Fundamental of HDD Technology (3)

Data Storage Technology Research Unit
Nakhon Pathom Rajabhat University

Assistant Prof. Piya Kovintavewat, Ph.D.
E-mail: piya@npru.ac.th
URL: http://home.npru.ac.th/piya

Outline

Head Parking and the Landing Zone

- Hard disks work by having the read/write heads fly over the surface of the disk platters.
- This floating action occurs only when the platters are spinning.
- When the platters are not moving, the air cushion dissipates, and the heads float down to contact the surfaces of the platters $\Rightarrow$ potential for damage
- While the platters and heads are designed with the knowledge in mind that this contact will occur, it still makes sense to avoid having this happen over an area of disk where there is data!

- For this reason, most disks set aside a special track that is designated to be where the heads will be placed for takeoffs and landings.
- Appropriately, this area is called the landing zone, and no data is placed there. The process of moving the heads to this designated area is called head parking.
- IBM has developed an alternative to conventional head parking.

Head Load/Unload Technology

- Instead of letting the heads fall down to the surface of the disk when the disk's motor is stopped, the heads are lifted completely off the surface of the disk while the drive is still spinning, using a special ramp.
- When the power is reapplied to the spindle motor, the process is reversed: the disks spin up, and once they are going fast enough to let the heads fly without contacting the disk surface, the heads are moved off the "ramp" and back onto the surface of the platters.
- IBM calls this load/unload technology.

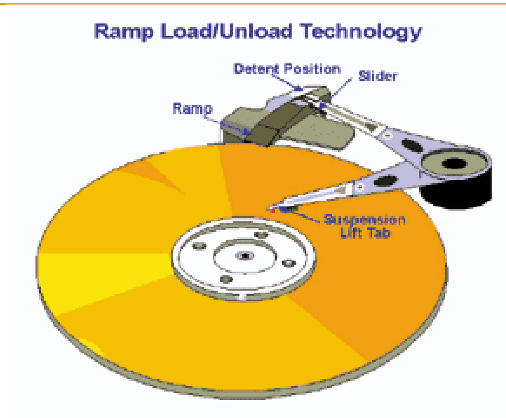

Diagram showing how IBM's load/unload ramp technology functions. (One head and surface is shown; there would be a different ramp for each head and surface the disk has.)

Track, Cylinder, and Sector

- Platter $\Rightarrow$ track $\Rightarrow$ sector (512 bytes)
- **Track**:
  - A concentric set of magnetic bits on the disk is called a track.
  - Each track is divided into 512 bytes (usually) sectors.
- **Sector:**
  - A part of each track defined with magnetic marking and an ID number.
  - Sectors have a sector header and an error correction code (ECC).
  - In modern drives, sectors are numbered sequentially.

Disk drive store over 100,000 bytes in a track --> Too big to deal with
We break each track into chunks called sectors :

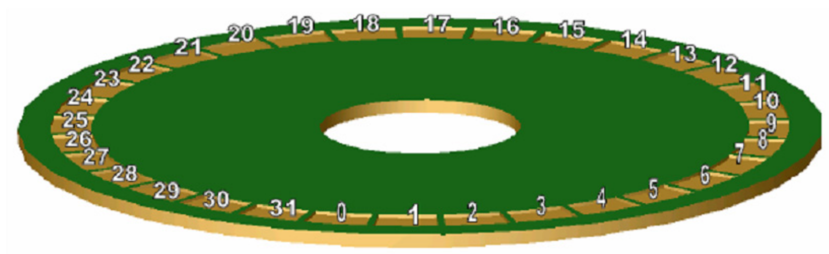

Most common sector Size = 512 Bytes (1024 and 2048 bytes common)
Typical Sectors Per Track = 50 to 256 (determined by bit density)
Breaking tracks into sectors used up some space - Formatting Efficiency (5% - 15%)

