CHAPTER 6

STORAGE MANEGEMENT

Introduction

- The computer system must provide secondary storage to back up main memory.
- Modern computer systems use disks as the primary on-line storage medium for information
- The **file system** provides the mechanism for on-line storage of and access to both data and programs residing on the disks.
- A **file** is a collection of related information defined by its creator.
  - Is mapped by the operating system onto physical devices.
  - Is normally organized into directories for ease of use.

Storage Devices

- Have vary in many aspects.
  - Some devices transfer a character or a block of characters at a time.
  - Some can be accessed only sequentially, others randomly.
  - Some transfer data synchronously, others asynchronously.
  - Some are dedicated, some shared.
  - They can be read-only or read–write.
  - They vary greatly in speed.
  - they are also the slowest major component of the computer.

Magnetic Disks

- provide the bulk of secondary storage for modern computer systems.
- Each disk platter has a flat circular shape, like a CD.
- Common platter diameters range from 1.8 to 3.5 inches.
- The two surfaces of a platter are covered with a magnetic material.
- We store information by recording it magnetically on the platters.

Magnetic Disks

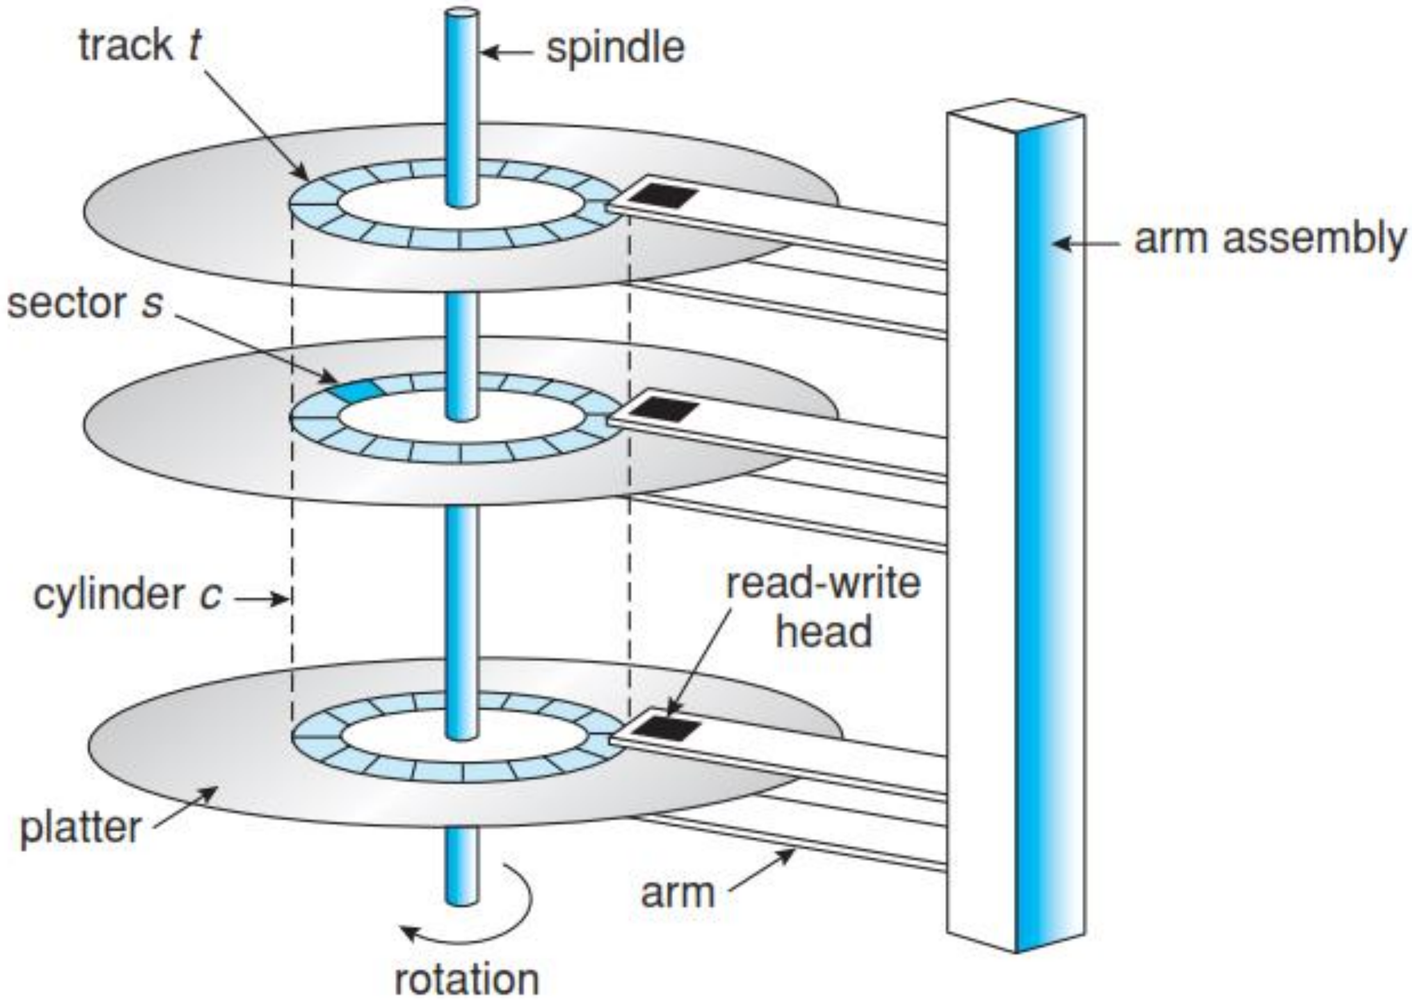

Magnetic Disks

- A read–write head “flies” just above each surface of every platter.
- The heads are attached to a disk arm that moves all the heads as a unit.
- The surface of a platter is logically divided into circular **tracks**, which are subdivided into **sectors**.
- The set of tracks that are at one arm position makes up a cylinder.
- There may be thousands of concentric cylinders in a disk drive,
- each track may contain hundreds of sectors.

