Disk Basics

WARNING! This Material NOT in Text

- "*Disk Subsystem*" – Interface to x86 – Responsible for Reading/Writing
- Interface Circuit is "*Disk Controller*"
  - Partially inside Fixed Disk Housing and on Motherboard/Card

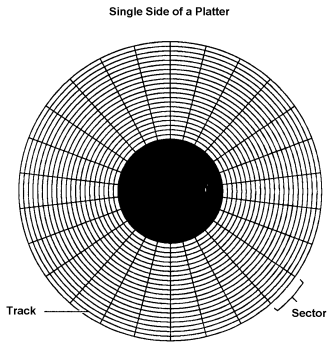

Disk Layout

- "*Platter*" – Circular Piece of Magnetic Material – Coated Aluminum
- One Platter Has 2 *"Sides"*
- Each Side is Divided into Concentric Rings Called *"Tracks"*
- Each Track is Divided into Arcs Called *"Sectors"*
- Each Sector Can Store Some Number of *"Bytes"* of Data
- *"Media Capacity"* Determines Amount of Storage per Sector

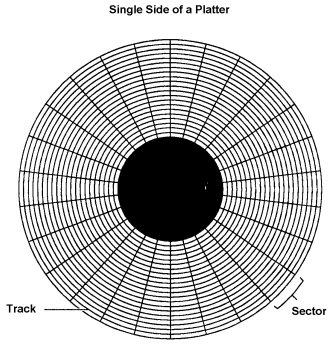

Physical Disk Coordinates

- A Platter can be *"Double Sided"* or *"Single Sided"*
- A Single Fixed Disk can Contain Multiple Platters
- Need Three-Dimensional Coordinate
- A Stack of Tracks On Top of Each Other is a *"Cylinder"*
- Each Side of a Platter Must Have a *"Read/Write Head"*

  Contains Induction Coil – Floats Over Surface
  (Except for Floppy Where it Actually Touches)

- Position of Read/Write Head Given by *"Cylinder"* Number
- Physical Location on Disk: **`<cylinder, head, sector>`**

Multi-Platter Disk Layout

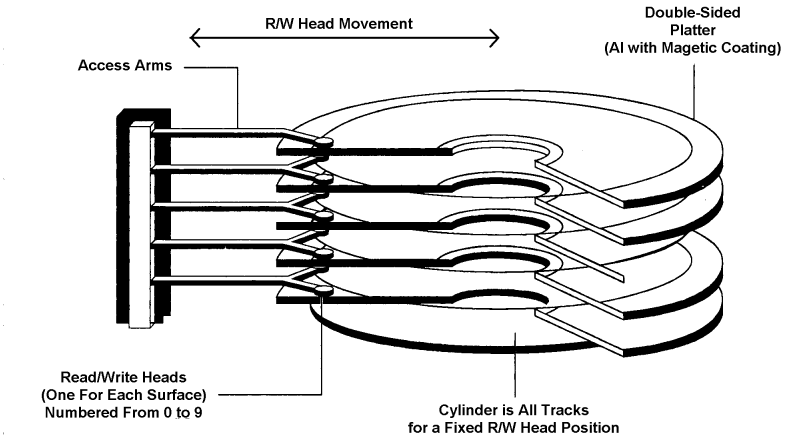

Disk Performance - Bandwidth

- *"Seek Time"* – (Average) Time Required to Position R/W Head

  Over a Track/Sector (5-15 ms)

- *"Rotational Latency"* – After Seek Wait for Sector to be Under R/W Head

  Typical Disks Rotate at 3600, 5400, 7200, 10800 RPM

  Average Rotational Latency = ½ Rotation (4-8ms)

- *"Transfer Time"* – Time for 1 Sector of Data to be Read/Written

  After Seek and Rotational Latency (25-50μs)

- *"Access Time"* is Sum of Above – Dominated by Seek and

  Rotational Latency

Disk Performance - Capacity

- *"Linear Bit Density"* – Number of Bits/cm that can be Stored Along a Track
- Linear Bit Density *(Typically 50,000-100,000 bits/cm)* Depends on:

  Purity of Magnetizable Media

  Air Quality

  Data Encoding

- *"Winchester Disk"* – Sealed at Factory to Preserve Air/Surface Quality

  -First "Winchester Disk" Was 30 MB Fixed and 30 MB Removable (IBM)

