

Input/Output – 5

Disks

CS 342 – Operating Systems

Ibrahim Korpeoglu

Bilkent University

Department of Computer Engineering

- Disks have many types

- Magnetic disks

- Hard disk
 - Floppy disk

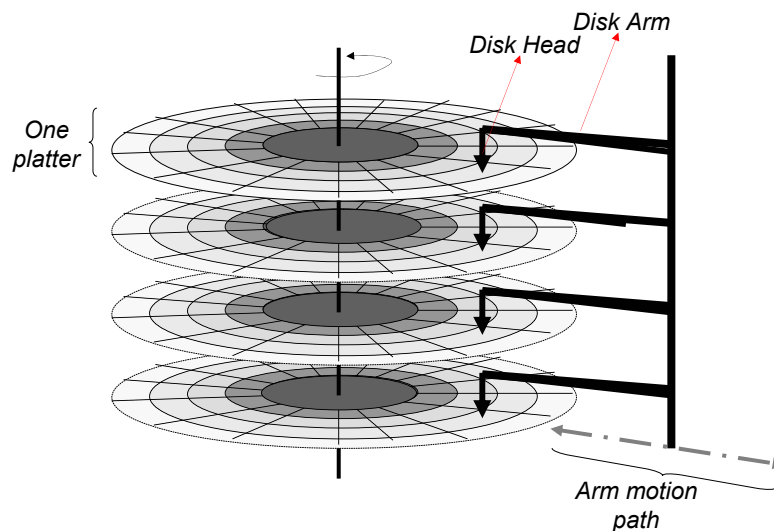
- Optical disks

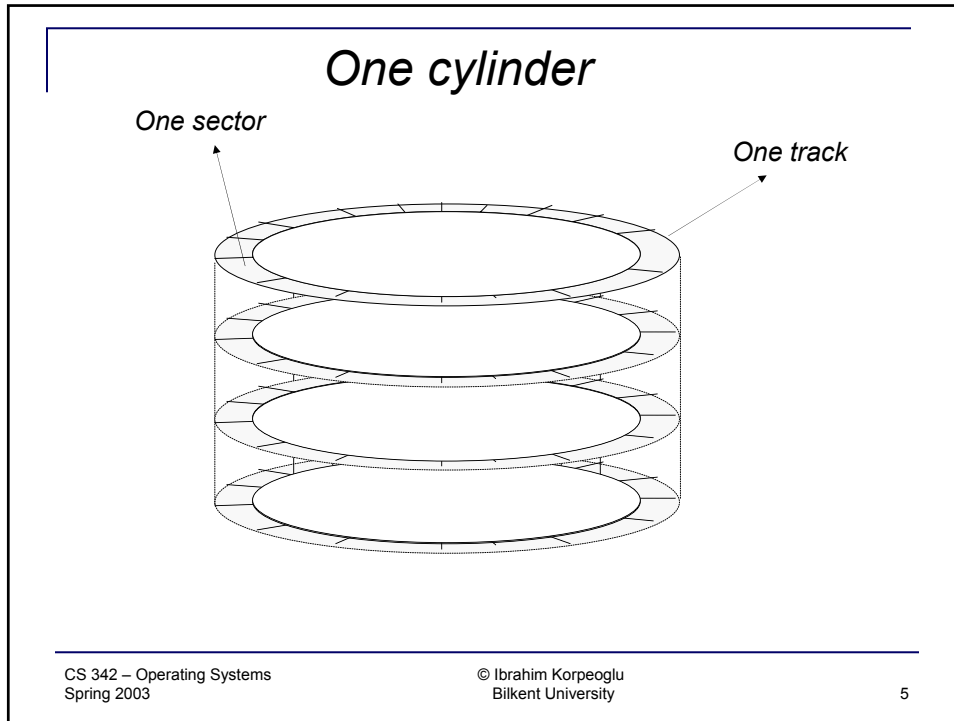
- CD-ROM
 - CD-Writable
 - CD-Recordable
 - DVDs

Magnetic Disks

- Organized into *cylinders*
 - Each cylinders contains *tracks* whose number is equal to the number of heads in the disk drive
 - Each track is divided into a number of sectors: 8-32 sectors/track.
 - The sector size is usually 512 bytes.
- There are IDE disks
- There are SCSI disks, etc.