Magnetic Disks

- Disk speed has two parts.
  - The **transfer rate** is the rate at which data flow between the drive and the computer.
  - The **positioning time**, or random-access time, consists of two parts:
    - the time necessary to move the disk arm to the desired cylinder, called the **seek time**,
    - and the time necessary for the desired sector to rotate to the disk head, called **the rotational latency**.
- Typical disks can transfer several megabytes of data per second, and they have seek times and rotational latencies of several milliseconds.

Magnetic Disks

- head crash is the head will sometimes damage the magnetic surface.
- Because the disk head flies on an extremely thin cushion of air (measured in microns),
- Although the disk platters are coated with a thin protective layer,.
- A head crash normally cannot be repaired; the entire disk must be replaced.

Disk Structure

- Modern magnetic disk drives are addressed as large one-dimensional arrays of logical blocks,
- logical block is the smallest unit of transfer.
  - The size of a logical block is usually 512 bytes,
- logical blocks is mapped onto the sectors of the disk sequentially.
  - Sector 0 is the first sector of the first track on the outermost cylinder.
  - The mapping proceeds in order through that track, then through the rest of the tracks in that cylinder,
  - and then through the rest of the cylinders from outermost to innermost.

Disk Scheduling

- One of the responsibilities of the operating system is to use the hardware efficiently.
- meeting this responsibility having fast access time and large disk bandwidth.
- The 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.
- We can improve both the access time and the bandwidth by managing the order in which disk I/O requests are serviced.