Drive Physical and Logical Organization

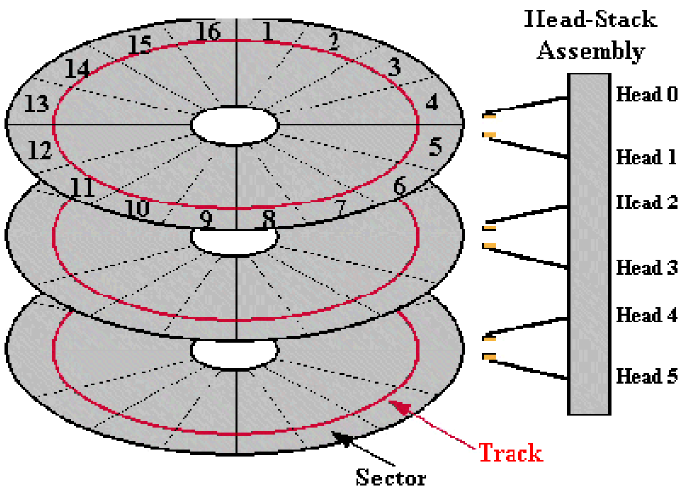

- Each platter uses two heads ⇒ the top and the bottom
- The heads are locked together on an assembly of head arms.
  - All the heads move in and out together
  - Impossible to have one head at track 0 and another at track 1000.
- The track location of the heads is commonly referred to as a cylinder number.
- A cylinder is basically the set of all tracks that all the heads are currently located at.

- **Cylinder:**
  - A group of tracks with the same radius is called a cylinder.
- **Data addressing:**
  - Two methods for Drive's data addressing:
    - CHS (cylinder-head-sector) ⇒ Used on most IDE drives
    - LBA (logical block address) ⇒ Used on SCSI and enhanced IDE drives.

Data Addressing

- CHS mode uses "cylinder, head, sector" to refer to a data block in HDDs
  - No translation done at the BIOS level
  - Limit to 1,024 cylinders, 16 heads and 63 sectors, or 504 MB
- Logical Block Addressing (LBA)
  - Each sector is instead assigned a unique "sector number", from 0, 1, 2, … (N-1)
    - N = number of sectors on the disk
  - Must be supported by BIOS, OS, and HDD
  - All newer hard disks support LBA ⇒ Support high capacity

Zoned Bit Recording

- One way that capacity and speed have been improved on hard disks over time is by improving the utilization of the larger, outer tracks of the disk.
- The first hard disks were rather primitive affairs and their controllers couldn't handle complicated arrangements that changed between tracks.
- As a result, every track had the same number of sectors. The standard for the first hard disks was 17 sectors per track.

- Of course, the tracks are concentric circles, and the ones on the outside of the platter are much larger than the ones on the inside.
- Since there is a constraint on how tight the inner circles can be packed with bits, they were packed as tight as was practically possible given the state of technology, and then the outer circles were set to use the same number of sectors by reducing their bit density.
- This means that the outer tracks were greatly underutilized, because in theory they could hold many more sectors given the same linear bit density limitations.

- To eliminate this wasted space, modern hard disks employ a technique called zoned bit recording (ZBR).
- Tracks are grouped into zones based on their distance from the center of the disk
  - Each zone is assigned a number of sectors per track.
- Normally, the outer tracks have more sectors per track than the inner tracks.
  - Pack more data onto the outer tracks
- To allow for more efficient use of the larger tracks on the outside of the disk.