A Hard Disk





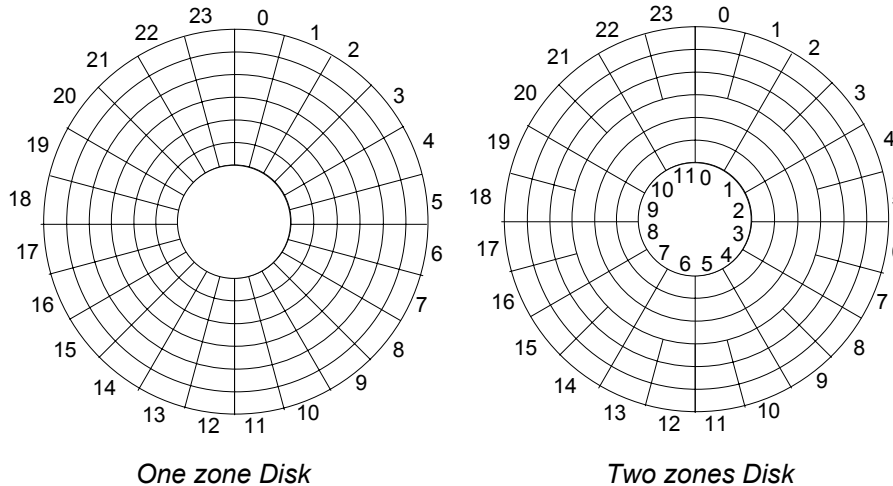
Disk Parameters

Parameter	WD 18300 Hard Disk
Number of cylinders	10601
Tracks per cylinder	12
Sectors per track	281 (average)
Sectors per disk	35742000
Bytes per sector	512
Disk capacity	18.3 GB
Seek time (adjacent cylinders)	0.8 msec
Seek time (average case)	6.9 msec
Rotation time	8.33 msec
Motor stop/start time	20 sec
Time to transfer 1 sector	17 μ sec

CS 342 – Operating Systems
 Spring 2003

 © Ibrahim Korpeoglu
 Bilkent University
 6

Physical Geometry



CS 342 – Operating Systems
Spring 2003

© Ibrahim Korpeoglu
Bilkent University

7

Disk Formatting

- After manufacturing, there is no information on the disk
 - Just empty bits.
- Each platter need to have
 - A low-level format
 - A high-level formatbefore disk can be used.

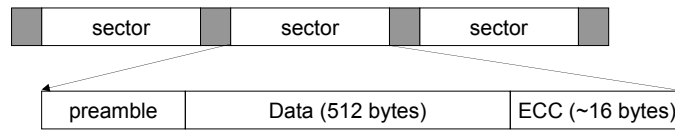
CS 342 – Operating Systems
Spring 2003

© Ibrahim Korpeoglu
Bilkent University

8

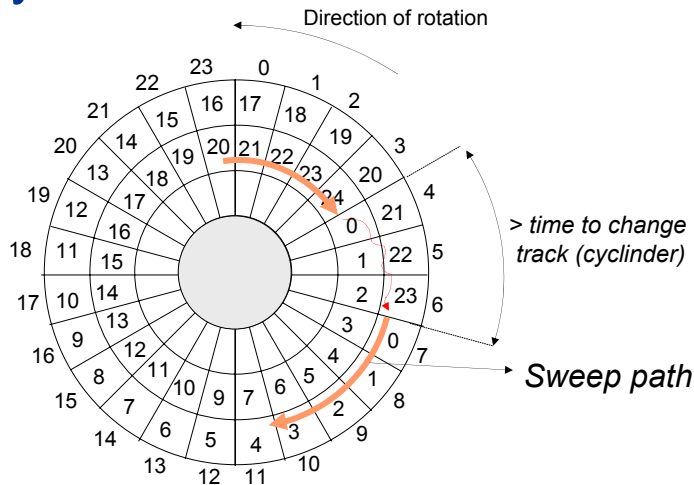
Disk Formatting

- Each track should have formatted in the following way:
 - Sectors and inter-sector-gap between them
 - A sector will have format like the following:
 - Preamble – data (512 bytes) – Checksum