Disk Scheduling

- Whenever a process needs I/O to or from the disk, The request specifies several pieces of information:
  - Whether this operation is input or output
  - What the disk address for the transfer is
  - What the memory address for the transfer is
  - What the number of sectors to be transferred is

Disk Scheduling

- For a multiprogramming system with many processes,
  - the disk queue may often have several pending requests.
  - Thus, when one request is completed, the operating system chooses which pending request to service next.
  - How does the operating system make this choice? Any one of several disk-scheduling algorithms can be used

FCFS Scheduling

- the first-come, first-served (FCFS) algorithm is The simplest way.
- example, a disk queue with requests for I/O to blocks on cylinders
  - 98, 183, 37, 122, 14, 124, 65, 67,
- If the disk head is initially at cylinder 53,
- it will first move from 53 to 98, then to 183, 37, 122, 14, 124, 65, and finally to 67, for a total head movement of 640 cylinders.

FCFS Scheduling

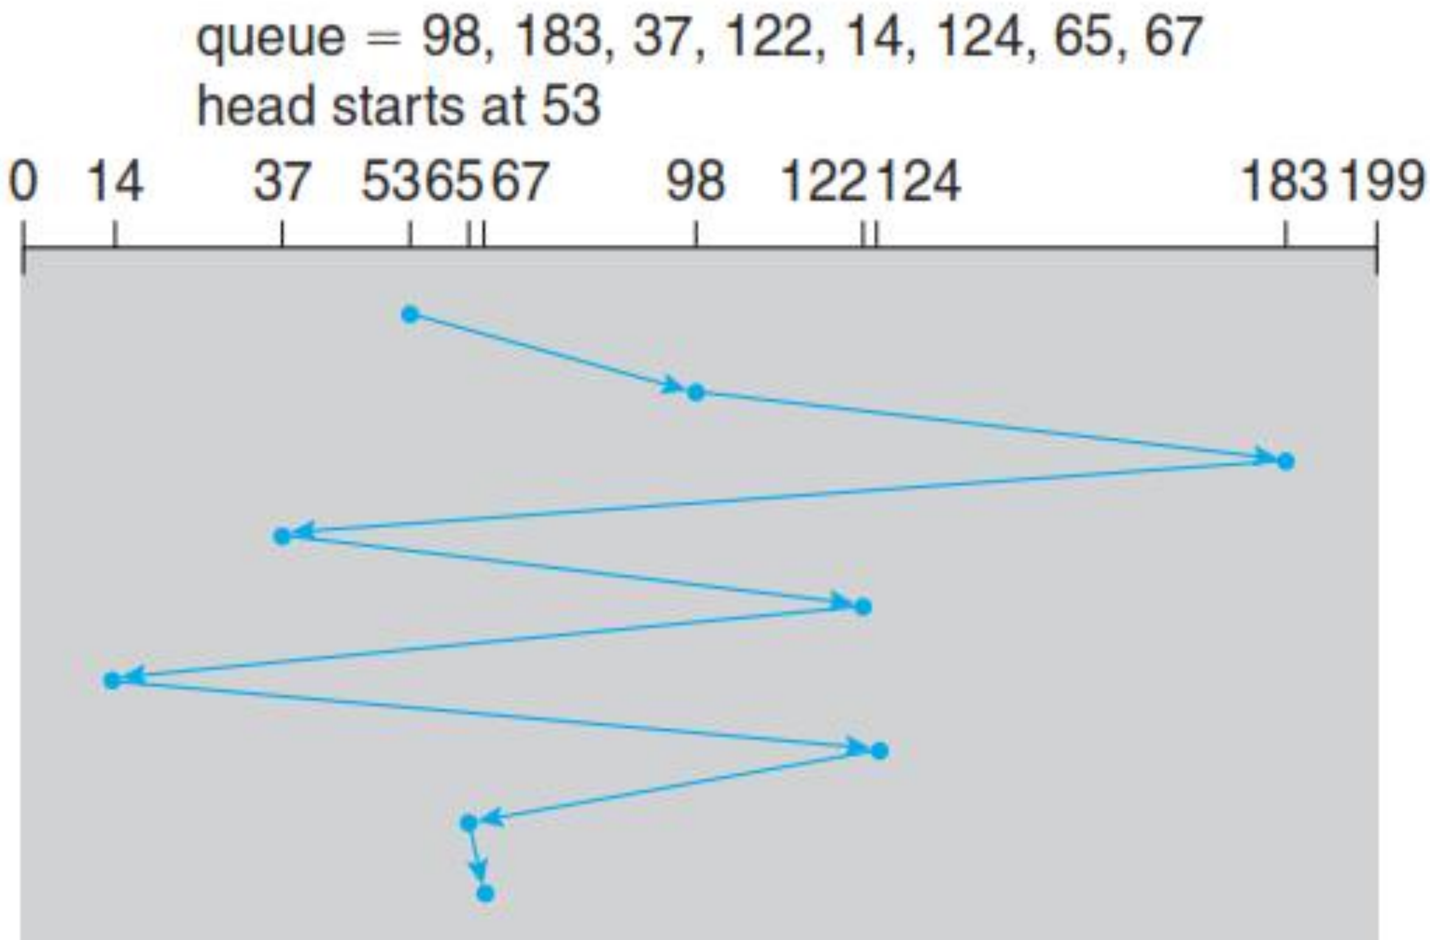

FCFS Scheduling

- The wild swing from 122 to 14 and then back to 124 illustrates the problem with this schedule.
- If the requests for cylinders 37 and 14 could be serviced together, before or after the requests for 122 and 124,
- the total head movement could be decreased
- This algorithm is intrinsically fair, but it generally does not provide the fastest service.

SSTF Scheduling

- shortest-seek-time-first (SSTF) algorithm selects the request with the least seek time from the current head position.
- In other words, SSTF chooses the pending request closest to the current head position.
- Previous example initial head position (53) closed at cylinder 65.
- Once we are at cylinder 65, the next closest request is at cylinder 67.
- Continuing, we service the request at cylinder 14, then 98, 122, 124, and finally 183.

