الملزمة العاشرة

Mass-Storage Systems

- Magnetic disks provide bulk of secondary storage of modern computers.
- **Transfer rate** is rate at which data flow between drive and computer.
- **Positioning time** (**random-access time**) is time to move disk arm to desired cylinder (**seek time**) and time for desired sector to rotate under the disk head (**rotational latency**).
- Drive attached to computer via I/O Bus. Busses vary, including EIDE, ATA, SATA, USB, Fiber Channel, SCSI.

Disk Structure

Disk drives are addressed as large 1-dimensional arrays of logical blocks, where the logical block is the smallest unit of transfer. The 1-dimensional array of logical blocks is mapped into the ***sectors*** of the disk sequentially. Sector 0 is the first sector of the first ***track*** on the outermost ***cylinder***. Mapping proceeds in order through that track, then the rest of the tracks in that cylinder, and then through the rest of the innermost cylinders from outermost to innermost.

H.W: Check your Hard Disk using EVEREST program.

Ex: SAMSUNG SP0411N hard disk has the following parameters:
- 77622 cylinders.
- 16 heads.
- 63 sector per track.
- 554 byte per sector.
- Rotational Speed: 7200 RPM.

- Seek Time: 19.5 ms.
- Interface: Ultra ATA/133.
- Buffer Size: 2MB.

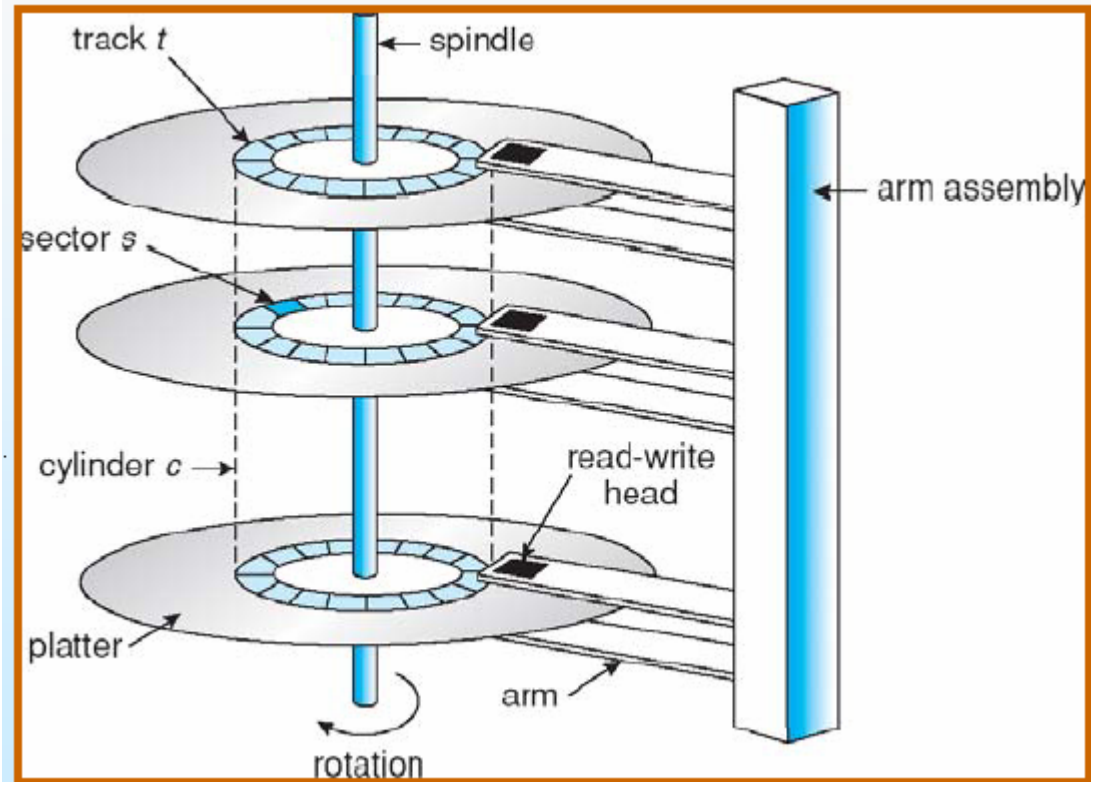

Disk Scheduling

The operating system is responsible for using hardware efficiently for the disk drives, this means having a fast access time and disk bandwidth.

