Chapter 12: Secondary-Storage Structure

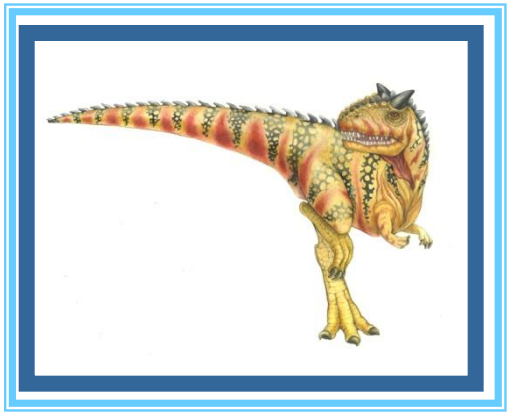

Chapter 12: Secondary-Storage Structure

- Overview of Mass Storage Structure
- Disk Structure
- Disk Attachment
- Disk Scheduling
- Disk Management

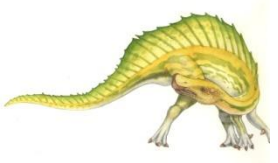

Objectives

- Describe the physical structure of secondary and tertiary storage devices and the resulting effects on the uses of the devices
- Discuss operating-system services provided for mass storage, including RAID and HSM

Overview of Mass Storage Structure

- Magnetic disks provide bulk of secondary storage of modern computers
  - Drives rotate at 60 to 200 times per second
  - Disk speed has two parts**:**
    - **Transfer rate** is rate at which data flow between drive and computer
    - **Positioning time** (**random-access time**):
      - **seek time** is time to move disk arm to desired cylinder .
      - **rotational latency** is time for desired sector to rotate under the disk head
- Disks can be removable, e.g. floppy disk , CD disk and DVD
- Disk Drive attached to computer via **I/O bus**
  - Busses vary, including **EIDE, ATA, SATA, USB, Fibre Channel, SCSI**
  - The data transfers on a bus are carried out by **controllers**
  - **Host controller** in computer uses bus to talk to **disk controller** built into drive or storage array

