

COMP 7970 Storage Systems

Magnetic Disks

Dr. Xiao Qin

Department of Computer Science and

Software Engineering

Auburn University

<http://www.eng.auburn.edu/~xqin>

xqin@auburn.edu

Devices: Magnetic Disks



- Purpose:
 - Long-term, nonvolatile storage
 - Large, inexpensive, slow level in the storage hierarchy
- Characteristics:
 - Seek Time (~ 8 ms avg)
 - positional latency
 - rotational latency
- Transfer rate
 - About a sector per ms (5-15 MB/s)
 - Blocks
- Capacity
 - Gigabytes
 - Quadruples every 3 years

Devices: Magnetic Disks

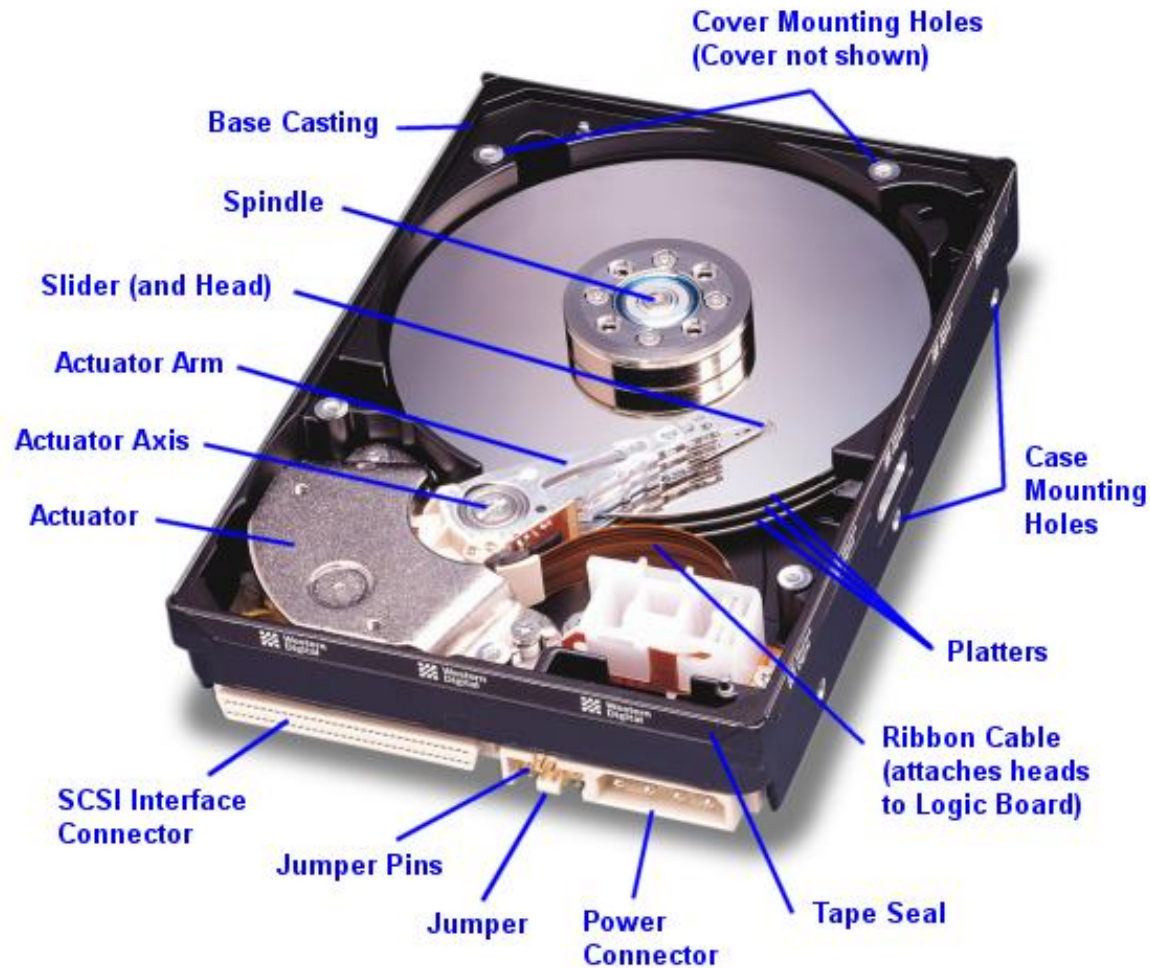
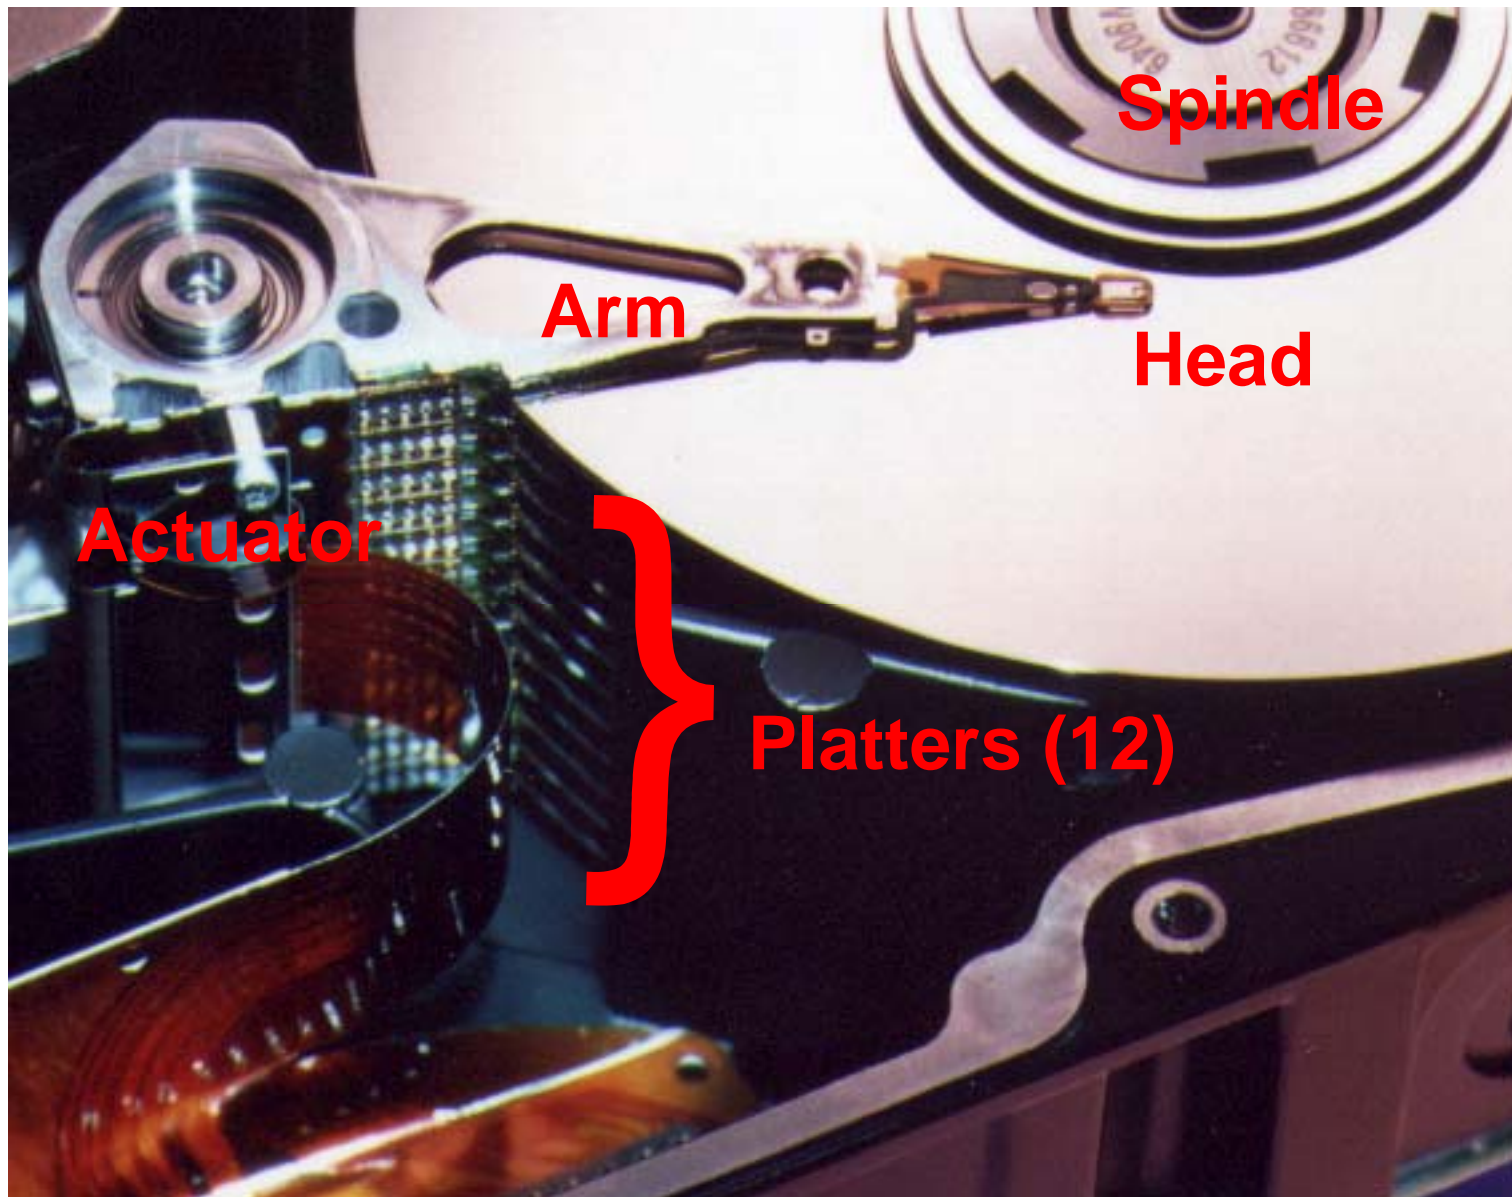
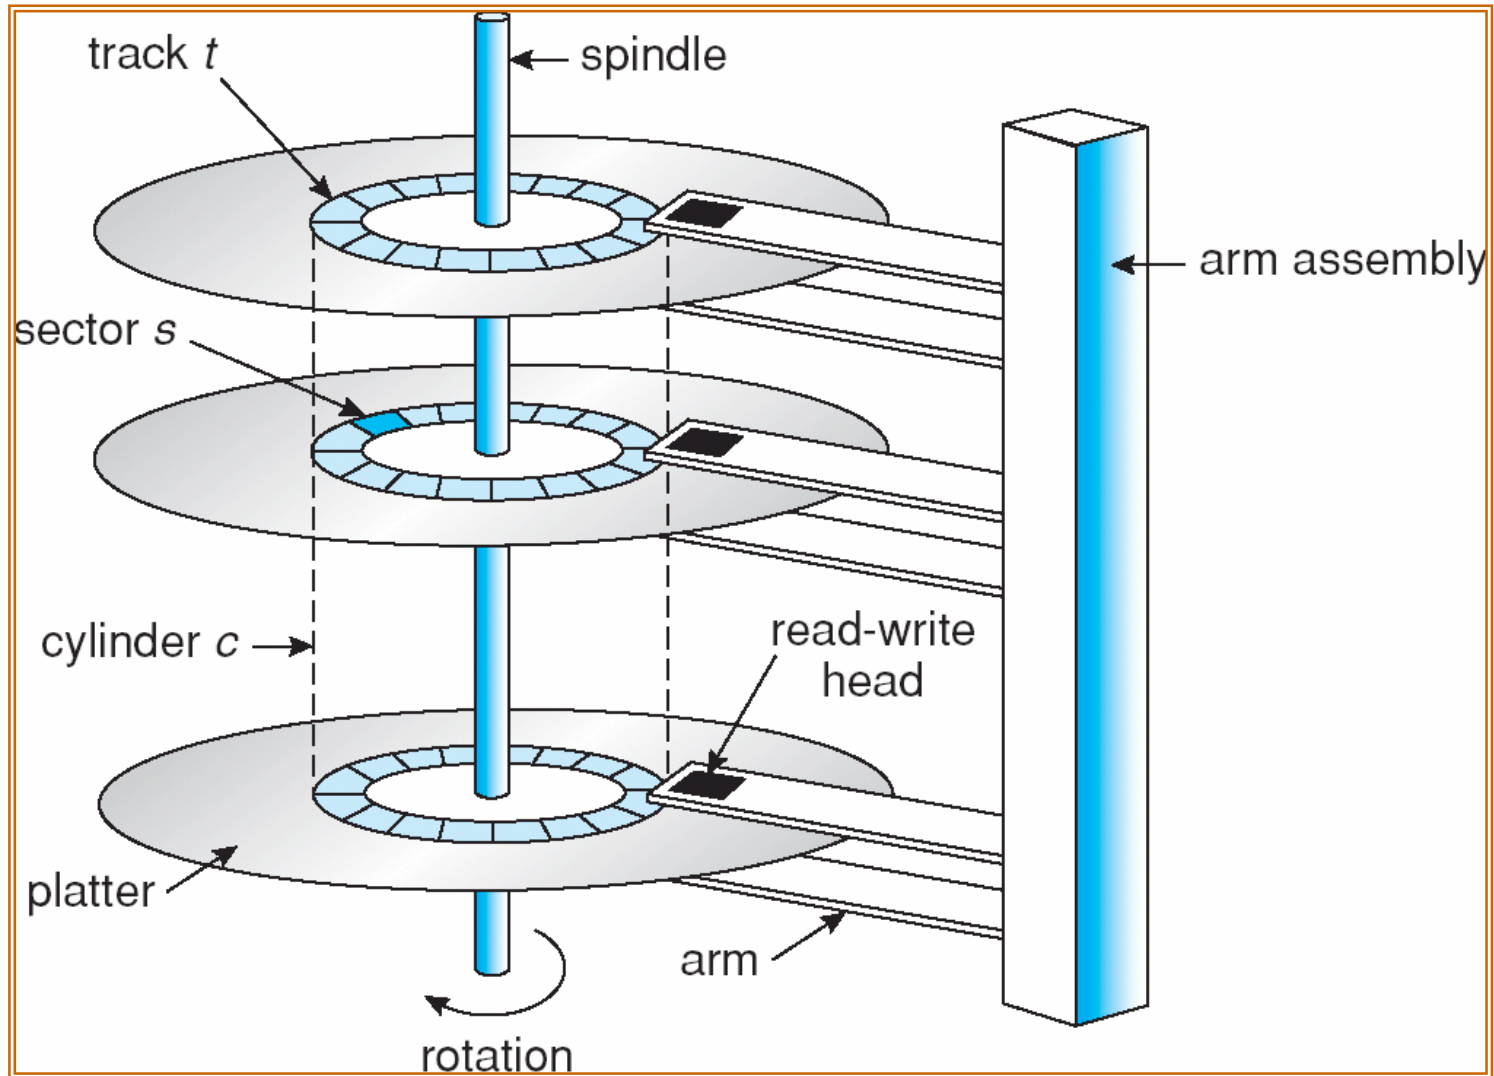


