OPERATING SYSTEMS · LECTURE 10

Storage Management

Primary vs secondary memory, mass storage, disk scheduling



Lecturer: Aidana Karibayeva

01 What is storage management?

Managing the devices that store user and system data, so that storage capacity is used well.

Key attributes

Performance	Reliability	Recoverability	Capacity

Advantages

- Simpler to manage capacity
- Saves time
- Better system performance
- More agility with virtualization and automation

Limitations

- Limited physical capacity; slower when full
- Complexity and cost
- Security; backup and recovery

02

Three phases of storage management

1 · INITIAL ALLOCATION

Each block is free or in use; free blocks are allocated during execution

→

2 · RECOVERY

Find objects that are no longer referenced and reclaim their memory

→

3 · COMPACTION & REUSE

Join small free pieces into large blocks

Partial compaction



Only **adjacent free blocks** are merged; active blocks stay in place.

Full compaction



All active blocks move to one end; **all pointers** must be updated.

03 Primary vs secondary memory

	PRIMARY (MAIN, INTERNAL)	SECONDARY (AUXILIARY, BACKUP)
Access	directly by the CPU	only through main memory
Speed	faster	slower
Cost	more expensive	cheaper
Capacity	about 16–32 GB	200 GB to terabytes
Volatility	volatile and non-volatile parts; data lost on power-off	non-volatile only; data kept on power-off
Examples	RAM, cache	HDD, SSD, flash, ROM, magnetic tape

04 Mass storage devices

Main purpose: **data backup and recovery**.

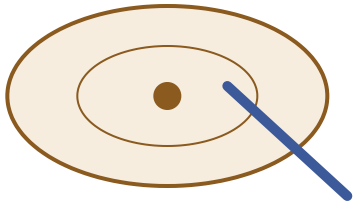

Magnetic disks

HDD, floppy, zip; written by magnetisation. First HDD: IBM 1956, 50 platters, only 5 MB

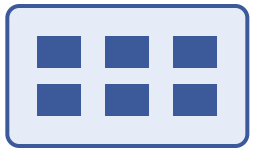

Solid state disks

Flash or battery-backed DRAM; no moving parts — much faster, but cost more, smaller, shorter life

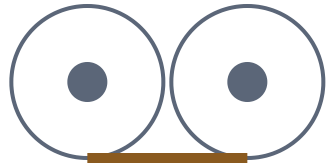

Magnetic tapes

Today mostly for backups; 20–200 GB, compression can double it

05 Disk structure

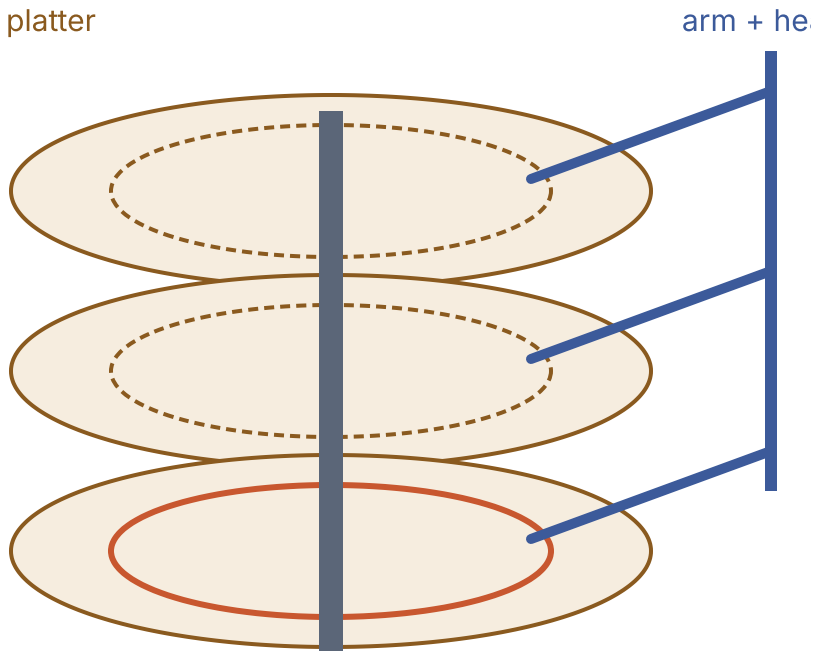

track · same track on all surfaces = cylinder

Track	concentric circle on a surface, divided into **sectors**
Cylinder	the same track on all surfaces
Block	smallest unit; disk = 1-D array of blocks