- Several algorithms exist to schedule the servicing of disk I/O requests.

We illustrate them with a request queue (0-199).

98, 183, 37, 122, 14, 124, 65, 67

Head pointer 53

1- FCFS

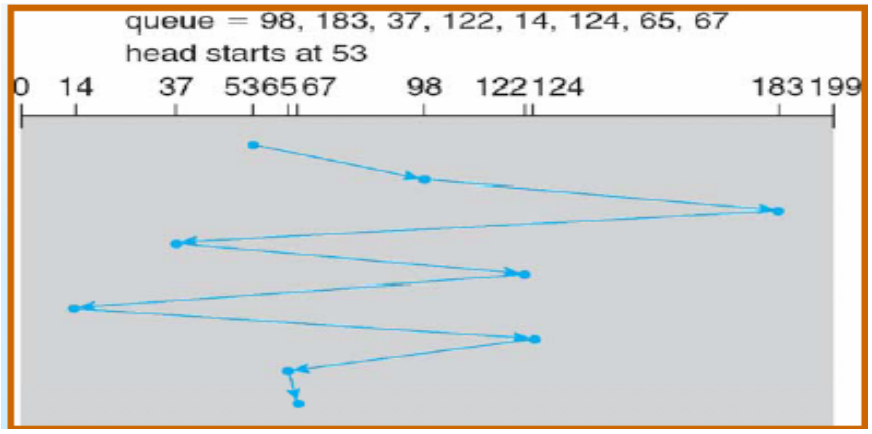

Illustration shows total head movement of 640 cylinders.

2- SSTF

Selects the request with the minimum seek time from the current head position.

- SSTF scheduling is a form of SJF scheduling; may cause starvation of some requests.
- Illustration shows total head movement of 236 cylinders.

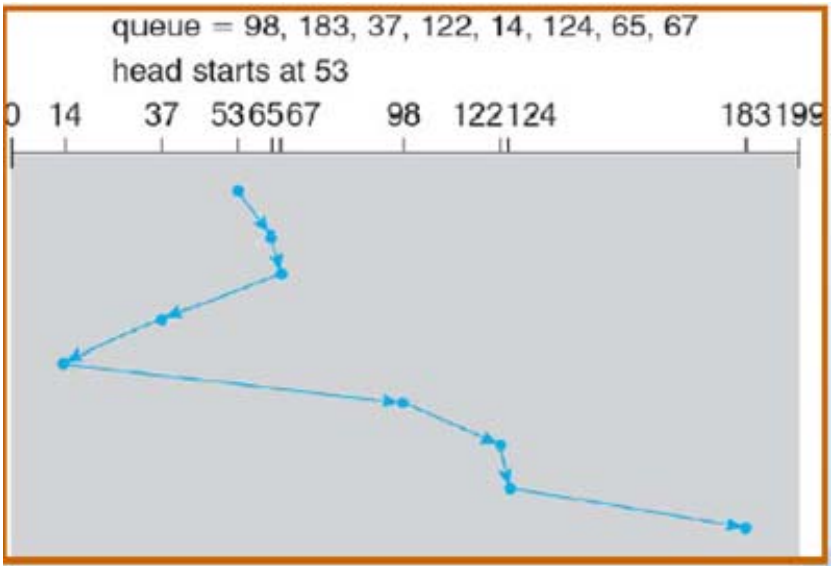

3- SCAN

The disk arm starts at one end of the disk, and moves toward the other end, servicing requests until it gets to the other end of the disk, where the head movement is reversed and servicing continues.

- Illustration shows total head movement of 208 cylinders.

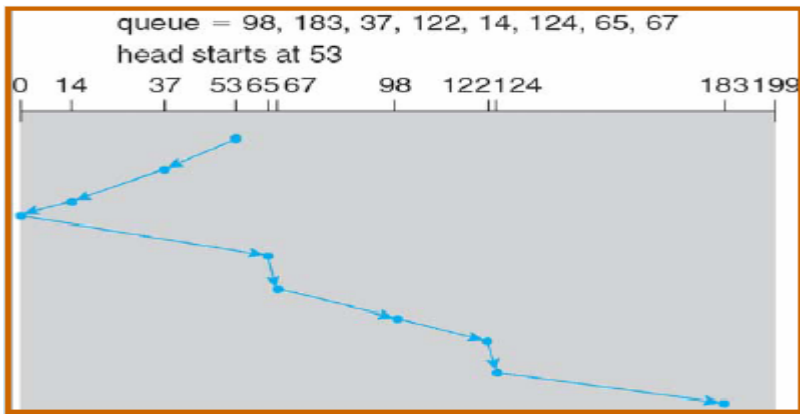

4- C-LOOK

- Version of C-SCAN.