Moving-head Disk Mechanism

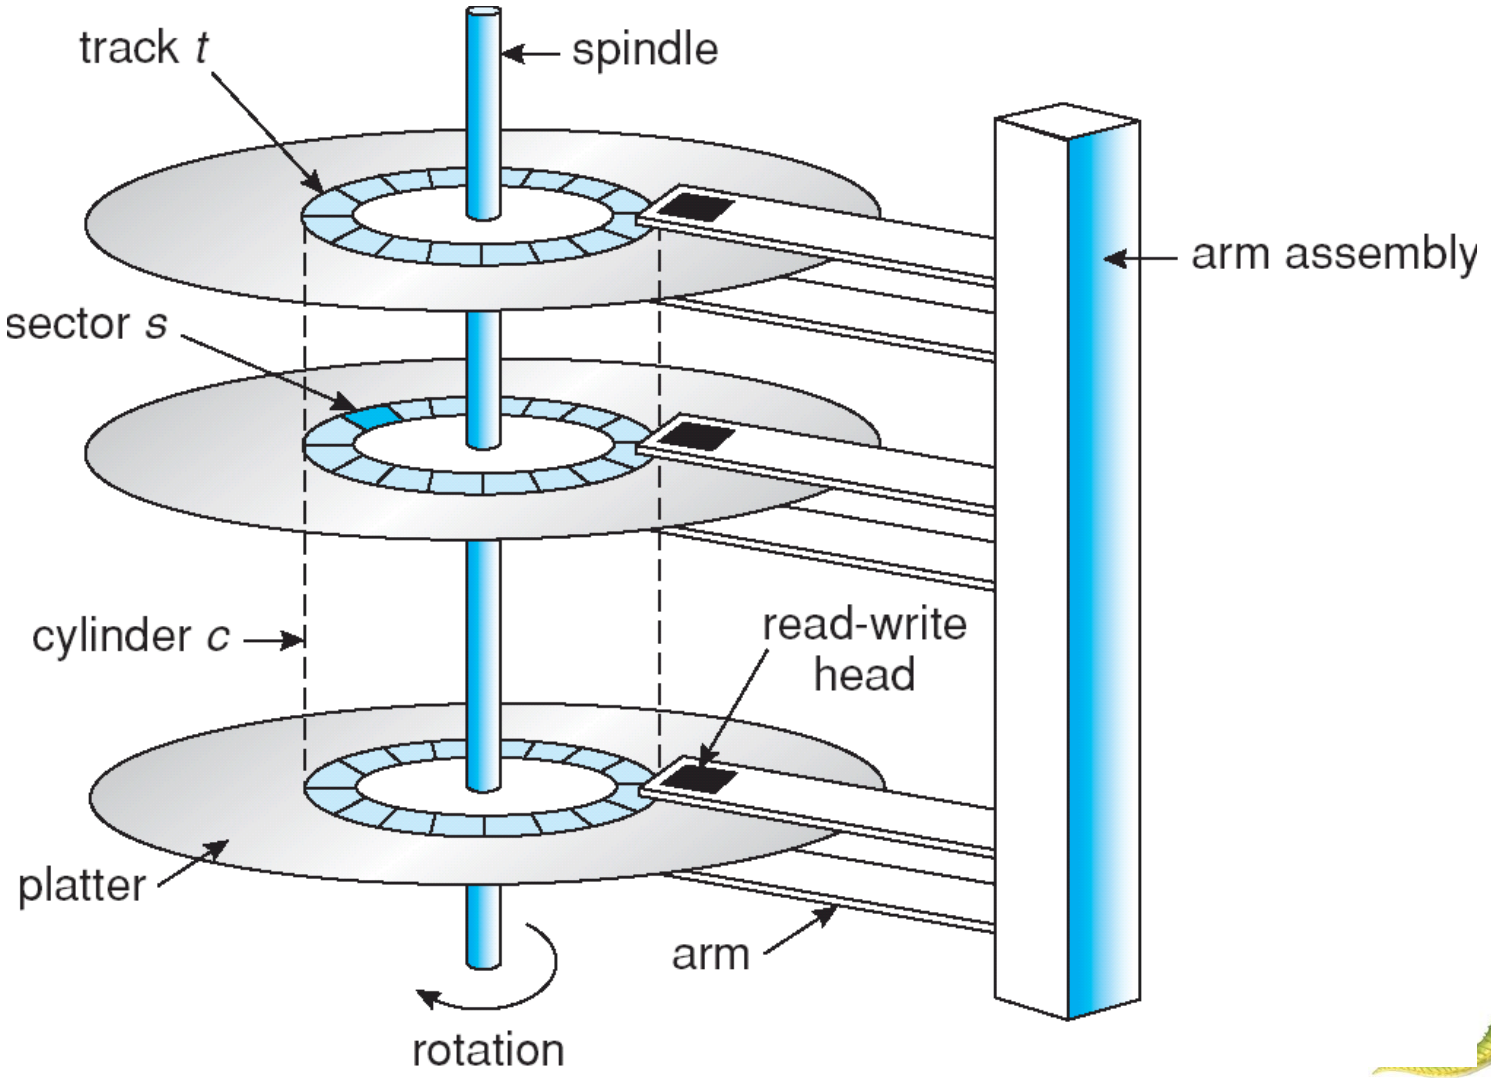

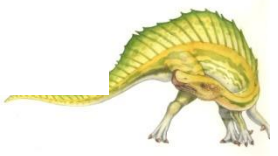

Overview of Mass Storage Structure (Cont.)

- Magnetic tape
  - Was early secondary-storage medium
  - Relatively permanent and holds large quantities of data
  - Access time slow
  - Random access ~1000 times slower than disk
  - Mainly used for backup, storage of infrequently-used data, transfer medium between systems
  - Kept in spool and wound or rewound past read-write head
  - Once data under head, transfer rates comparable to disk
  - 20-200GB typical storage

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 size of a logical block is usually 512 bytes, but some disks can be low-level-formatted to have a different logical block size , such as 1,024 bytes.
- 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.