SSTF Scheduling

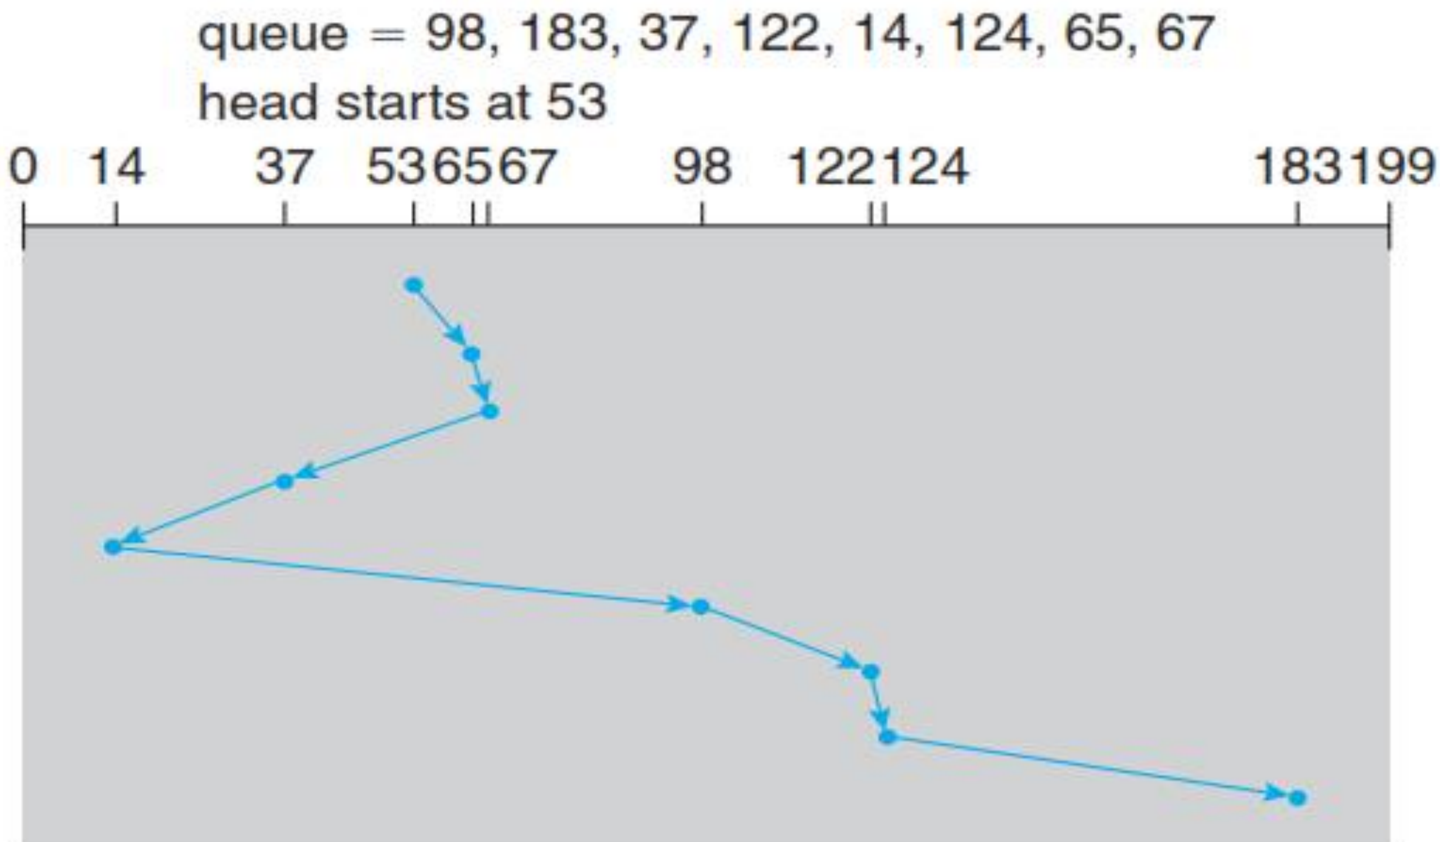

This scheduling method results in a total head movement of only 236 cylinders

SSTF Scheduling

- Although the SSTF algorithm is a substantial improvement over the FCFS algorithm, it is not optimal.
- In the example, we can do better by moving the head from 53 to 37, even though the latter is not closest, and then to 14, before turning around to service 65, 67, 98, 122, 124, and 183.
- This strategy reduces the total head movement to 208 cylinders.

SCAN Scheduling

- the disk arm starts at one end of the disk and moves toward the other end,
- servicing requests as it reaches each cylinder, until it gets to the other end of the disk.
- At the other end, the direction of head movement is reversed, and servicing continues.
- The SCAN algorithm is sometimes called the elevator algorithm

SCAN Scheduling

- Previous example:
  - we need to know the direction of head movement in addition to the head's current position.
  - Assuming that the disk arm is moving toward 0 and that the initial head position is again 53,
  - the head will next service 37 and then 14. At cylinder 0, the arm will reverse and will move toward the other end of the disk, servicing the requests at 65, 67, 98, 122, 124, and 183.

