CH06 External Memory

- Magnetic Disk
- RAID
- Optical Memory
- Magnetic Tape

TECH
Computer Science

Types of External Memory

- Magnetic Disk
  - **RAID**
  - **Removable**
- Optical
  - **CD-ROM**
  - **CD-Writable (WORM)**
  - **CD-R/W**
  - **DVD**
- Magnetic Tape

Magnetic Disk

- Metal or plastic disk coated with magnetizable material (iron oxide…rust)
- Range of packaging
  - **Floppy**
  - **Winchester hard disk**
  - **Removable hard disk**

Data Organization and Formatting

- Concentric rings or tracks
  - **Gaps between tracks**
  - **Reduce gap to increase capacity**
  - **Same number of bits per track (variable packing density)**
  - **Constant angular velocity**
- Tracks divided into sectors
- Minimum block size is one sector
- May have more than one sector per block

Disk Data Layout

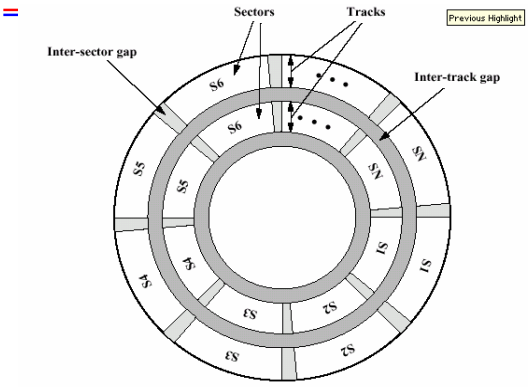

Fixed/Movable Head Disk

- Fixed head
  - **One read write head per track**
  - **Heads mounted on fixed ridged arm**
- Movable head
  - **One read write head per side**
  - **Mounted on a movable arm**

Fixed and Movable Heads

Removable or Not //

- Removable disk
  - **Can be removed from drive and replaced with another disk**
  - **Provides unlimited storage capacity**
  - **Easy data transfer between systems**
- Nonremovable disk
  - **Permanently mounted in the drive**

Floppy Disk

- 8", 5.25", 3.5"
- Small capacity
  - **Up to 1.44Mbyte (2.88M never popular)**
- Slow (disk rotate at 300 and 600 rpm, average delay 100/2 and 200/2 ms.)
- Universal
- Cheap

Winchester Hard Disk (1)

- Developed by IBM in Winchester (USA)
- Sealed unit
- One or more platters (disks)
- Heads fly on boundary layer of air as disk spins (crash into disk!)
- Very small head to disk gap
- Getting more robust

Winchester Hard Disk (2)

- Universal
- Cheap
- Fastest external storage (typically rotate 3600 rpm, newer faster, average rotational delay 8.3 ms.)
- Getting larger all the time
  - **Multiple Gigabyte now usual**

Removable Hard Disk

- ZIP
  - **Cheap**
  - **Very common**
  - **Only 100M**
- JAZ
  - **Not cheap**
  - **1G**
- L-120 (a: drive)
  - **Also reads 3.5" floppy**
  - **Becoming more popular?**

Finding Sectors

- Must be able to identify start of track and sector
- Format disk
  - **Additional information not available to user**
  - **Marks tracks and sectors**

ST506 format (old!)

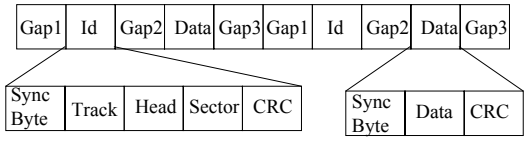

- Foreground reading
  - **Find others**

Characteristics

- Fixed (rare) or movable head
- Removable or fixed
- Single or double (usually) sided
- Single or multiple platter
- Head mechanism
  - **Contact (Floppy)**
  - **Fixed gap**
  - **Flying (Winchester)**

Multiple Platter

- One head per side
- Heads are joined and aligned
- Aligned tracks on each platter form cylinders
- Data is striped by cylinder
  - **reduces head movement**
  - **Increases speed (transfer rate)**

Speed

- Seek time
  - **Moving head to correct track**
- (Rotational) latency
  - **Waiting for data to rotate under head**
- Access time = Seek + Latency
- Transfer rate T = (number of bytes to be transferred)/(rotation speed)/(number of bytes on a track) = b/(rN)
- total access time $T_a = T_s + 1/(2r) + b/(rN)$

Sequential organization vs. random access e.g.