Disk Attachment

- Computers access disk storage in two ways:
  1. Host-attached storage : via I/O ports talking to I/O busses, common on small systems.
  2. Network-attached storage: via a remote host in a distributed file system.

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.
- 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.
- 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.

- Improving both the access time and the bandwidth by managing the order in which disk I/O requests are serviced.

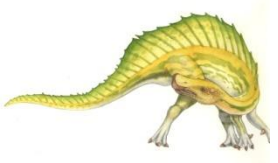

Disk Scheduling (cont.)

- Improving both the access time and the bandwidth by managing the order in which disk I/O requests are serviced.
- If process needs I/O to or from the disk, it issues a system call to OS. The request have the following information:
  - Is the operation input or output?
  - What is the disk address for the transfer ?
  - What is the memory address for the transfer ?
  - What is the number of sectors to be transferred?
  - If the desired disk drive and controller are
    - available, the request can be serviced immediately.
    - busy , the request waiting in the disk queue.

Disk Scheduling (Cont.)

- In multiprogramming system , the disk queue have several pending requests. When one is completed, the OS choose which request is the next?
- Several algorithms exist to schedule the servicing of disk I/O requests.
- We illustrate them with a request queue :

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

Head pointer 53

FCFS

Illustration shows total head movement of 640 cylinders.