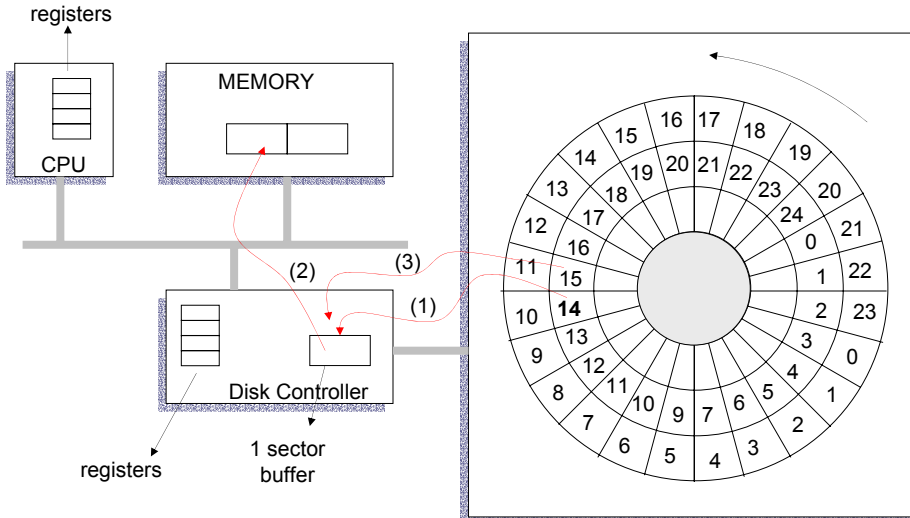


-certain bit pattern
-cylinder number
-sector number

Cylinder Skew



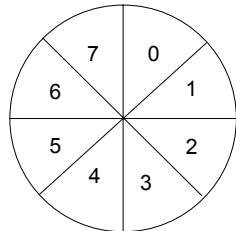
Interleaving



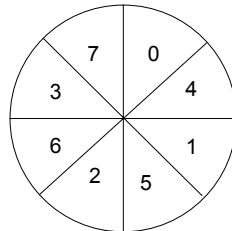
Interleaving

- If we have ONE sector buffer in the disk controller:
 - After we have transferred one sector of data from hard-disk to controller buffer:
 - We will copy the controller buffer to the memory
 - This may take time.
 - During this time, the disk-head will pass the start of the next sector.
 - Therefore, the next logical sector, should not be the next physical sector in hard-disk.
 - There should be some interleaving.

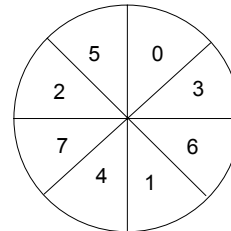
Interleaving



Non-interleaving



Single Interleaving



Double Interleaving

Disk Arm Scheduling

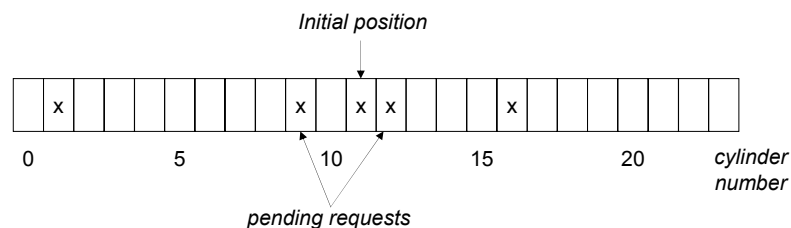
- How long does it take to read a block (sector)
 - Seek time
 - Move arm to the correct cylinder
 - Rotational delay
 - Wait until correct sector comes under head.
 - Actual data transfer time
- For most systems, the seek time dominates the other two times.

Disk Arm Scheduling

- Assume the disk is heavily loaded.
 - There are a lot requests for disk blocks that are arriving to the hard *disk driver*.
 - The drivers queues the requests.
 - The driver knows for each block request, where in hard disk the block should be stored (the cylinder number).
 - Since moving between cylinders (seek time) is costly, we should try to minimize these seek times.

Disk Arm Scheduling

- Disk driver
 - Keeps a table of requests.
 - Table is indexed by the cylinder number
 - Blocks Requests for the same cylinder are put in a linked list.



Disk Arm Scheduling Algorithms

- Several Algorithms Exist
 - First Come First Served (FCFS)
 - Shortest Seek First (SSF)
 - Elevator Algorithm