  - **e.g. a hard disk has average seek time of 20 ms, a transfer rate of 1 M byte/s, and 512 byte sectors with 32 sectors per track. Need to read a file consisting 256 sectors for a total of 128 K bytes. What is the total time for the transfer?**
- Case 1: Sequential Organization (256 sectors on 8 tracks x 32 sectors/tracks)
  - **Average seek time = 20.0 ms**
  - **Rotational delay = 8.3 ms**
  - **Read 32 sections (one track) = 16.7 ms**
  - **total time to read first track = 45 ms**
  - **Total time = 45 ms + 7*(8.3 + 16.7) ms = 0.22 s**

Time required for random access on highly fragmented organization

- Case 2: random access rather than sequential access
  - **Average seek time = 20.0 ms**
  - **Rotational delay = 8.3 ms**
  - **Read 1 sector = 16.7/32 = 0.5 ms**
  - **time to read one sector = 28.8 ms**
  - **Total time = 256 * 28.8 ms = 7.37 s**
- De-fragment you hard disk!

Optical Storage CD-ROM //

- Originally for audio
- 650Mbytes giving over 70 minutes audio
- Polycarbonate coated with highly reflective coat, usually aluminum
- Data stored as pits
- Read by reflecting laser
- Constant packing density
- Constant linear velocity

Constant Angular Velocity vs. Constant Linear Velocity

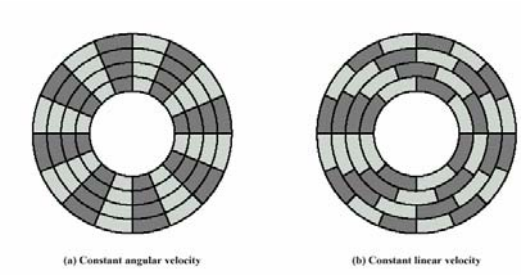

CD-ROM Drive Speeds

- Audio is single speed
  - **Constant linear velocity**
  - **1.2 ms^{-1}**
  - **Track (spiral) is 5.27km long**
  - **Gives 4391 seconds = 73.2 minutes**
  - **Date 176.4 K bytes/s total capacity 774.57 M Bytes**
- Other speeds are quoted as multiples
- e.g. 24x ~= 4 M Bytes/s (data transfer rate)
- The quoted figure is the maximum the drive can achieve

CD-ROM Format

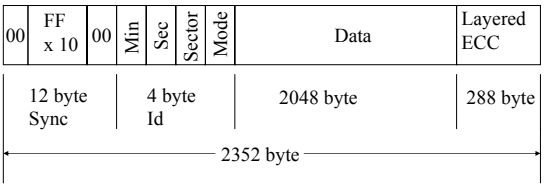

- Mode 0=blank data field
- Mode 1=2048 byte data+error correction
- Mode 2=2336 byte data

Random Access on CD-ROM

- Difficult
- Move head to rough position
- Set correct speed
- Read address
- Adjust to required location
- (Yawn!)

CD-ROM for & against

- Large capacity (?)
- Easy to mass produce
- Removable
- Robust

- Expensive for small runs
- Slow
- Read only

Other Optical Storage

- CD-Writable
  - **WORM**
  - **Now affordable**
  - **Compatible with CD-ROM drives**
- CD-RW
  - **Erasable**
  - **Getting cheaper**
  - **Mostly CD-ROM drive compatible**

DVD - what's in a name?

- Digital Video Disk
  - **Used to indicate a player for movies**
    - Only plays video disks
- Digital Versatile Disk
  - **Used to indicate a computer drive**
    - Will read computer disks and play video disks
- Dogs Veritable Dinner
- Officially - nothing!!!

DVD - technology

- Multi-layer
- Very high capacity (4.7G per layer)
- dual-layer (single-sided ?) hold 8.5 Gbytes ~> 4hr movie
- Full length movie on single disk
  - **Using MPEG compression**
- Finally standardized (honest!)
- Movies carry regional coding
- Players only play correct region films

DVD - Writable

- Loads of trouble with standards
- First generation DVD drives may not read first generation DVD-W disks
- First generation DVD drives may not read CD-RW disks
- Wait for it to settle down before buying!