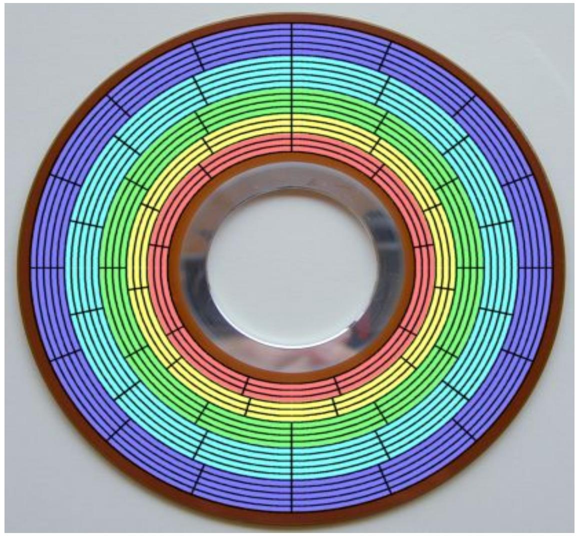

Zone Bit Recording

Zone Table

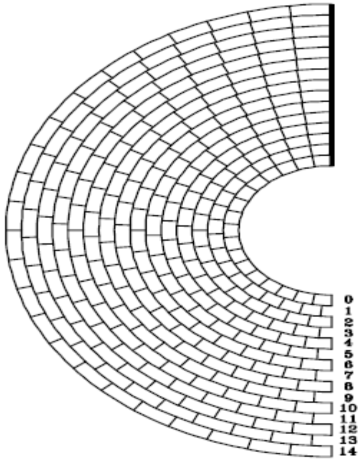

Zone Table

Zone #	Track # Begin		Track # End	# Sect.
0	1050	–	1124	34
1	975	–	1049	35
2	900	–	974	36
3	825	–	899	37
4	750	–	824	38
5	675	–	749	39
6	600	–	674	40
7	525	–	599	41
8	450	–	524	42
9	375	–	449	43
10	300	–	374	44
11	225	–	299	45
12	150	–	224	46
13	75	–	149	47
14	0	–	74	48