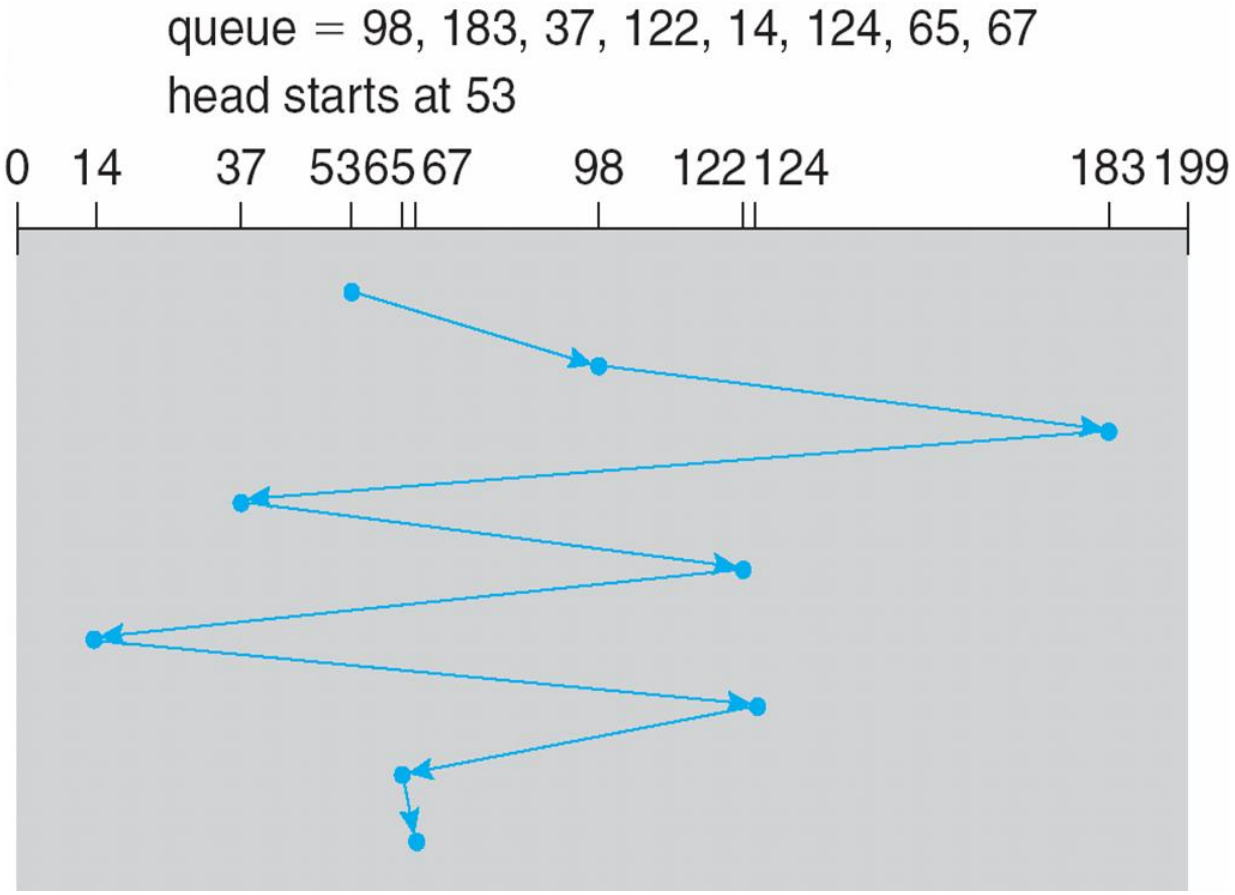

SSTF

- **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.

SSTF (Cont.)

