector. The disk in this case would interleave the sectors so that sufficient time is given for the data to be ready by skipping the next one or two sectors as seen in Figure 0-13.

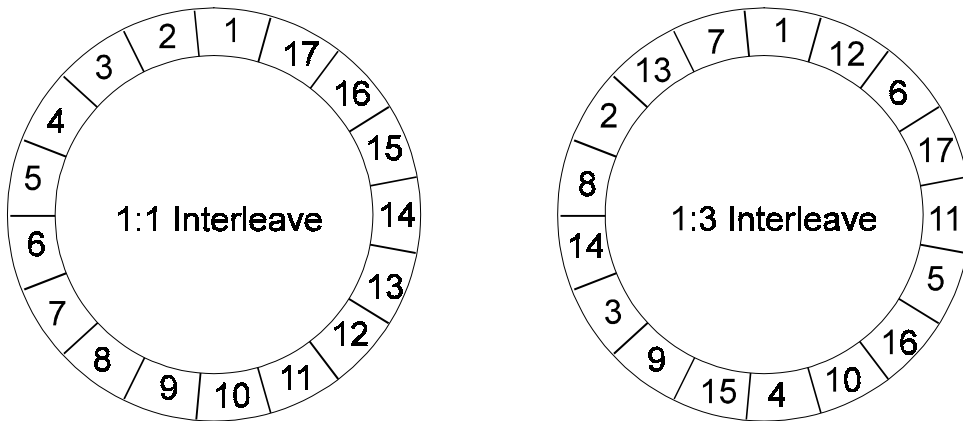


Figure 0-13. Interleave scheme in disk recording.

10.3.11 Cylinder skewing

Once a read/write finishes reading from one track, the head must stepped to another (usually adjacent) track. This stepping process, no matter how rapid does require some finite amount of time. When the head tried to step directly from the end of one track to the beginning of another, the head arrives too late to catch the new track's index. Cylinder skewing technique is intended to improve hard-drive performance by reducing the disk time lost during normal head steps, by offsetting the starting points of each track.




Figure 0-14 Cylinder skewing

10.4 Track Access

10.4.1 Open Loop System




Figure 0-15. Open Loop System - Stepper Motor

The open loop system using a stepper motor are found only in floppy disks where the track density is relatively low, e.g. 96 tpi. is used on the 1.44 MB minidiskette.

Advantages:	Disadvantages:
Low Initial Cost	Poor Tolerance for Track Distortion
Low Upkeep Cost (Servo Data on Media not Required)	Low Track Densities
Simple, Reliable Design	Sensitive to mechanical alignment and temperature variations.

10.4.2 Closed-Loop Servo Systems

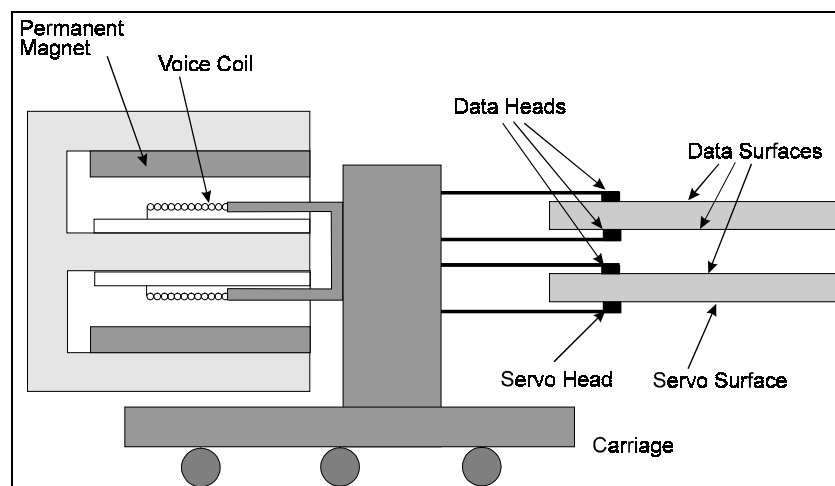


Figure 0-16. Closed-loop voice-coil positioner (ST4096)

Higher track densities requires some kind of feedback to be applied to the positioning motor to locate the head onto the selected track. Lower cost solutions used stepper motors, but for improved performance specially designed linear or voice-coil motors are normally used, as shown in Figure 0-15.