SCAN Scheduling

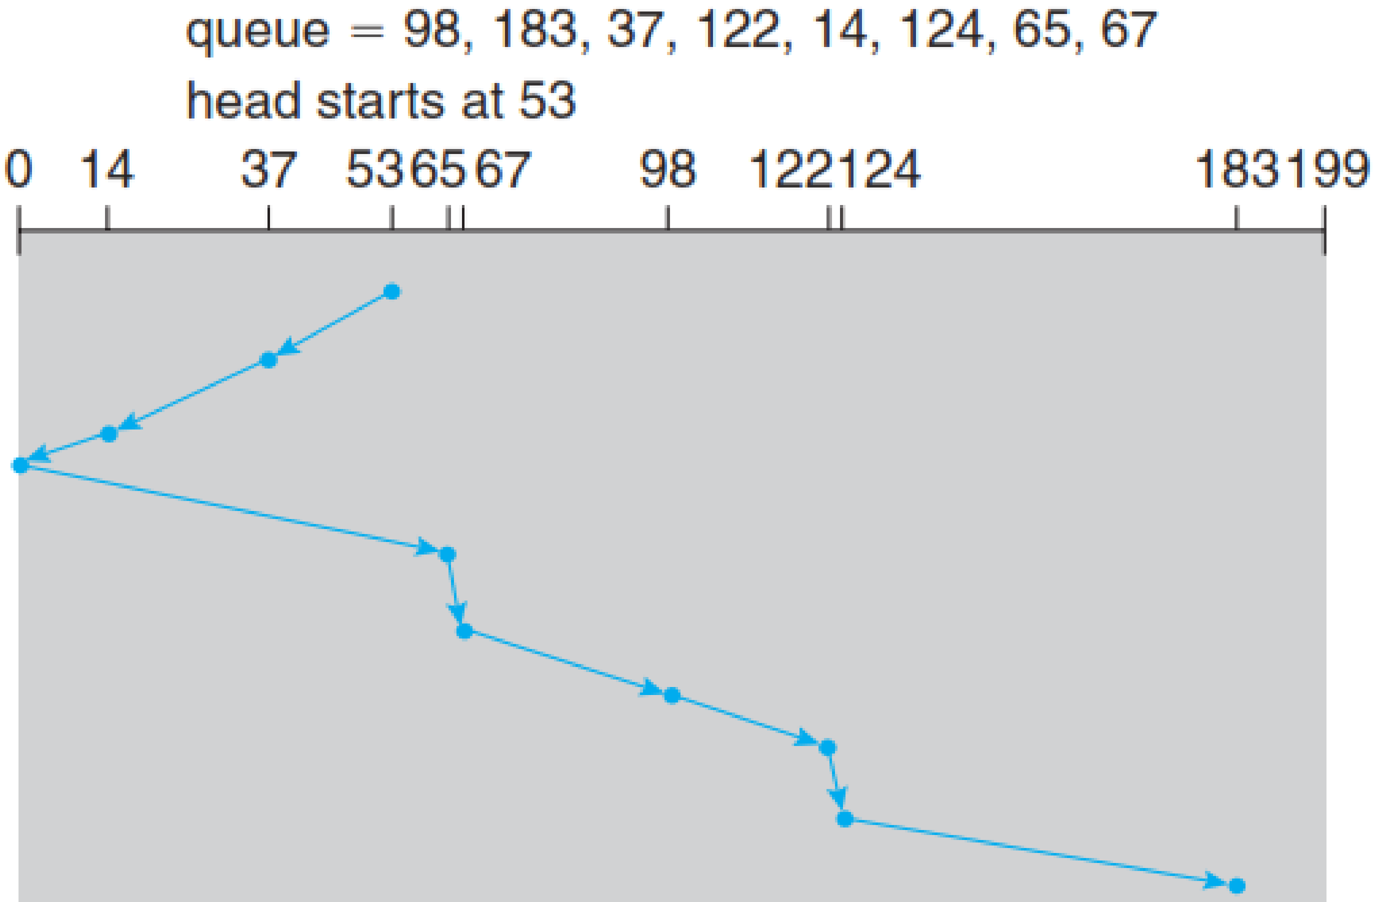

C-SCAN Scheduling

- Circular SCAN (C-SCAN) scheduling is Like SCAN, C-SCAN moves the head from one end of the disk to the other, servicing requests along the way.
  - When the head reaches the other end, however, it immediately returns to the beginning of the disk without servicing any requests on the return trip