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

head starts at 53

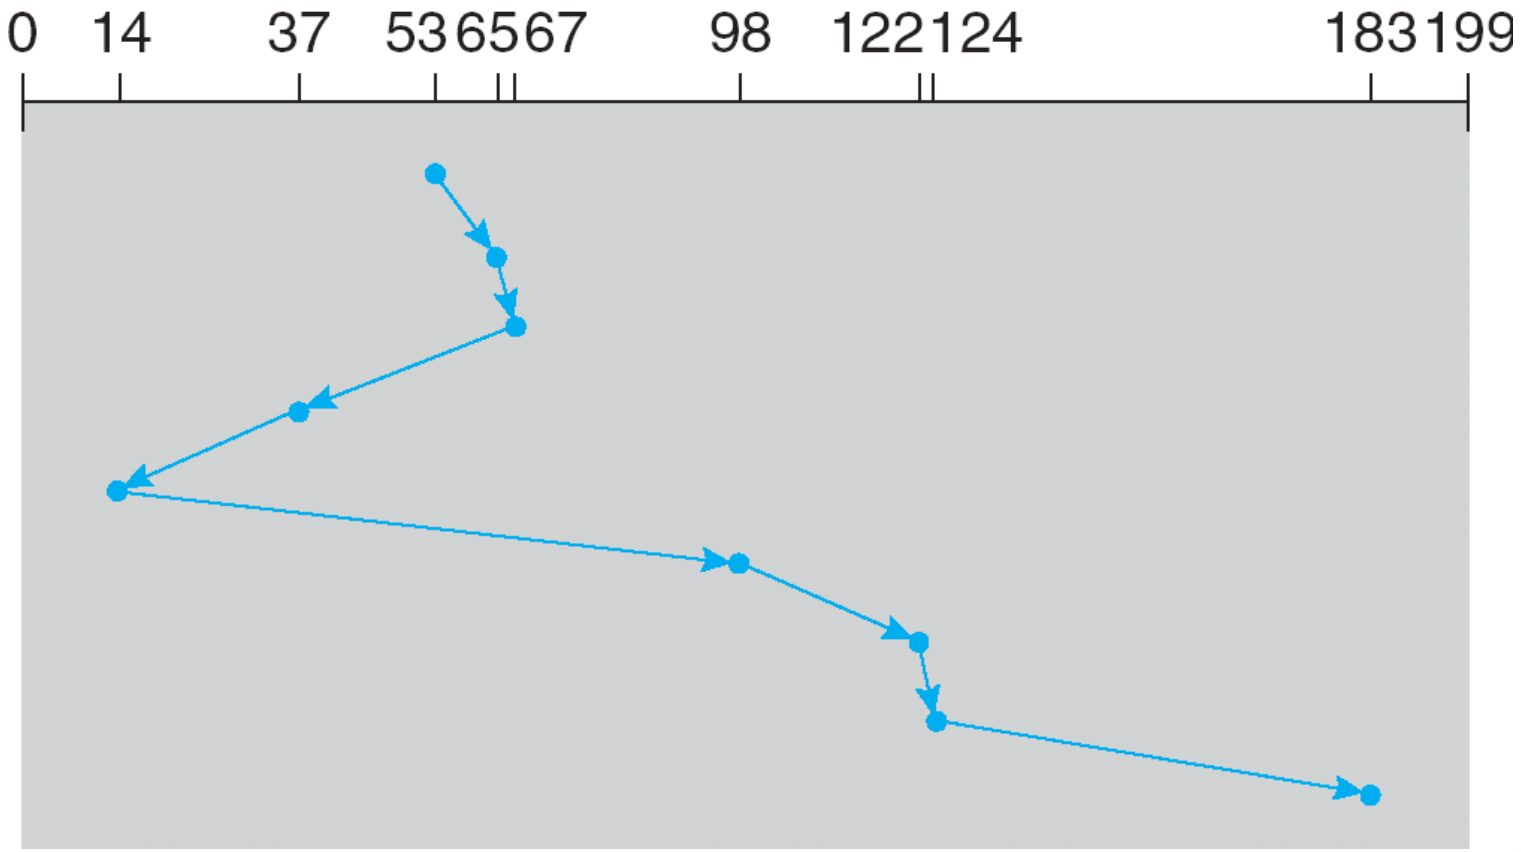

SCAN

- 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.
- Sometimes called the *elevator algorithm*.
- We need to know the direction of head movement in addition to the head's current position.
- Illustration shows total head movement of 208 cylinders.

SCAN (Cont.)

