


# Magnetic disk in floppy disk



# Secondary Storage Devices

- **Two major types of storage devices:**
  - 1. Direct Access Storage Devices (DASDs)**
    - **Magnetic Disks**
      - Hard disks (high capacity, low cost per bit)**
      - Floppy disks (low capacity, slow, cheap)**
    - **Optical Disks**
      - CD-ROM = (Compact disc, read-only memory)**
  - 2. Serial Devices**
    - **Magnetic tapes (very fast sequential access)**

# Magnetic Disks


- **Bits of data (0's and 1's) are stored on circular magnetic platters called disks.**
- **A disk rotates rapidly (& never stops).**
- **A disk head reads and writes bits of data as they pass under the head.**
- **Often, several platters are organized into a disk pack (or disk drive).**

# Bits of Data

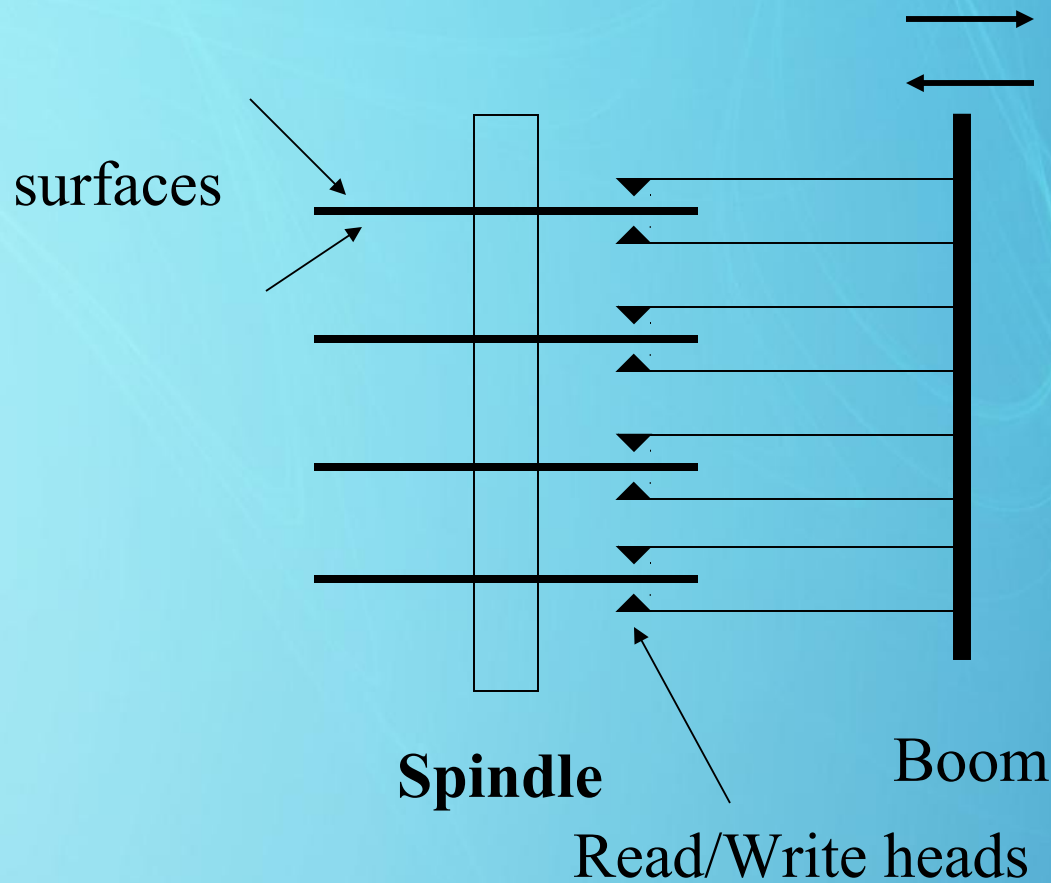
From Computer Desktop Encyclopedia  
© 1998 The Computer Language Co. Inc.

□ 0 bit    ■ 1 bit

recording channels (tracks)



# A Disk Drive




Disk drive with 4 platters and 8 surfaces

From Computer Desktop Encyclopedia  
Reproduced with permission.  
© 1997 Singapore Technologies



From Computer Desktop Encyclopedia  
© 1998 The Computer Language Co. Inc.