Advantages:	Disadvantages:
Accommodate Higher Track Densities/ Storage Capacities Tolerate Media Distortion Can Offer Faster Positioning	Higher Initial Cost Require Media Containing Servo Data More Complex Design

In the closed-loop system shown in Figure 0-16, the head carriage is actuated by the linear (“voice coil”) motor. The position of the head is sensed by reading the magnetic servo pattern pre-recorded during drive manufacturing. Specially built servo-writers are used to individually write these servo tracks on each production disk.

The actuator, together with the control electronics, senses the servo signals, processing them to determine the off-track error and adjusts the position of the head. Microprocessor- and DSP-based control circuits are used to provide the fast and accurate response needed in high track density drives.

The two types of servo systems used are described below.

10.4.2.1 Embedded Servo System

For drives with just a few media surfaces, the servo information is often embedded into the data tracks. Two basic organisations are used, the wedge servo in which the servo signals are confined to a wedge or sector of the tracks or the servo signals can be distributed throughout the tracks by embedded the signals into every sector. In either case a portion of the surface is not available for data storage. Embedding the servo throughout the tracks gives finer control and is easier on the control electronics. whereas in the wedge servo scheme, the head is basically “running free” outside of the servo data sector.

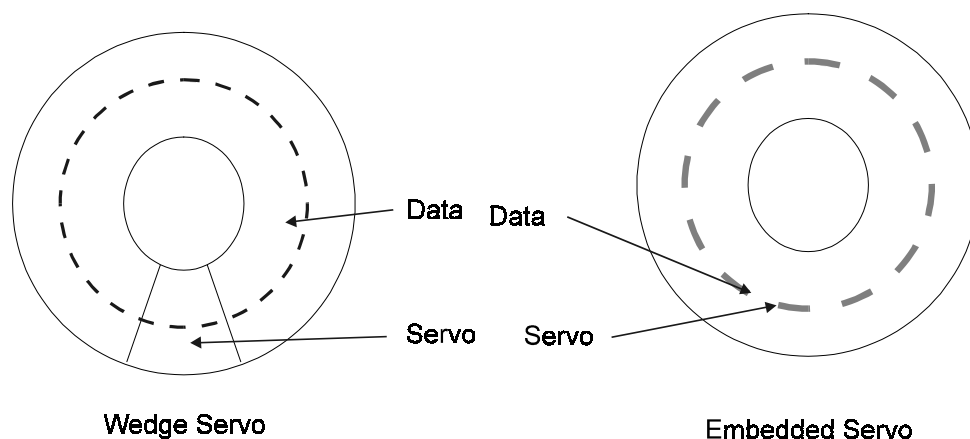


Figure 0-17. Embedded servo tracks in disk media.

Figure 0-18. gives a close-up picture of the embedding of servo data as part of each data record. The servo pairs “A” and “B” are sensed by the read head as it passes over the record. These two signals are compared for balance. Any difference between them is used as the feedback error signal to correct the position of the head. When signals “A” and “B” balance, the head is said to be “on-track”.