Capacity = heads × tracks per surface × sectors per track × bytes per sector

Seek time	move heads to the cylinder (slowest part)
Rotational latency	wait for the sector to reach the head
Transfer rate	speed of moving data to the computer

06 Connecting a disk

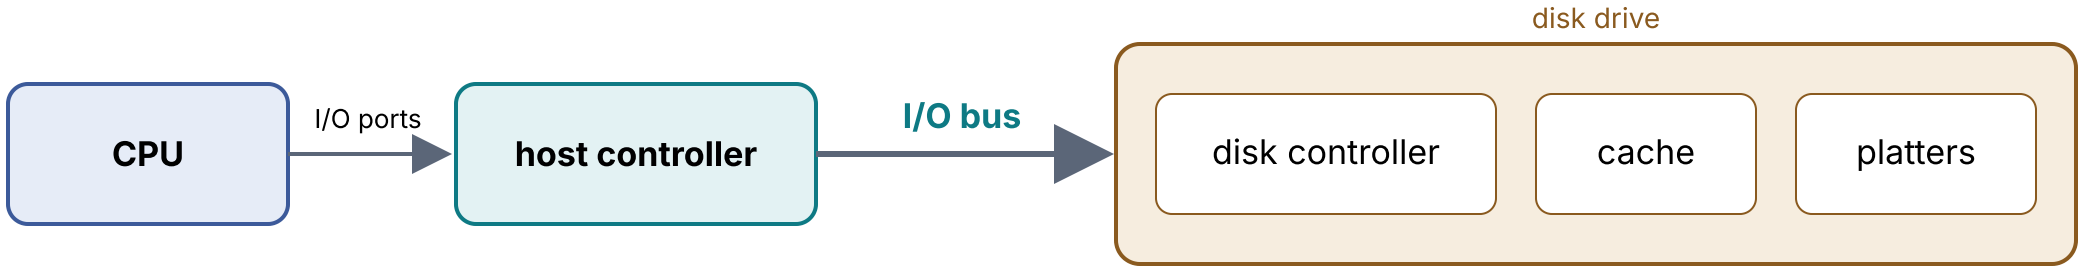

Common bus interfaces

`EIDE` `ATA` `SATA` `USB` `Fibre Channel` `SCSI`

> Heads fly on a thin air cushion; touching the surface causes a **head crash**.

07 Disk scheduling

The main task of secondary storage management: choose the order of disk requests to reduce **head movement**.