# Looking at the Surface

# Organization of Disks

- **Disk contains concentric tracks.**
- **Tracks are divided into sectors**
- **A sector is the smallest addressable unit in a disk.**


# Accessing Data

- When a program reads a byte from the disk, the operating system locates the surface, track and sector containing that byte, and **reads the entire sector into a special area in main memory called buffer.**
- The bottleneck of a disk access is moving the read/write arm. So it makes sense to store a file in tracks that are below/above each other in different surfaces, rather than in several tracks in the same surface.

# Cylinders

- **A cylinder is the set of tracks at a given radius of a disk pack.**
  - i.e. a cylinder is the set of tracks that can be accessed without moving the disk arm.
- **All the information on a cylinder can be accessed without moving the read/write arm.**

# Cylinders



# Estimating Capacities

- **Track capacity = # of sectors/track \* bytes/sector**
- **Cylinder capacity = # of tracks/cylinder \* track capacity**
- **Drive capacity = # of cylinders \* cylinder capacity**
- **Number of cylinders = # of tracks in a surface**

# Exercise

- Store a file of 20000 records on a disk with the following characteristics:

# of bytes per sector = 512

# of sectors per track = 40

# of tracks per cylinder = 12

# of cylinders = 1331

**Q1.** How many cylinders does the file require if each data record requires 256 bytes?

**Q2.** What is the total capacity of the disk?

# Clusters

- Another view of sector organization is the one maintained by the O.S.'s **file manager**.
- It views the file as a series of **clusters** of sectors.
- File manager uses a **file allocation table (FAT)** to map logical sectors of the file to the physical clusters.

# Extents

- If there is a lot of room on a disk, it may be possible to make a file consist entirely of contiguous(near by)clusters. Then we say that the file is **one extent**. (very good for sequential processing)
- If there isn't enough contiguous space available to contain an entire file, the file is divided into two or more noncontiguous parts. Each part is an extent.

# Fragmentation

- **Internal fragmentation: loss of space within a sector or a cluster.**
  - 1) **Due to records not fitting exactly in a sector: e.g. Sector size is 512 and record size is 300 bytes. Either**
    - store one record per sector, or
    - allow records *span(duration)* sectors.
  - 2) **Due to the use of clusters: If the file size is not a multiple of the cluster size, then the last cluster will be partially used.**

# Choice of cluster size

- **Some operating systems allow system administrator to choose cluster size.**
- **When to use large cluster size?**
- **What about small cluster size?**

# Organizing Tracks by Block

- **Disk tracks may be divided into user-defined blocks rather than into sectors.**
- **Blocks can be fixed or variable length.**
- **A block is usually organized to hold an integral number of logical records.**
- **Blocking Factor = number of records stored in a block.**
- **No internal fragmentation, no record spanning two blocks.**
- **In block-addressing scheme each block of data is accompanied by one or more *sub blocks* containing extra information about the block.**

# Non-data Overhead

- **Both blocks and sectors require non-data overhead (written during formatting)**
- **On sector addressable disks this information involves sector address, track address, and condition (usable/defective). Also pre-formatting involves placing gaps and synchronization marks.**
- **On block-organized disk, more information is needed and the programmer should be aware of some of this information.**

# Exercise

- Consider a *block-addressable* disk with the following characteristics:
  - Size of track 20,000 bytes.
  - Nondata overhead per block = 300 bytes.
  - Record size = 100 byte.
- **Q) How many records can be stored per track if blocking factor is**
  - a) 10
  - b) 60

# The Cost of a Disk Access

- The time to access a sector in a track on a surface is divided into 3 components:

## Time Component

## Action

Seek Time

Time to move the read/write arm to the correct cylinder

Rotational delay (or latency)

Time it takes for the disk to rotate so that the desired sector is under the read/write head

Transfer time

Once the read/write head is positioned over the data, this is the time it takes for transferring data

# Seek time

- Seek time is the time required to move the arm to the correct cylinder.
- Largest in cost.

## Typically:

- 5 ms (miliseconds) to move from one track to the next (track-to-track)
- 50 ms maximum (from inside track to outside track)
- 30 ms average (from one random track to another random track)

# Average Seek Time ( $s$ )

- Since it is usually impossible to know exactly how many tracks will be traversed in every seek, we usually try to determine the average seek time ( $s$ ) required for a particular file operation.
- If the starting and ending positions for each access are random, it turns out that the average seek traverses one third of the total number of cylinders.
- Manufacturer's specifications for disk drives often list this figure as the average seek time for the drives.
- Most hard disks today have  $s$  of less than 10 ms, and high-performance disks have  $s$  as low as 7.5 ms.