  -.30-30 Winchester

- *"Unformatted Capacity"* – Disregards Preamble, ECC and Intersector Gap

  -Marketing Hype

Capacity Along a Sector

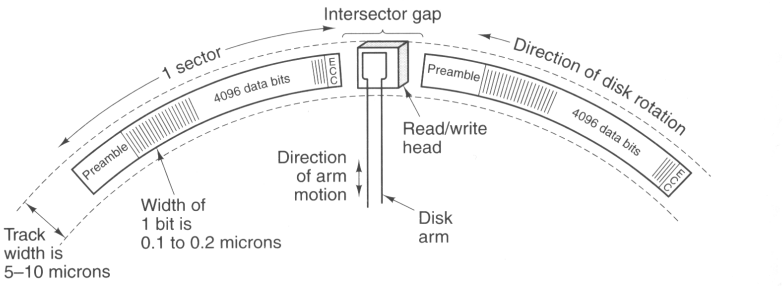

- Preamble – Used to Synchronize R/W Head

  "formatting" a disk is "writing" all preambles on it

- 4096 bits – Actual Data Storage Area (512 bytes is typical)
- ECC – Error Correcting Code – Usually a Reed-Solomon Code

Increasing Capacity - Cylinder Zones

Linear Bit Density – $c=2\pi r$

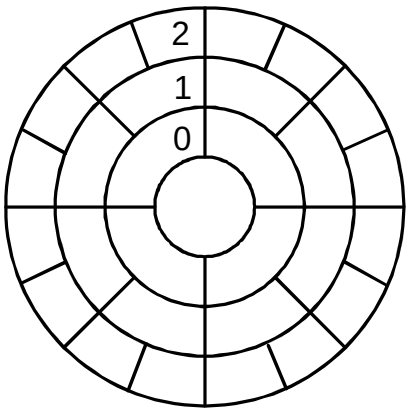

- *"Zones"* are Groups of Tracks
- Each Zone has Same Number Sectors/Track
- Attempts to Make Each Sector Have Same Linear Distance
- Increase Overall Capacity
- More Complex Seek Process

EXAMPLE

ZONE 0: 4 Sectors/Track
ZONE 1: 8 Sectors/Track
ZONE 2: 16 Sectors/Track

Typically 10-40 Zones Per Disk Drive

Data Encoding in a Sector

- **MFM (Modified Frequency Modulation)**

  Common technique used to encode the magnetic fluxes recorded on a drive into data. Still used on floppy drives and most original XT and AT systems. Notice that most drive types supported in CMOS have 17 sectors per track. This is the standard density for MFM encoding.

- **RLL (Run Length Limited)**

  Encoding method that allows 50% more information to be recorded on a track than MFM. Actually accomplished by recording more fluxes for every byte, but packing them more tightly onto the surface. Often called 2,7 RLL because the recording scheme involves patterns with no more than 7 successive zeros and no less than two.

- **ARLL (Advanced Run Length Limited)**

  More complex yet powerful derivatives of the RLL scheme. Include 1,7 and 3,9 encoding.

Floppy Disks

Most Common 5.25" and 3.5"

Parameter	Low Denstiy	High Density	Low Density	High Density
Size (inches)	5.25	5.25	3.5	3.5
Capacity (bytes)	360k	1.2M	720k	1.44M
Tracks	40	80	80	80
Sectors/Track	9	15	9	18
Heads	2	2	2	2
RPM	300	360	300	300
Rate (kbps)	250	500	250	500
Enclosure	flexible	flexible	rigid	rigid

Disk Controllers

- Each Disk Drive Has a Controller Associated With It

  Purpose of Controller:

  1) Accept Commands from Software (READ, WRITE, FORMAT)
  2) Control Arm Motion
  3) Detect/Correct Errors
  4) Convert parallel data into Serial Bit-Stream (vice-versa)
  5) Remap bad Sectors to Spares (each track has spares)
  6) Caching Sectors

- IDE – "Integrated Drive Electronics" – Most Controller Functionality is Inside the Drive Enclosure

- EIDE – "Extended IDE"

- SCSI – "Small Computer System Interface" – Actually a Bus Standard – Disk Media is Same

Disk Controllers

- **ST506/412**

  Standard interface used on XT and AT drives and controllers. Originally developed by Seagate Technologies to support their ST506 (5 MB) and ST412 (10 MB) drives. The entire controller mechanism is located on a controller card and communications between the drive and controller flow over 2 ribbon cables - one for drive control and one for data.