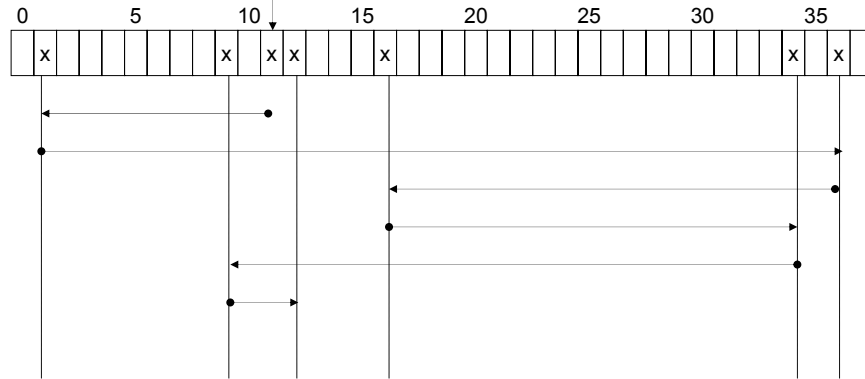
FCFS

- Serve the request in the order they arrive (don't look to the cylinder numbers)
- Example:
 - Assume the current position is cylinder 11.
 - The following requests arrive in the given time order: 1, 36, 16, 34, 9, 12
 - The request will be served in the same order.
 - The head will go to:
 - Cylinder 1 first (10 cylinder motion)
 - Cylinder 36 next (35 cylinder motion)
 - Cylinder 16 next (20)
 - Cylinder 34 next (18)
 - Cylinder 9 next (25)
 - Cylinder 12 next (3)
 - Total of 111 cylinders are skipped. Cost = 111

FCFS

New request arrival sequence: 1, 36, 16, 34, 9, 12

Initial head position



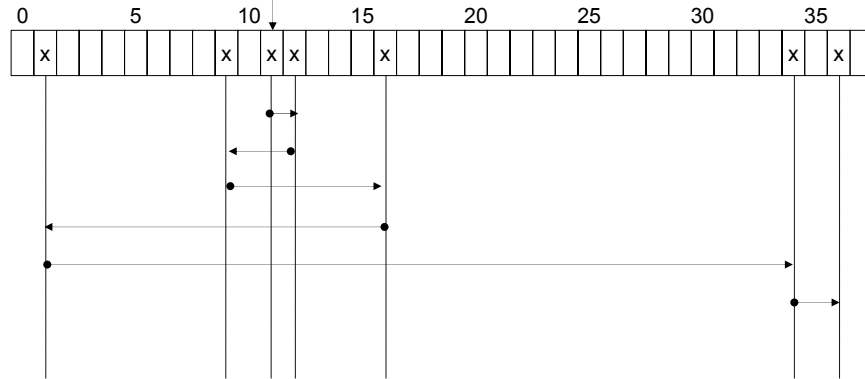
Shortest Seek First

- Satisfy the request that need shortest seek from the current position.
- In the previous example, this algorithms moved the head to cylinders in the order given below.
 - 12, 9, 16, 1, 34, and 36.
- Moving the head to these cylinders require:
 - 1, 3, 7, 15, 33, and 2 arm motions (that many cylinders are skipped at every request).
 - Total cost = 61.

SSF

New request arrival sequence: 1, 36, 16, 34, 9, 12

Initial head position



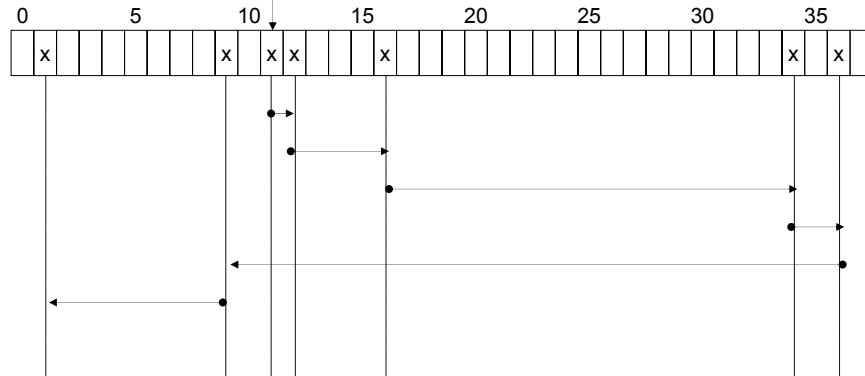
Elevator Algorithm

- Do like what an elevator does.
 - 1. Move in one direction until all request in that direction is served
 - 2. Then change direction.
 - 3. Move in the new direction until all request in that direction are server.
 - Repeat steps 1, 2, and 3.
- In the previous example
 - Assume initial position was 11 and initial direction was up (going to higher numbered cylinders).

Elevator Algorithm

New request arrival sequence: 1, 36, 16, 34, 9, 12

Initial head position



Elevator Algorithm

- The request are server in the following order:
 - 12, 16, 34, 36, 9, 1
- Arm motions (# of cylinder skips)
 - 1, 4, 18, 2, 27, 8
 - Total cost = 60 cylinders.
- *Performs good and is fair.*