ZBR TABLE EXAMPLE

DATA RECORDING ON HARD DISK DRIVE

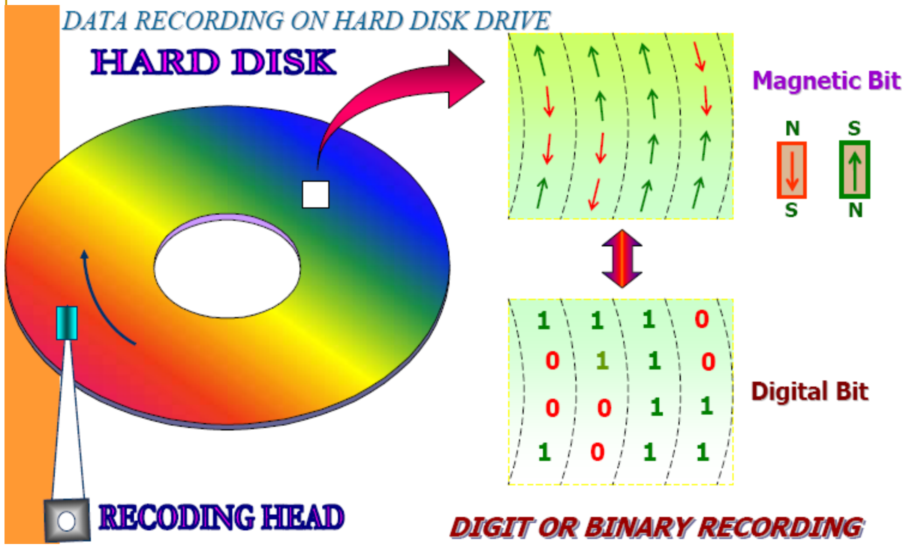

AREAL DENSITY OF RECORDING HEADS

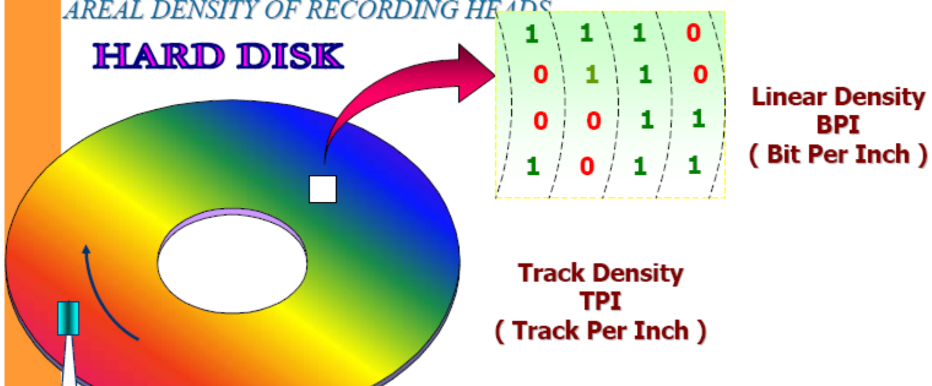

Areal Density = TPI X BPI
(AD) = Bit Per Square Inch
= bpsi

e.g. GMR head AD is about 20-100 Gbpsi

AREAL DENSITY v.s. DATA STORAGE

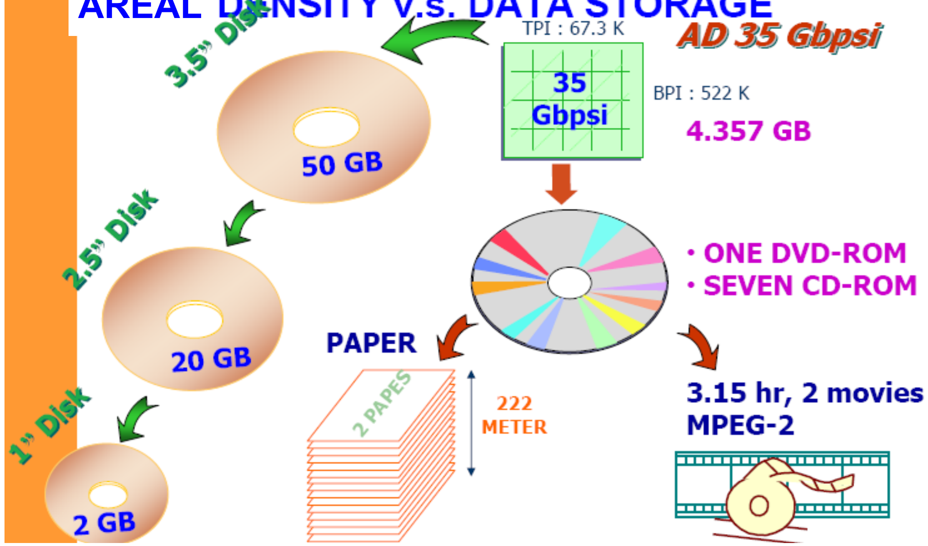

Track

Track = A strip of data written on a magnetic film

Each bit's value is sampled at regular interval:

1	when magnetic transition presents
0	when magnetic transition does not present

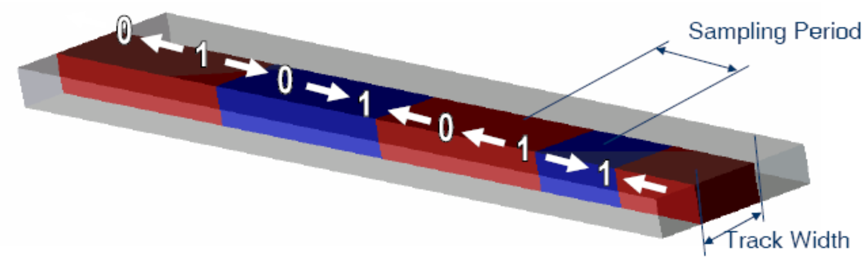

Track Density

Track Density = Number of tracks that fit in one inch (TPI)

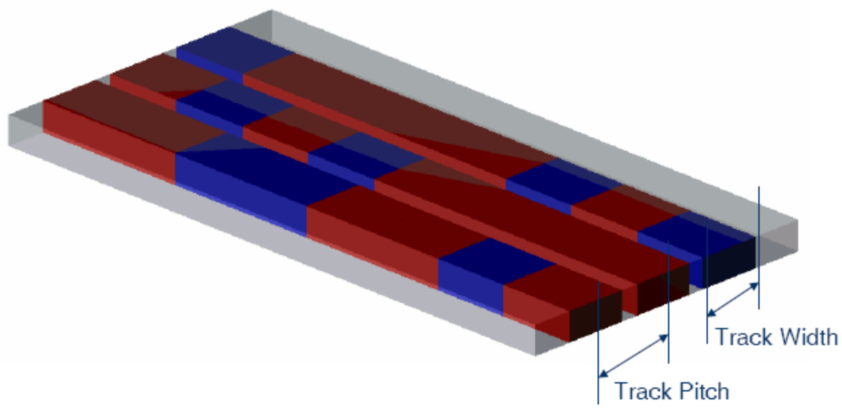

Bit Density (Linear Density)