# Latency (rotational delay)

- Latency is the time needed for the disk to rotate so the sector we want is under the read/write head.
- Hard disks usually rotate at about 5000rpm, which is one revolution per 12 msec.
- Note:
  - Min latency = 0
  - Max latency = Time for one disk revolution
  - **Average latency ( $r$ )** =  $(\text{min} + \text{max}) / 2$   
=  $\text{max} / 2$   
= time for  $\frac{1}{2}$  disk revolution
- Typically 6 – 8 ms average

# Transfer Time

- Transfer time is the time for the read/write head to pass over a block.
- The transfer time is given by the formula:

$$\text{Transfer time} = \frac{\text{number of bytes transferred}}{\text{number of bytes on a track}} \times \text{rotation time}$$

- e.g. if there are 63 sectors per track, the time to transfer one sector would be 1/63 of a revolution.

# Exercise

Given the following disk:

- 20 surfaces  
800 tracks/surface  
25 sectors/track  
512 bytes/sector
- 3600 rpm (revolutions per minute)
- 7 ms track-to-track seek time  
28 ms avg. seek time  
50 ms max seek time.

Find:

- Average latency
- Disk capacity
- Time to read the entire disk, one cylinder at a time

# Sequential Reading

- Given the following disk:
  - $s = 16$  ms
  - $r = 8.3$  ms
  - Block transfer time = 8.4 ms
- a) Calculate the time to read 10 sequential blocks
- b) Calculate the time to read 100 sequential blocks

# Random Reading

Given the same disk,

- a) Calculate the time to read 10 blocks randomly
- b) Calculate the time to read 100 blocks randomly

# Fast Sequential Reading

- **We assume that blocks are arranged so that there is no rotational delay in transferring from one track to another within the same cylinder. This is possible if consecutive track beginnings are staggered (like running races on circular race tracks)**
- **We also assume that the consecutive blocks are arranged so that when the next block is on an adjacent cylinder, there is no rotational delay after the arm is moved to new cylinder**
- ***Fast sequential reading*: no rotational delay after finding the first block.**

# Consequently ...

Reading  $b$  blocks:

i. Sequentially:

$$\underbrace{s + r + b * btt}$$

insignificant for large files

$$\Rightarrow b * btt$$

ii. Randomly:

$$b * (s + r + btt)$$

# Exercise

- **Specifications of a 300MB disk drive:**
  - Min seek time = 6ms.
  - Average seek time = 18ms
  - Rotational delay = 8.3ms
  - transfer rate = 16.7 ms/track or 1229 bytes/ms
  - Bytes per sector = 512
  - Sectors per track = 40
  - Tracks per cylinder = 12
  - Tracks per surface = 1331
  - Interleave factor = 1
  - Cluster size= 8 sectors
  - Smallest extent size = 5 clusters

**Q) How long will it take to read a 2048Kb file that is divided into 8000 256 byte records?**

- Access the file sequentially**
- Access the file randomly**


# **Secondary Storage Devices: CD-ROM**

# Physical Organization of CD-ROM

- Compact Disk – read only memory (write once)
- Data is encoded and read optically with a laser
- Can store around 600MB data
- Digital data is represented as a series of **Pits** and **Lands**:
  - Pit = a little depression, forming a lower level in the track
  - Land = the flat part between pits, or the upper levels in the track

# Organization of data

- Reading a CD is done by shining a laser at the disc and detecting changing reflections patterns.
  - 1 = change in height (land to pit or pit to land)
  - 0 = a “fixed” amount of time between 1’s



- Note : we cannot have two 1’s in a row!  
=> uses Eight to Fourteen Modulation (EFM) encoding table.

# Properties

- Note that: Since 0's are represented by the **length of time** between transitions, we must travel at **constant linear velocity (CLV)** on the tracks.
- Sectors are organized along a spiral
- Sectors have same linear length
- Advantage: takes advantage of all storage space available.
- Disadvantage: has to change rotational speed when seeking (slower towards the outside)

# Addressing

- 1 second of play time is divided up into 75 **sectors**.
- Each sector holds 2KB
- 60 min CD:  
 $60\text{min} * 60 \text{ sec/min} * 75 \text{ sectors/sec} =$   
 $270,000 \text{ sectors} = 540,000 \text{ KB} \sim 540 \text{ MB}$
- A **sector** is addressed by:  
Minute:Second:Sector  
e.g. 16:22:34