FCFS

in order of arrival

SSTF

closest request first

SCAN

to the disk end, then back

LOOK

to the last request, then back

Example for all algorithms

Head starts at 50; cylinders 0–99

`10` `95` `23` `78` `80`

08 FCFS and SSTF

FCFS — first come, first served

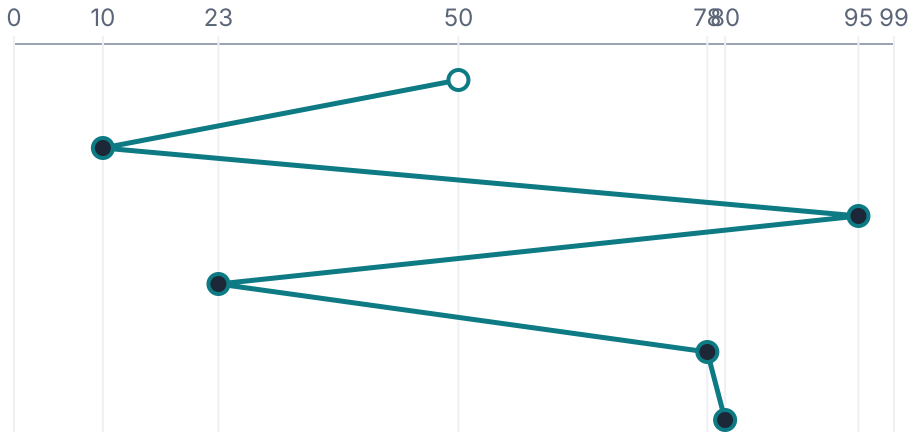

`50 → 10 → 95 → 23 → 78 → 80`

Fair, but not the fastest. Movement: **254** cylinders

SSTF — shortest seek time first

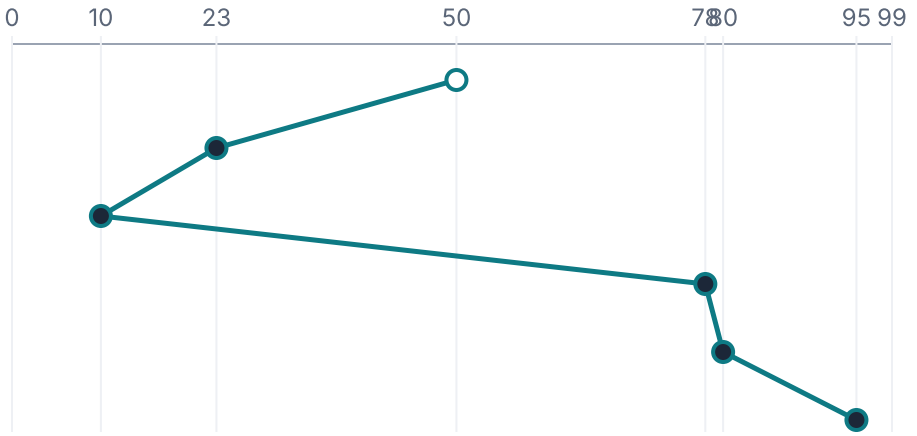

`50 → 23 → 10 → 78 → 80 → 95`

Nearest first; may **starve** far requests. Movement: **125**

09

SCAN and LOOK

SCAN — elevator

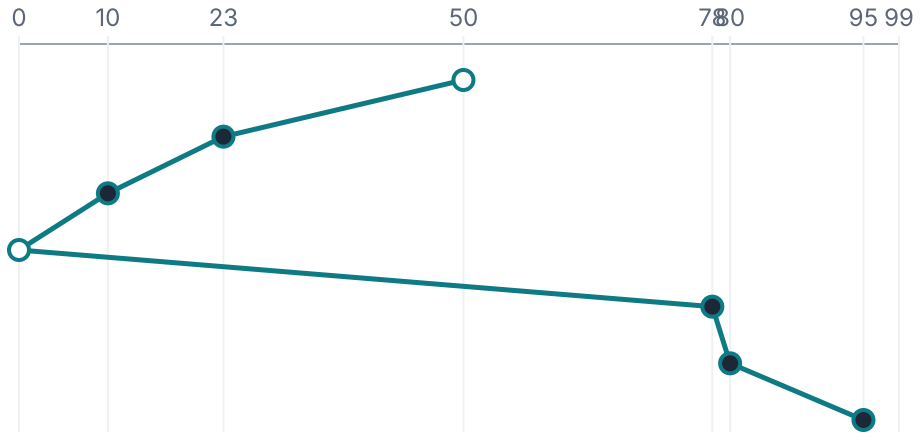

`50 → 23 → 10 → 0 → 78 → 80 → 95`

Goes to the **end of the disk**, then turns. Movement: **145**

LOOK — practical SCAN

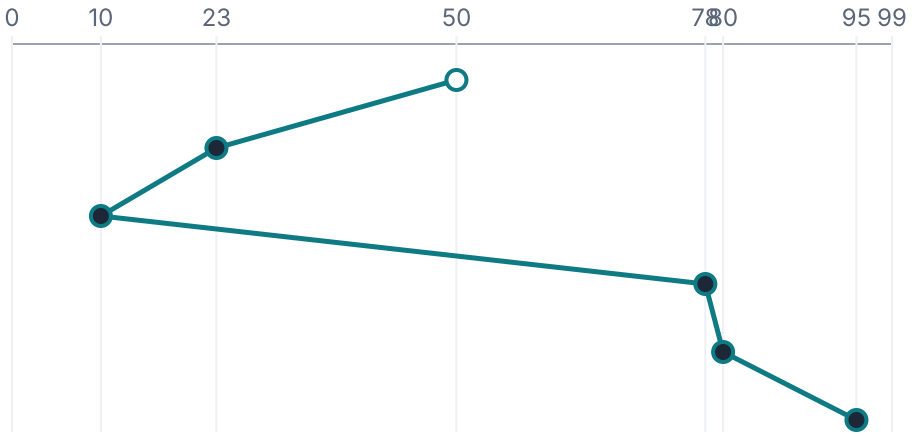

`50 → 23 → 10 → 78 → 80 → 95`

Turns at the **last request**. Movement: **125**

10 Comparison

ALGORITHM	ORDER SERVED	TOTAL HEAD MOVEMENT	NOTE
FCFS	`10, 95, 23, 78, 80`	**254**	simple and fair
SSTF	`23, 10, 78, 80, 95`	**125**	risk of starvation
SCAN	`23, 10, (0), 78, 80, 95`	**145**	goes to the disk end
LOOK	`23, 10, 78, 80, 95`	**125**	stops at the last request

Head starts at 50, moving towards 0 for SCAN and LOOK.