Arm only goes as far as the last request in each direction, then reverses direction immediately, without first going all the way to the end of the disk.

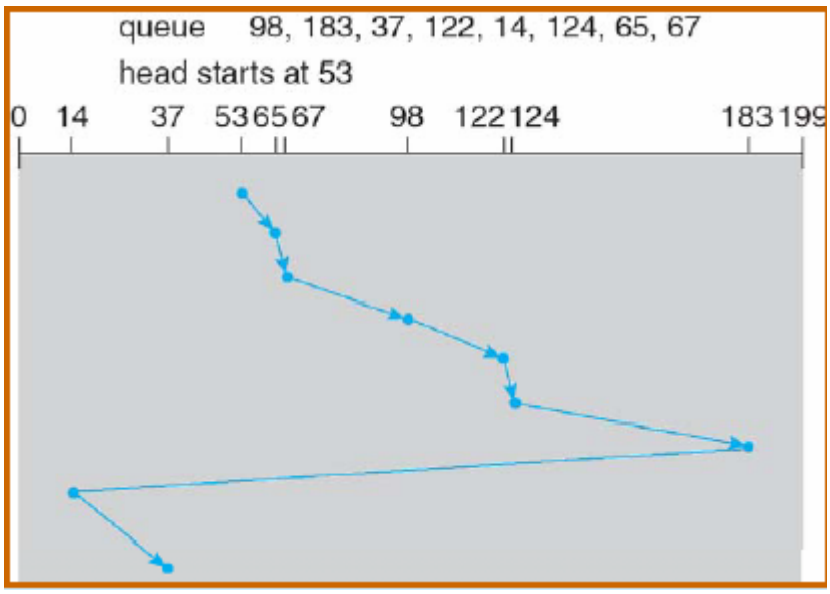

RAID (Redundant Array of Independent Disks)

A category of disk drives that employ two or more drives in combination for fault tolerance and performance.

- RAID divided into six different levels.

RAID-0 (striping)

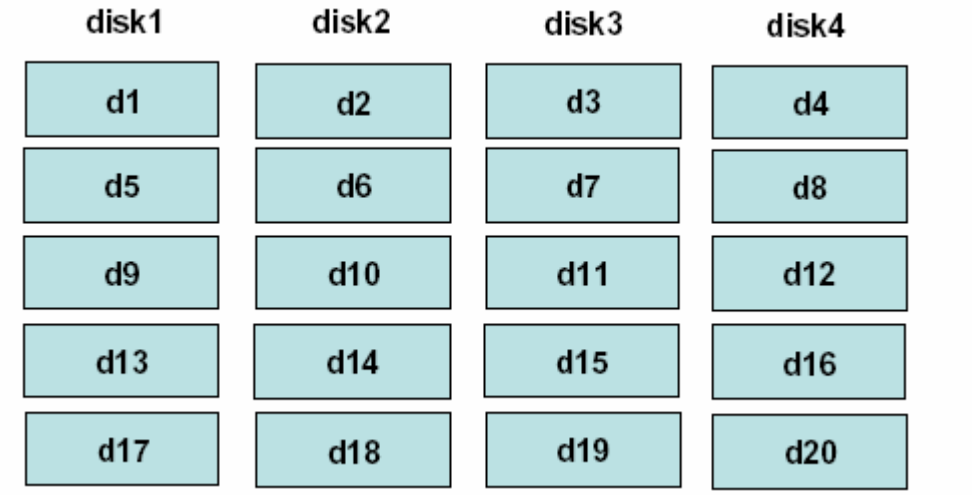

•RAID-0 implements a striped disk array, the data is broken down into blocks and each block is written to a separate disk drive.

•Not a fault-tolerant but more efficient in data read.

RAID-1 (Mirroring)

- Duplicate each disk.
- 100% redundancy of data means no rebuild is necessary in case of a disk failure, just a copy to the replacement disk.