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

head starts at 53

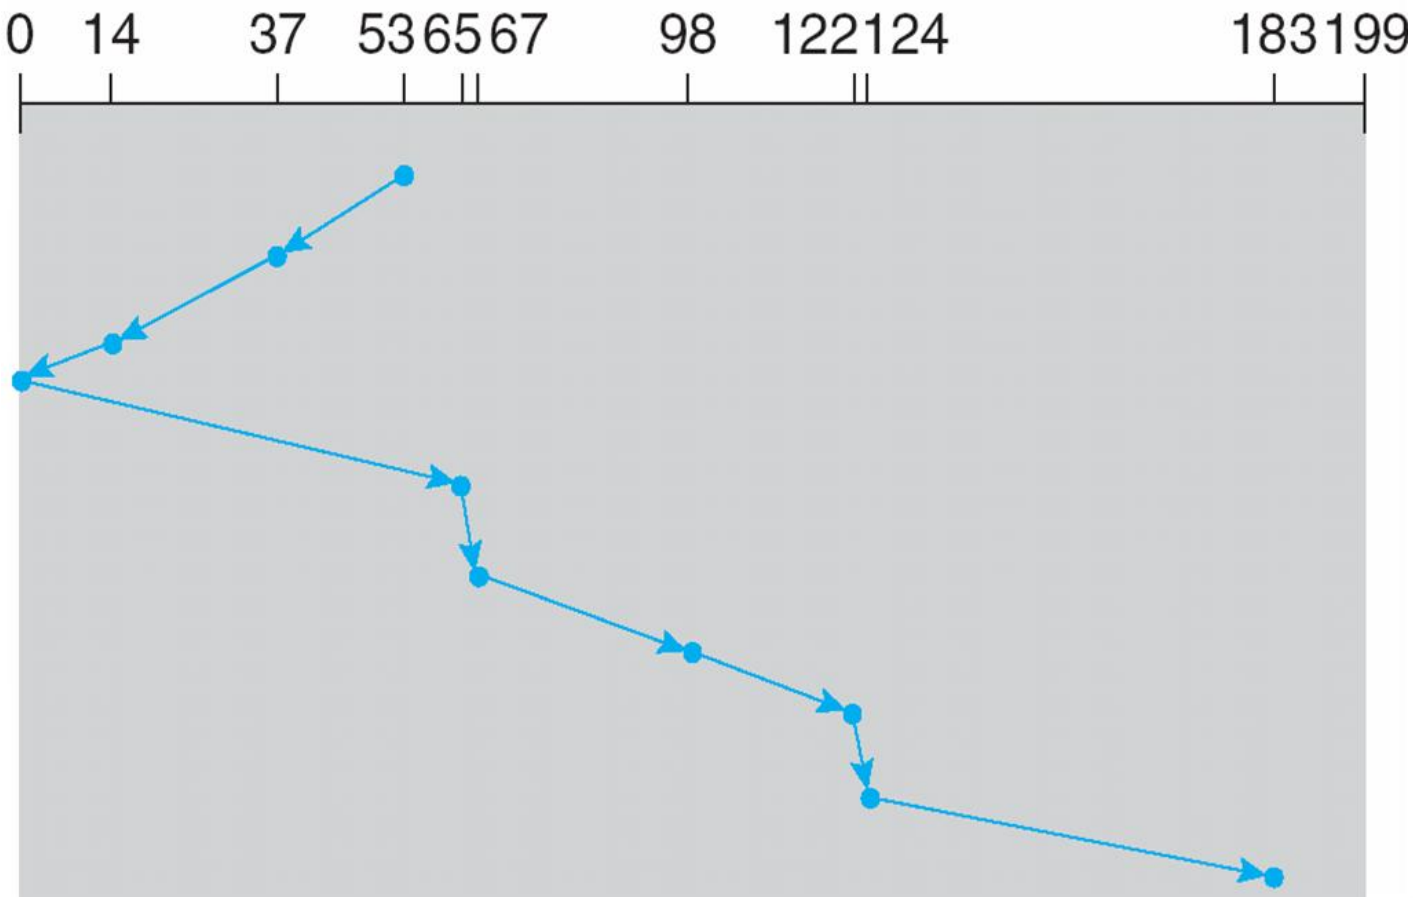

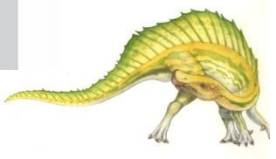

C-SCAN

- **Circular SCAN (C-SCAN)** scheduling is a variant of SCAN designed to 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.

C-SCAN (Cont.)