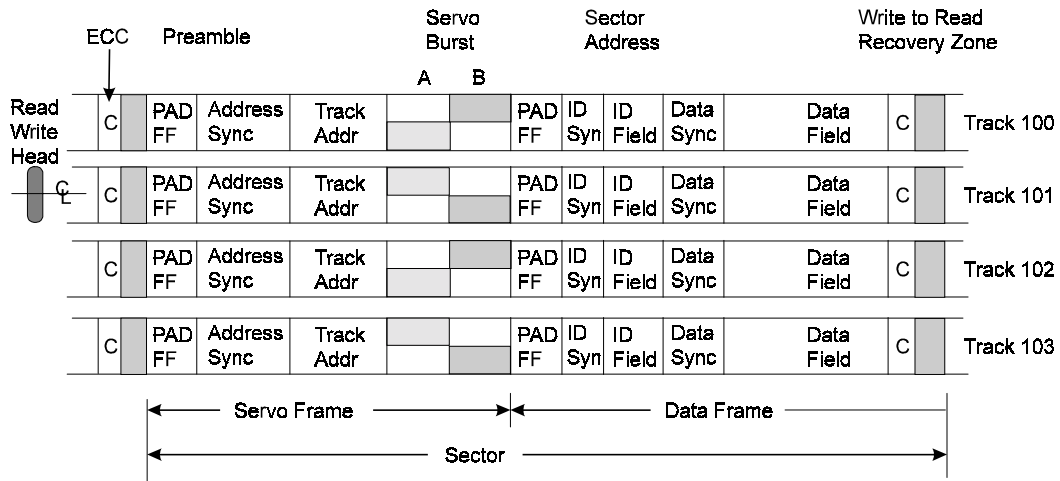


Figure 0-18. Layout of adjacent tracks showing embedded servo bursts.

10.4.2.2 Dedicated Servo System

In this scheme, a complete disk surface is dedicated for servo data. This has the advantage of even faster positioning and more accurate track following due to the high servo sampling available. In turn this enables higher track density. On the converse, a complete media surface is lost to servo data. A separate servo head is required, although this however can be optimised for servo use.

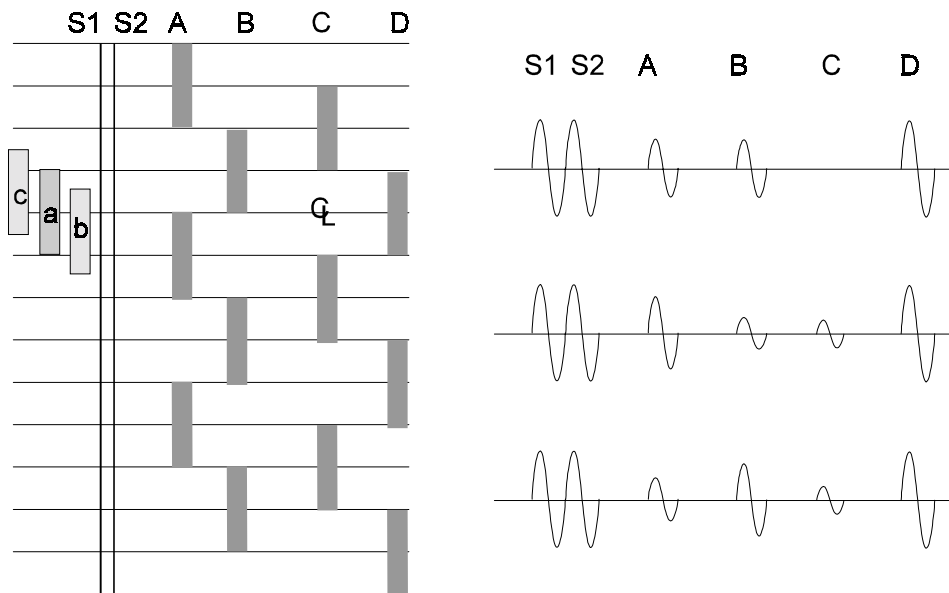


Figure 0-19. Dedicated servo signals.

As shown in Figure 0-19 above, four servo signals are used. In this case, The head is “on-track” when signals A & B balance and the error signal C is zero. Note that when head is completely off-track A & B can balance but D is zero and C is large.