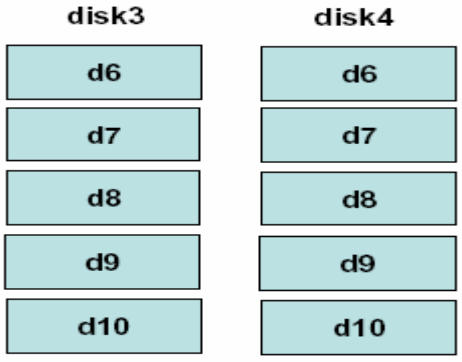

RAID-3

• The data block is subdivided ("striped") and written on the data disks. Stripe parity is generated on Writes, recorded on the parity disk and checked on Reads.

• RAID Level 3 requires a minimum of 3 drives to implement

• Taking longer to recover failure than RAID-1 but spending less on redundant storage

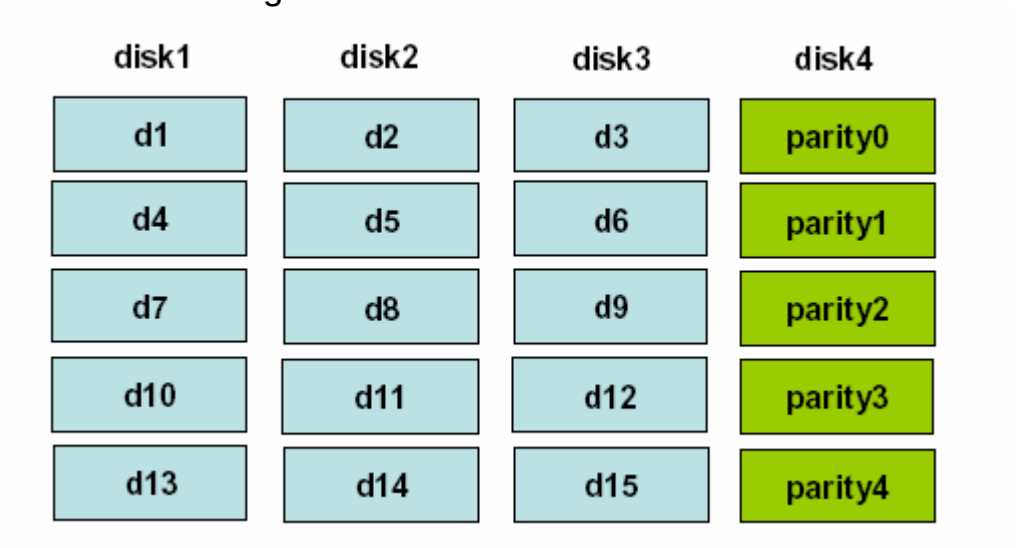

RAID 5 - Disk Striping with Parity

•Each entire data block is written on a data disk; parity for blocks in the same rank is generated on Writes, recorded in a distributed location and checked on Reads.

•RAID Level 5 requires a minimum of 3 drives to implement

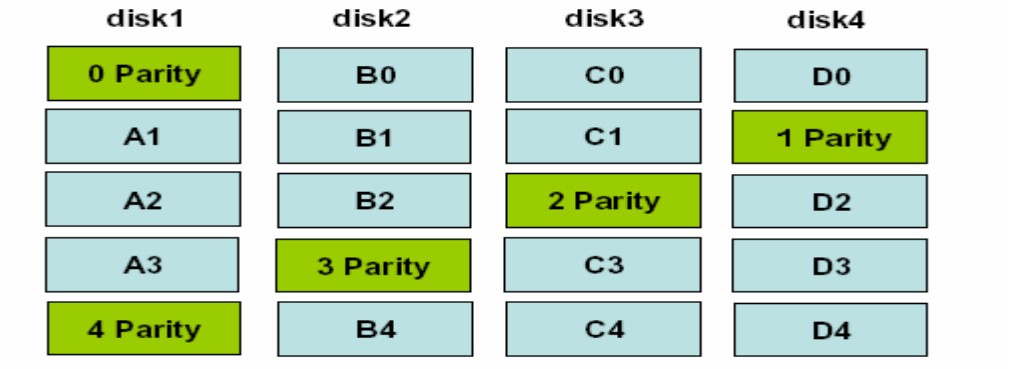

Storage Area Network (SAN)

• A high-speed sub network of shared storage devices.

• A storage device is a machine that contains nothing but a disk or disks for storing data.

• A SAN's architecture works in a way that makes all storage devices available to all servers on a LAN or WAN.