Bit Density = Number of bits that fit in one inch of track (BPI)

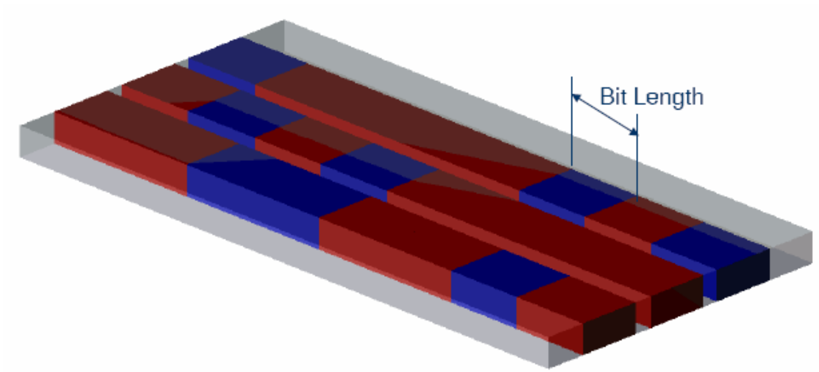

Areal Density

Areal Density = The amount of data that can be stored in 1 square inch

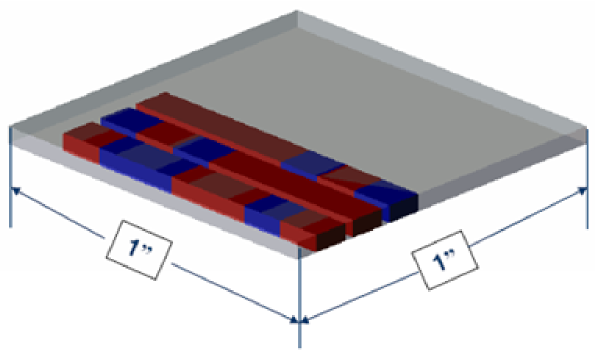

Areal density = BPI x TPI (Gbits/in^2)

Calculate Data Rate

Data Rate = Bit Density x Velocity
[bits/sec] [bits/inch] [inch/sec]

DR = BPI x (RPM/60) x $2\pi r$

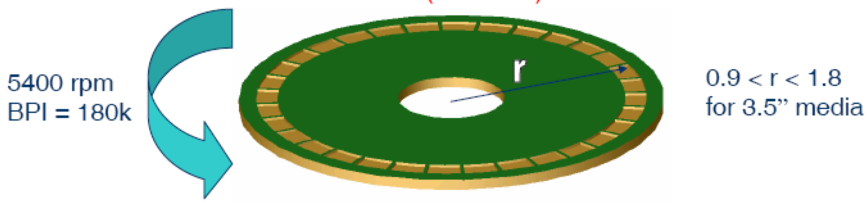

Circumference at outer edge = $2\pi r$ = 11.3 inches

Velocity = Circumference x Rotation Rate

Velocity = 11.3 inches/rev x 5400 rpm x (1/60) min/sec

Velocity = 1017 inches/sec

Data Rate = 180 Kbpi x 1017 inches/sec = 183 Mbps

- Transfer rate As of 2008:
  - The data transfer rate at the inner zone ranges from 44.2 MB/s to 74.5 MB/s
  - The data transfer rate at the outer zone ranges from 74.0 MB/s to 111.4 MB/s
- **Remark:**
  - One interesting side effect of this design (i.e., ZBR) is that the raw data transfer rate (or the media transfer rate) of the disk when reading the outside cylinders is much higher than when reading the inside ones.

- This is because the outer cylinders contain more data, but the angular velocity of the platters is constant regardless of which track is being read.
- Since hard disks are filled from the outside in, the fastest data transfer occurs when the drive is first used.
- Sometimes, people benchmark their disks when new, and then many months later, and are surprised to find that the disk is getting slower!
- In fact, the disk most likely has not changed at all, but the second benchmark may have been run on tracks closer to the middle of the disk.