- **ESDI (Enhanced Small Device Interface)**

  Developed by Maxtor in the early 1980's as an upgrade and improvement to the ST506 design. While the drive does not have an embedded controller, one of the most critical functions , encoding-decoding, is performed on the drive. This allows for faster communications and higher drive capacities. Uses the same cabling as ST506 interface, but carries different signals on each line.

- **SCSI (Small Computer System Interface)**

  Based on an original design by Shugart Associates, SCSI is not specifically a drive interface, but a method of allowing different devices to communicate with a PC. For hard drives the entire controller is built onto the drive PCB, allowing for very high speed transfers to and from the drive. Fully interpreted, parallel data is then transferred to and from the PC by way of a single cable through a bus interface that has configured the device as a hard drive.

- **IDE (Integrated Drive Electronics)**

  A technology pioneered by Compaq and Conner that embedded a controller onto the hard disk PCB while maintaining compatibility with the register level commands sent by the computer's INT 13 routines. IDE drives are configured and appear to the computer like standard ST506 drives.

Disk Controllers

- **ATA (AT Attachment)**
  Implementation of the IDE design with a 16 bit AT style controller on board the drive.
- **XTA (XT Attachment)**
  Rarely used implementation of IDE with an integrated 8 bit XT controller.
- **ATA-2**
  Enhancement to the AT Attachment standard to provide for considerable performance improvement and more sophisticated drive identification.
- **EIDE (Enhanced IDE) and FAST-ATA**
  Various implementations of the ATA-2 standard as marketed by Western Digital (EIDE) and Seagate/Quantum (FAST-ATA).
- **UDMA-33 and UDMA-66 (Ultra-ATA)**
  Utilizes DMA channels and other tricks to increase peak theoretical transfer rates.

MOST MODERN SYSTEMS UTILIZE UDMA-33 and UDMA-66
33 Means a Peak Data Transfer Rate of 33 MHz
66 Means a Peak Data Transfer Rate of 66 MHz

IDE

- Evolved from Conventions Used by Original Fixed Disk in IBM PC XT
  10 MB Seagate – 4 heads, 306 cylinders, 17 sectors/track
  Xebec Disk Controller on Separate Plug-In Card on Motherboard
  Utilized BIOS Interrupt Services Exclusively
- BIOS Conventions Specify Location by <head, cylinder, sector>
  head and track numbers begin with 0
  sector begins with 1
- 4 Bits for Head, 6 Bits for Sector, 10 Bits for Cylinder
  Maximum Size 16 heads, 64 sectors, 1024 heads
  At 512 bytes/sector this Translates to a 528 MB Drive!
- First Drives Above 528 MB Limit had 4 heads, 32 sectors, 2000 cylinders
  Tricks Used to Fool Disk Controller
  Remapped Actual Geometry to Virtual Geometry
  Caused Trouble with OS

EIDE

- New Addressing Scheme to Overcome 528 MB Limit: LBA
- LBA – *Logical Block Address*
  - *Numbers Sectors from 0 through 2^{24}-1*
  - *Controller Must Convert Absolute Sector Number into <hd, cyl, sec>*
- Other Improvements
  - *EIDE Controller Can Control 4 Separate Drives (with IDE only 2)*
  - *Higher Transfer Rates*
  - *Can Also Control CD-ROM*

Disk Controller Data Transfer Rates

NAME	Data Bits	Bus MHz	MB/sec
SCSI-1	8	5	5
SCSI-2	8	5	5
IDE (ATA)	8	?	2.1-8.3
SCSI-2	16	5	10
Fast SCSI-2	8	10	10
EIDE (ATA-2)	16	?	11.1-16.6
Fast SCSI-2	16	10	20
Fast & Wide SCSI-2	16	10	20
UDMA-33	16	?	33.3
Ultra SCSI	16	20	40
UDMA-66	16	?	66.6
Ultra-2 SCSI	16	40	80