

Lesson 26, 27, 28

Objectives

- Disk Structure
- Disk Scheduling
- Disk management

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 cylinders from outermost to innermost.

The operating system is responsible for using hardware efficiently — for the disk drives, this means having a fast access time and disk bandwidth. Access time has two major components.

- *Seek time* is the time for the disk are to move the heads to the cylinder containing the desired sector.
- *Rotational latency* is the additional time waiting for the disk to rotate the desired sector to the disk head.

Criteria for disk management:

- Minimize seek time: $\text{Seek time} \approx \text{seek distance}$
- Disk bandwidth is the total number of bytes transferred, divided by the total time between the first request for service and the completion of the last transfer.


Disk Scheduling

Different algorithms are used for scheduling the disk head so that seek time may be mitigated.

First come first serve


Serve the requests in the order of their arrival.

Example:




Shortest Seek Time First (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.




Scan Algorithm (Elevator Algorithm)

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. Sometimes it is called the elevator algorithm. Illustration shows total head movement of 208 cylinders.




Circular Scan (C-Scan)

Provides a more uniform wait time than SCAN. The head moves from one end of the disk to the other servicing requests as it goes. When it reaches the other end, however, it immediately returns to the beginning of the disk, without servicing any requests on the return trip. Treats the cylinders as a circular list that wraps around from the last cylinder to the first one.



Circular Look (C-Look)

It's a 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.



Selection of the Algorithm

- SSTF is common and has a natural appeal
- SCAN and C-SCAN perform better for systems that place a heavy load on the disk.
- Performance depends on the number and types of requests.
- Requests for disk service can be influenced by the file allocation method.
- The disk-scheduling algorithm should be written as a separate module of the operating system, allowing it to be replaced with a different algorithm if necessary.
- Either SSTF or LOOK is a reasonable choice for the default algorithm.