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

head starts at 53

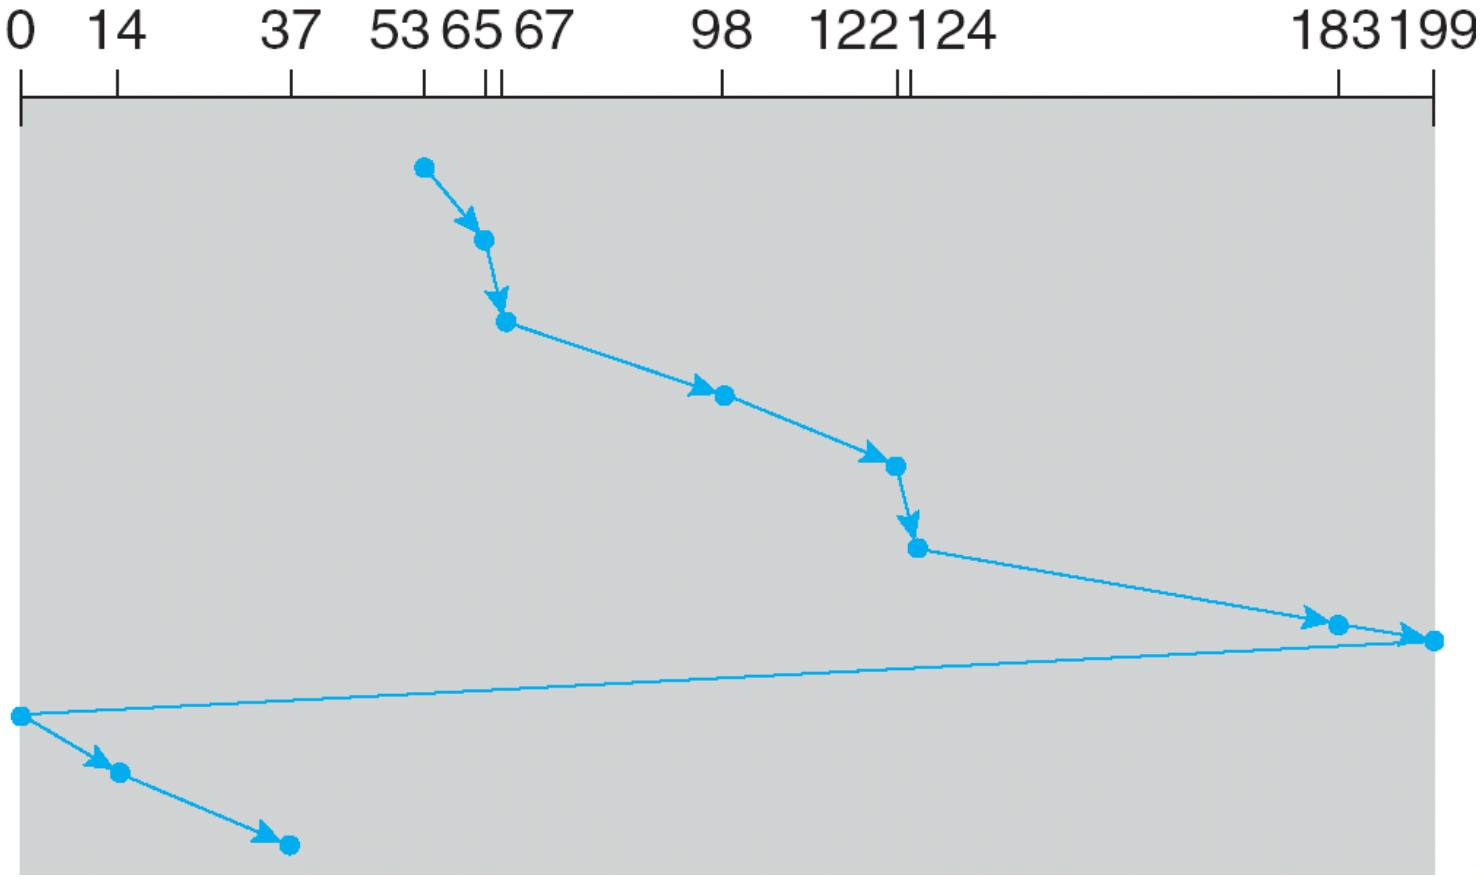

C-LOOK

- 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.

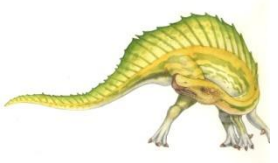

C-LOOK (Cont.)