C-SCAN Scheduling

queue = 98, 183, 37, 122, 14, 124, 65, 67
head starts at 53

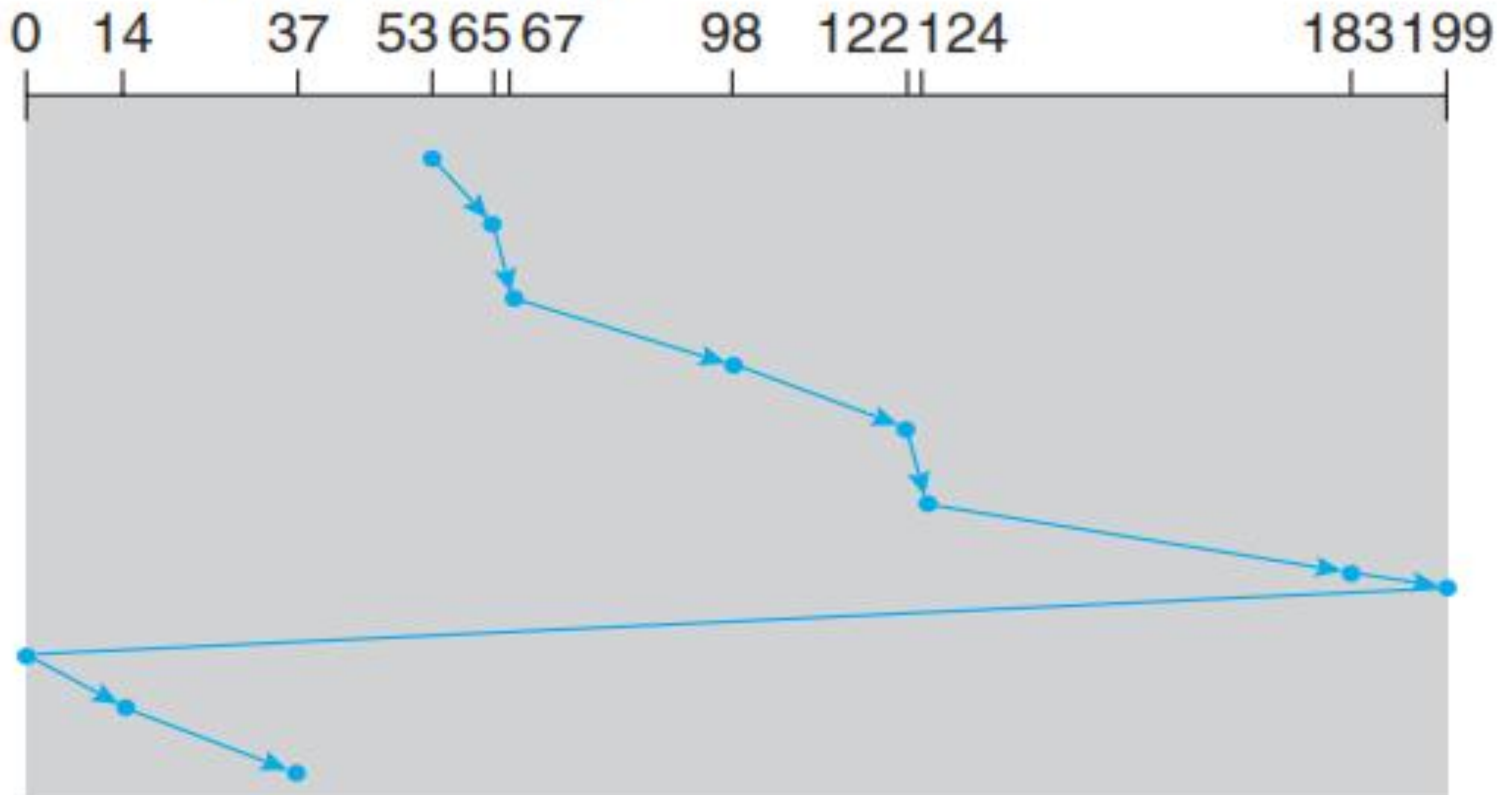

LOOK Scheduling

- the arm goes only as far as the final request in each direction.
- Then, it reverses direction immediately, without going all the way to the end of the disk.
- Versions of SCAN and C-SCAN that follow this pattern are called LOOK and C-LOOK scheduling,
  - because they look for a request before continuing to move in a given direction