Foreground Reading

- Check out optical disk storage options
- Check out Mini Disk

Magnetic Tape

- Serial access
- Slow
- Very cheap
- Backup and archive

Digital Audio Tape (DAT)

- Uses rotating head (like video)
- High capacity on small tape
  - **4Gbyte uncompressed**
  - **8Gbyte compressed**
- Backup of PC/network servers

RAID

- Redundant Array of Independent Disks
- Redundant Array of Inexpensive Disks
- 6 levels in common use
- Not a hierarchy
- Set of physical disks viewed as single logical drive by O/S
- Data distributed across physical drives
- Can use redundant capacity to store parity information

RAID Levels 0, 1, 2

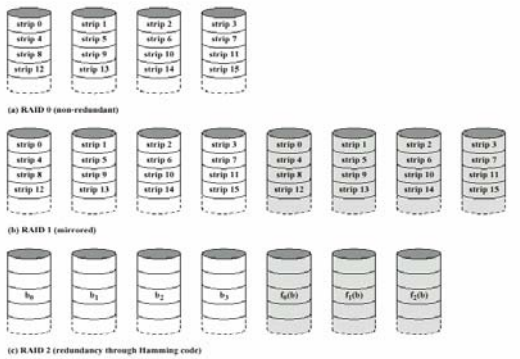

RAID Levels 3, 4

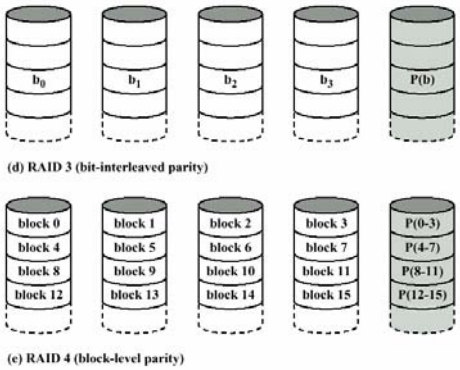

RAID Levels 5, 6

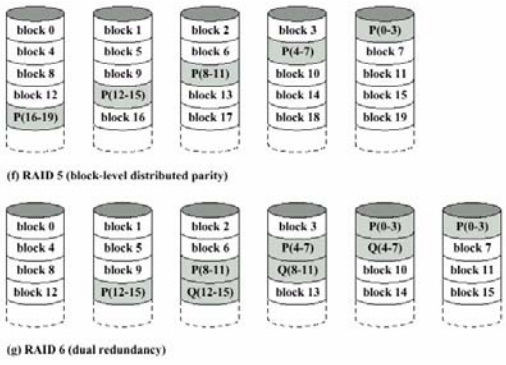

RAID 0

- No redundancy
- Data striped across all disks
- Round Robin striping
- Increase speed
  - **Multiple data requests probably not on same disk**
  - **Disks seek in parallel**
  - **A set of data is likely to be striped across multiple disks**

RAID 1

- Mirrored Disks
- Data is striped across disks
- 2 copies of each stripe on separate disks
- Read from either
- Write to both
- Recovery is simple
  - **Swap faulty disk & re-mirror**
  - **No down time**
- Expensive

RAID 2

- Disks are synchronized
- Very small stripes
  - **Often single byte/word**
- Error correction calculated across corresponding bits on disks
- Multiple parity disks store Hamming code error correction in corresponding positions
- Lots of redundancy
  - **Expensive**
  - **Not used**

RAID 3

- Similar to RAID 2
- Only one redundant disk, no matter how large the array
- Simple parity bit for each set of corresponding bits
- Data on failed drive can be reconstructed from surviving data and parity info
- Very high transfer rates

RAID 4

- Each disk operates independently
- Good for high I/O request rate
- Large stripes
- Bit by bit parity calculated across stripes on each disk
- Parity stored on parity disk

RAID 5

- Like RAID 4
- Parity striped across all disks
- Round robin allocation for parity stripe
- Avoids RAID 4 bottleneck at parity disk
- Commonly used in network servers

RAID 6

- Two different parity calculations are carried out and
- stored in separate blocks on different disks.
- Able to regenerate data even if two disks containing user data fail