Increasing Areal Density

- HDD capacity is measured by Areal Density
- To hold more data, need more TPI and BPI
- As capacity increases, bit sizes must decrease
  - More difficult to write and read the magnetic signal
- Increasing the areal density (BPI and/or TPI) is difficult
  - Require many technological advances and changes to various components of HDD

- When the data is packed closer and closer
  - **Cause** ⇒ interference
  - **Solution** ⇒ Reduce the strength of the magnetic signals stored on the disk
    - Must ensure that the signals are stable and the heads are sensitive and close enough to the surface to pick them up
- Every few years a read/write head technology breakthrough enables a significant jump in density.
  - That's why HDD has been doubling in size so frequently.

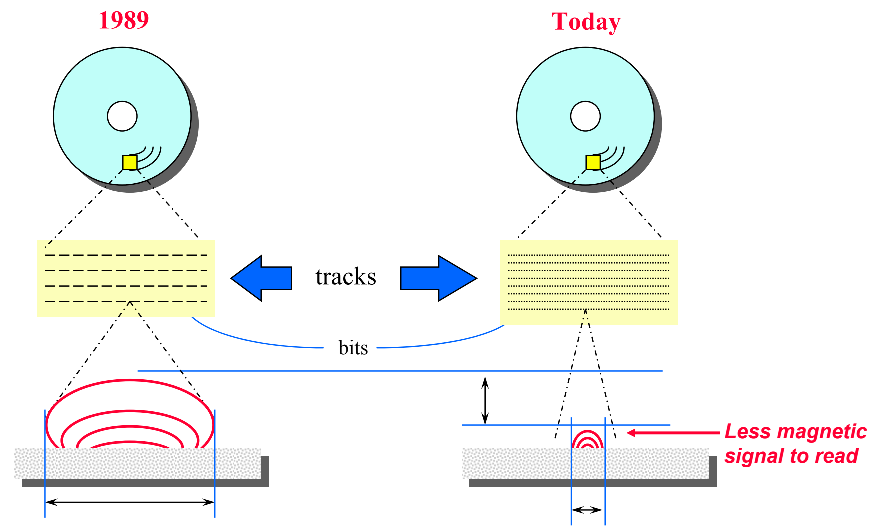

Current Areal Density

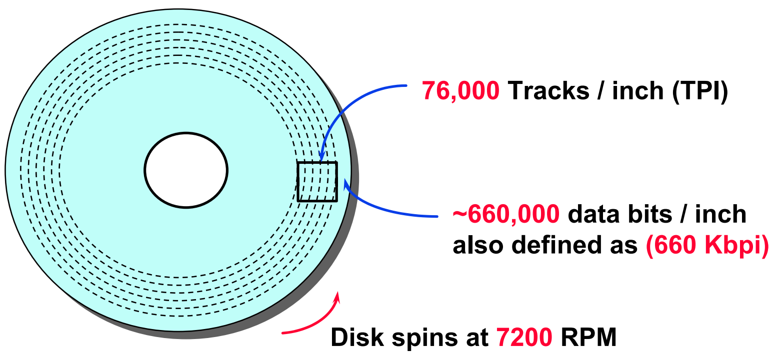

- In 1 inch, the read/write head can differentiate 660,000 data bits
- Also in 1 inch, it can place about 76,000 tracks

Hard Disk Capacity Growth

1990 | 2004

20 MB Technical Specs...

Physical Parameters:

Cylinders	= 782
Heads	= 2
Sectors/track	= 27
Sectors/drive	= 42,228
RPM	= 4500

Data Density:

Bits/inch (BPI)	= 22Kb
Tracks/inch(TPI)	= 1021

250 GB (250,000 MB!)

Physical Parameters:

Cylinders	= 74,686
Heads	= 6
Sectors/track	= 506 – 945
Sectors/drive	= 488,397,168
RPM	= 7200

Data Density:

Bits/inch (BPI)	= 660Kb
Tracks/inch(TPI)	= 76,000

In the same 3.5" form factor, the hard disk can now store >12,500 times more data