C-LOOK Scheduling

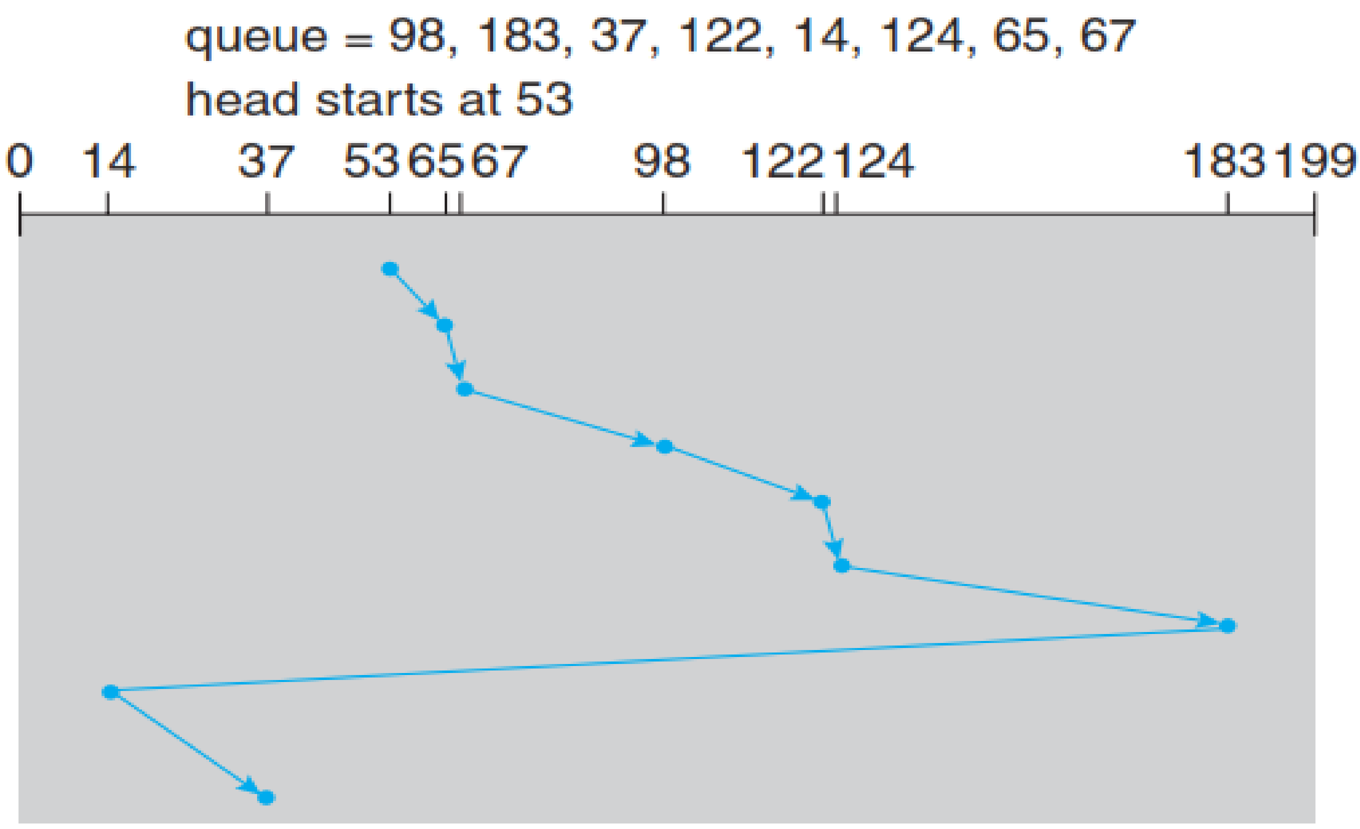

Selection of a Disk-Scheduling Algorithm

- SSTF is common and has a natural appeal because it increases performance over FCFS.
- SCAN and C-SCAN perform better for systems that place a heavy load on the disk, because they are less likely to cause a starvation problem.
- Optimal disk-scheduling can be replaced with a different algorithm if necessary.
  - Because of the complexities, the disk-scheduling algorithm should be written as a separate module of the operating system.
- Either SSTF or LOOK is a reasonable choice for the default algorithm.

References

- “ORGANIZATIONAL BEHAVIOR”(collected and edited by prof. dr. Štefan Ivanko) UNIVERSITY OF LJUBLJANA FACULTY OF PUBLIC ADMINISTRATION