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

head starts at 53

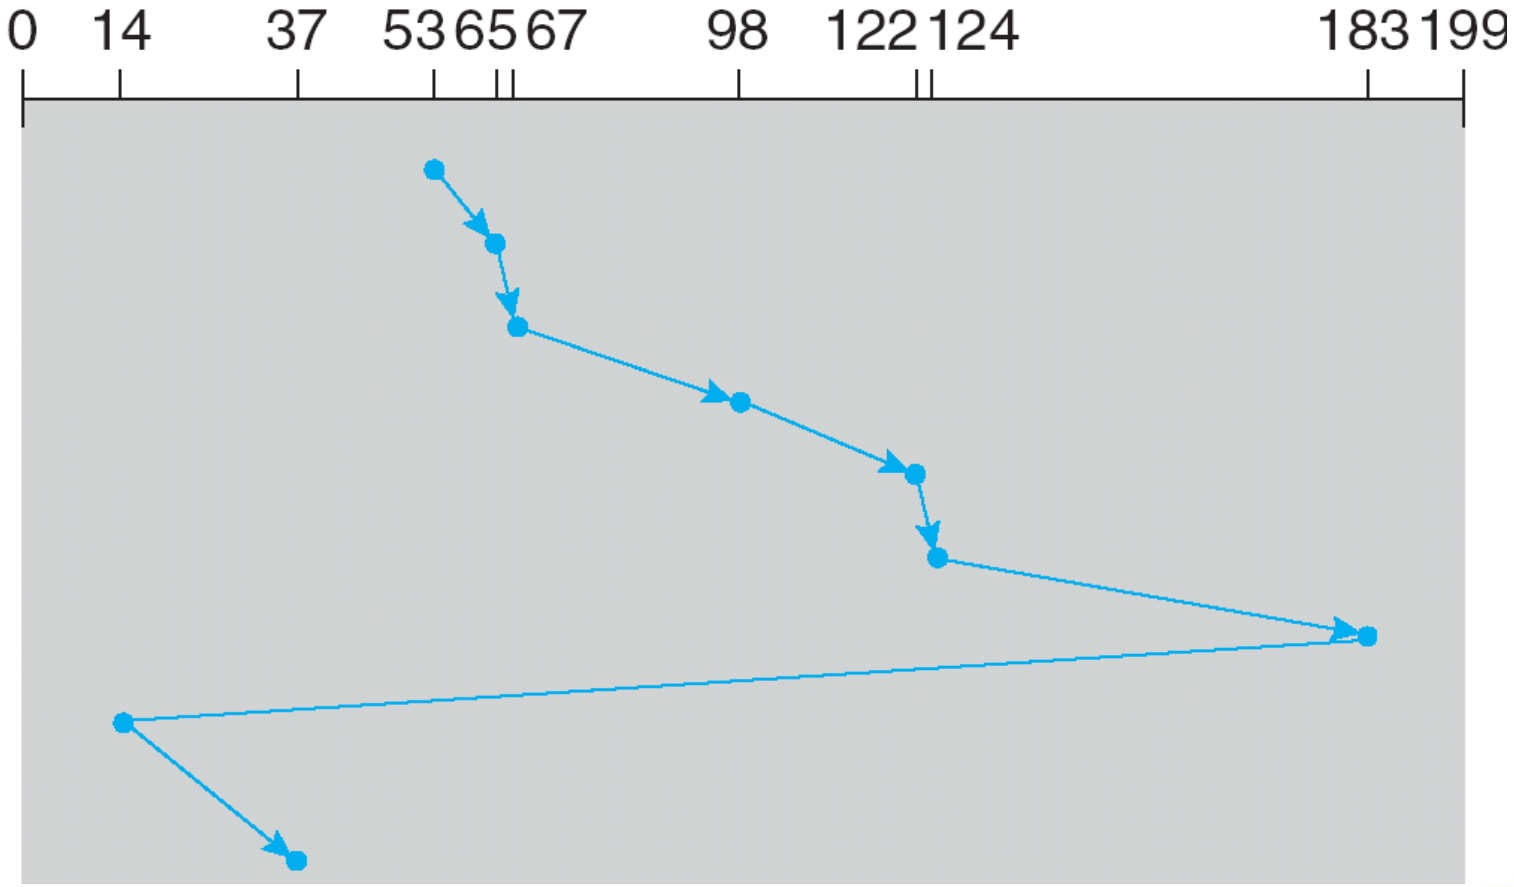

Selecting a Disk-Scheduling Algorithm

- SSTF is common and has a natural appeal, because it increase performance over FCFS.
- SCAN and C-SCAN perform better for systems that place a heavy load on the disk, because they cause less starvation problem.
- Performance depends on the number and types of requests.
- 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.

Disk Management

Disk Formatting:

- *Low-level formatting*, or *physical formatting*
  - Dividing a disk into sectors that the disk controller can read and write.
  - Filling the disk with a special data structure for each sector, which is consists of :
    - Header and Trailer : contain information used by the disk controller
    - Data area (512 bytes in size)

- To use a disk to hold files, the operating system still needs to record its own data structures on the disk.
  - *Partition* the disk into one or more groups of cylinders.
  - *Logical formatting* or “making a file system” that means OS stores the initial file system data structure onto the disk e.g. initial empty directory.

End of Chapter 12