10.5 Position Error Signal (PES)

PES is a signal proportional to the relative difference of the positions of the centre of the servo head and the nearest track centre. Thus the position error signal is a periodic function of x for stationary and ideal track centres. The position error signal contains two sources of motion: Motion of the actuator and; Motion of the disk surface itself.

The pattern used on the servo surface is designed in concert with a demodulation scheme, such that when read back, the signals infer head position relative to the nearest track centre. Two basic types of demodulation are employed: Peak detection and Area detection. Both peak detection and area detection are sensitive to the amplitude of the read back signal from the servo head. Area detection is less sensitive to disk surface defects and noise. Because many parameters affect a wide range of readback signal amplitude, AGC technique is usually employed to prevent unwanted variations in the PES signal gain.

10.5.1 Output signals of a typical position error channel

To keep track of the absolute position of the actuator, cylinder pulses are generated which indicate that a track boundary was traversed. This information is used in the seek operation.




Figure 0-20: signals of a typical position error channel

Output signals of a typical position error channel. (a) Ideal triangular output waveform, the zero crossings of which represent track centres; (b) PES ramps derived from ideal PES waveforms. PES ramps have a slope with a constant sign; (c) cylinder pulses indicate servo head is at half-track point.

10.6 Recording Formats

The layout for a typical sector is shown below.

The location and ID information for each sector is developed when the drive is formatted. After formatting, only data and ECC bytes are updated during writing. If sector ID information is accidentally overwritten or corrupted, the data recorded in the afflicted sector becomes unreadable.

When a drive identifies a location, it generates a CRC code which it compares to the CRC code recorded on the disk. If the two CRC codes match, the address is assumed to be valid, and the disk operation can continue. otherwise, an error has occurred and the entire sector is considered invalid.

Up to 512 bytes (in the case of DOS) can be written or read from the data field. The data is processed to derived eleven bytes of ECC error-checking code using Reed-Solomon encoding.

If data is being read, the derived ECC is compared to the recorded ECC. When the codes match, data is assumed to be valid.

When writing data, the derived ECC will be written on the disk.

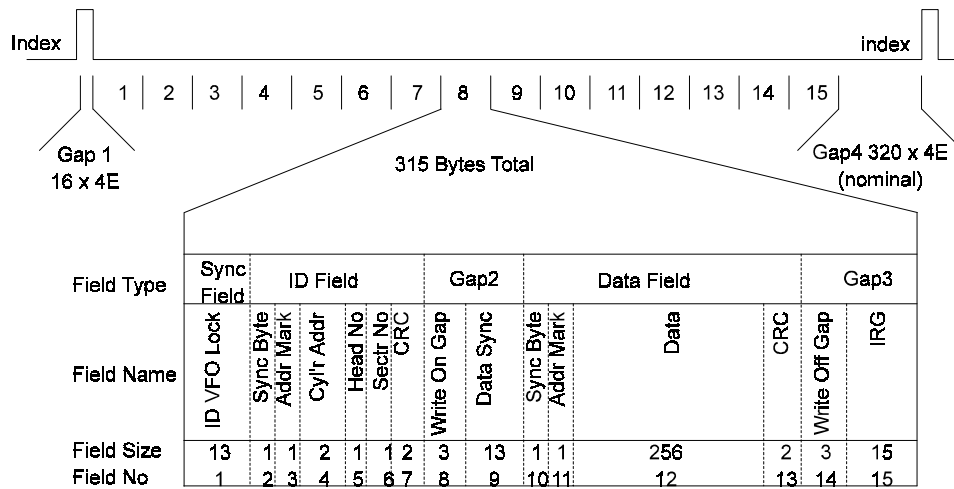


Figure 0-21. 256 Bytes/Sector Format

10.7 Disk Controllers and Interfaces

10.7.1 Drive Electronics

IDE

ESDI

SCSI

IPI

Cache

RAID

10.7.2 Drive Functions

10.7.3 Controller Functions