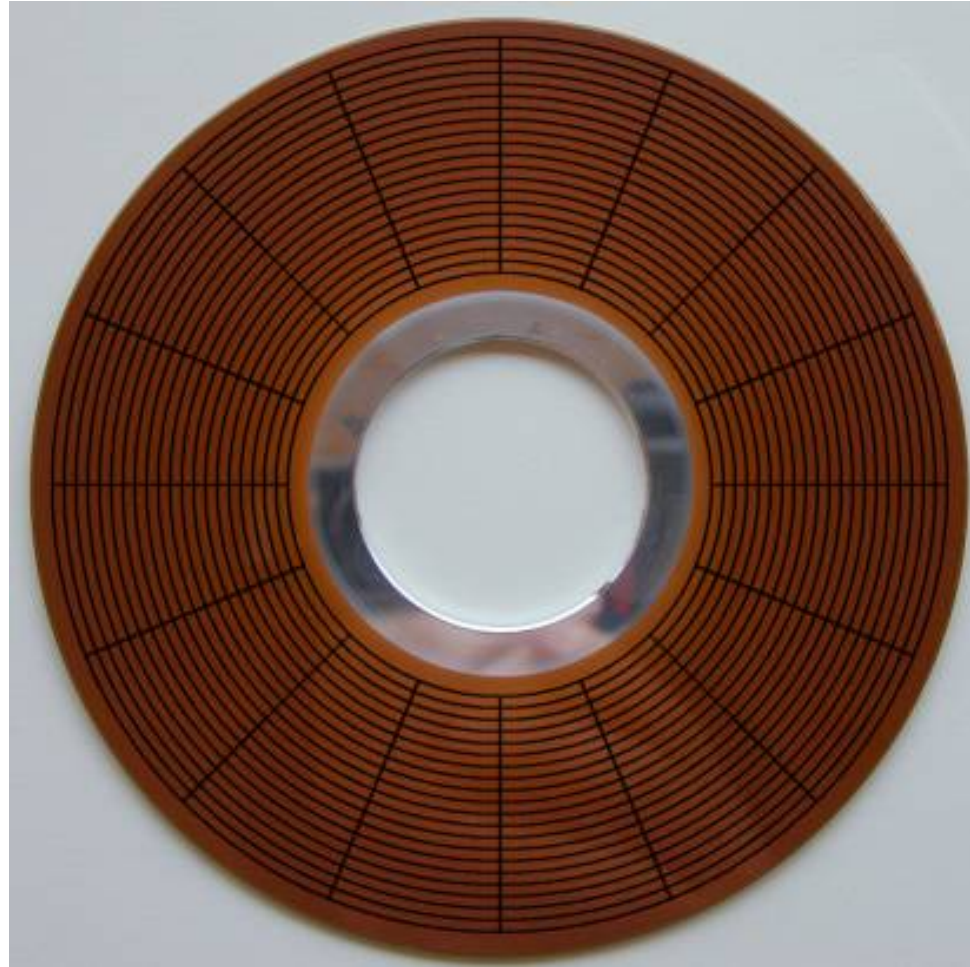
Photo of Disk Head, Arm, Actuator



Moving-head Disk Mechanism



Devices: Magnetic Disks



How many cylinders? 16,383 Cylinders

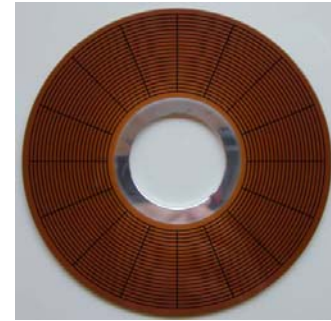
Magnetic Tape

- 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
- Common technologies are 4mm, 8mm, 19mm, LTO-2 and SDLT



IBM 2401 magnetic tape unit

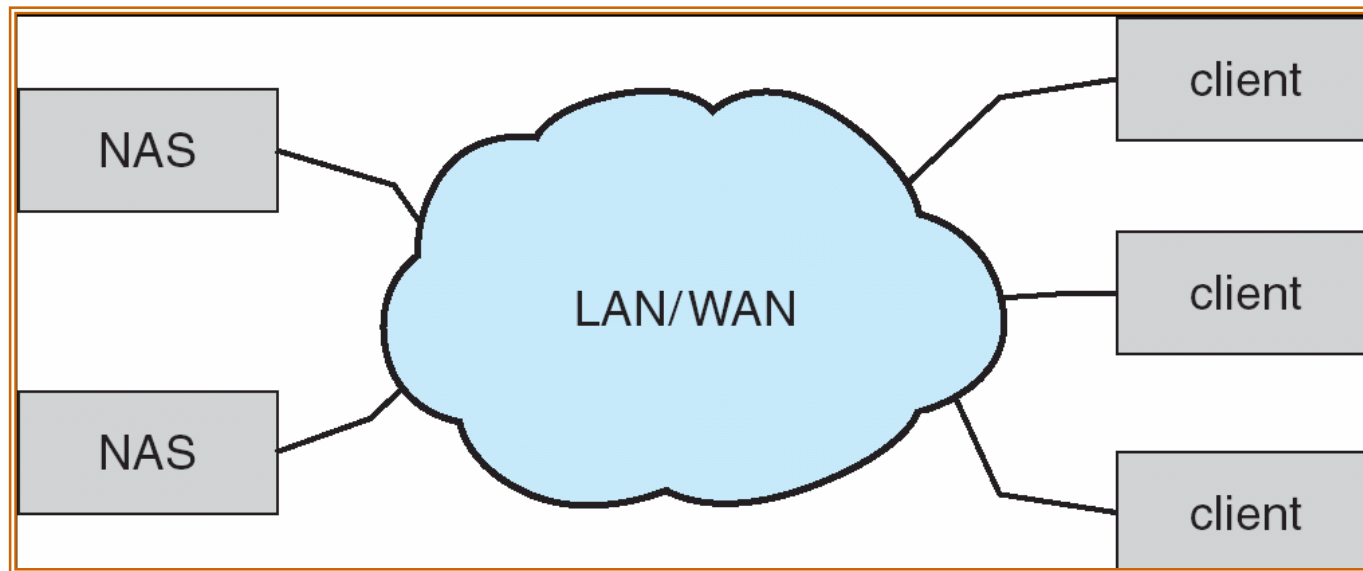
Disk Structure



- 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.

Network-Attached Storage

- Network-attached storage (**NAS**) is storage made available over a network rather than over a local connection (such as a bus)
- NFS and CIFS are common protocols
- Implemented via remote procedure calls (RPCs) between host and storage
- New iSCSI protocol uses IP network to carry the SCSI protocol



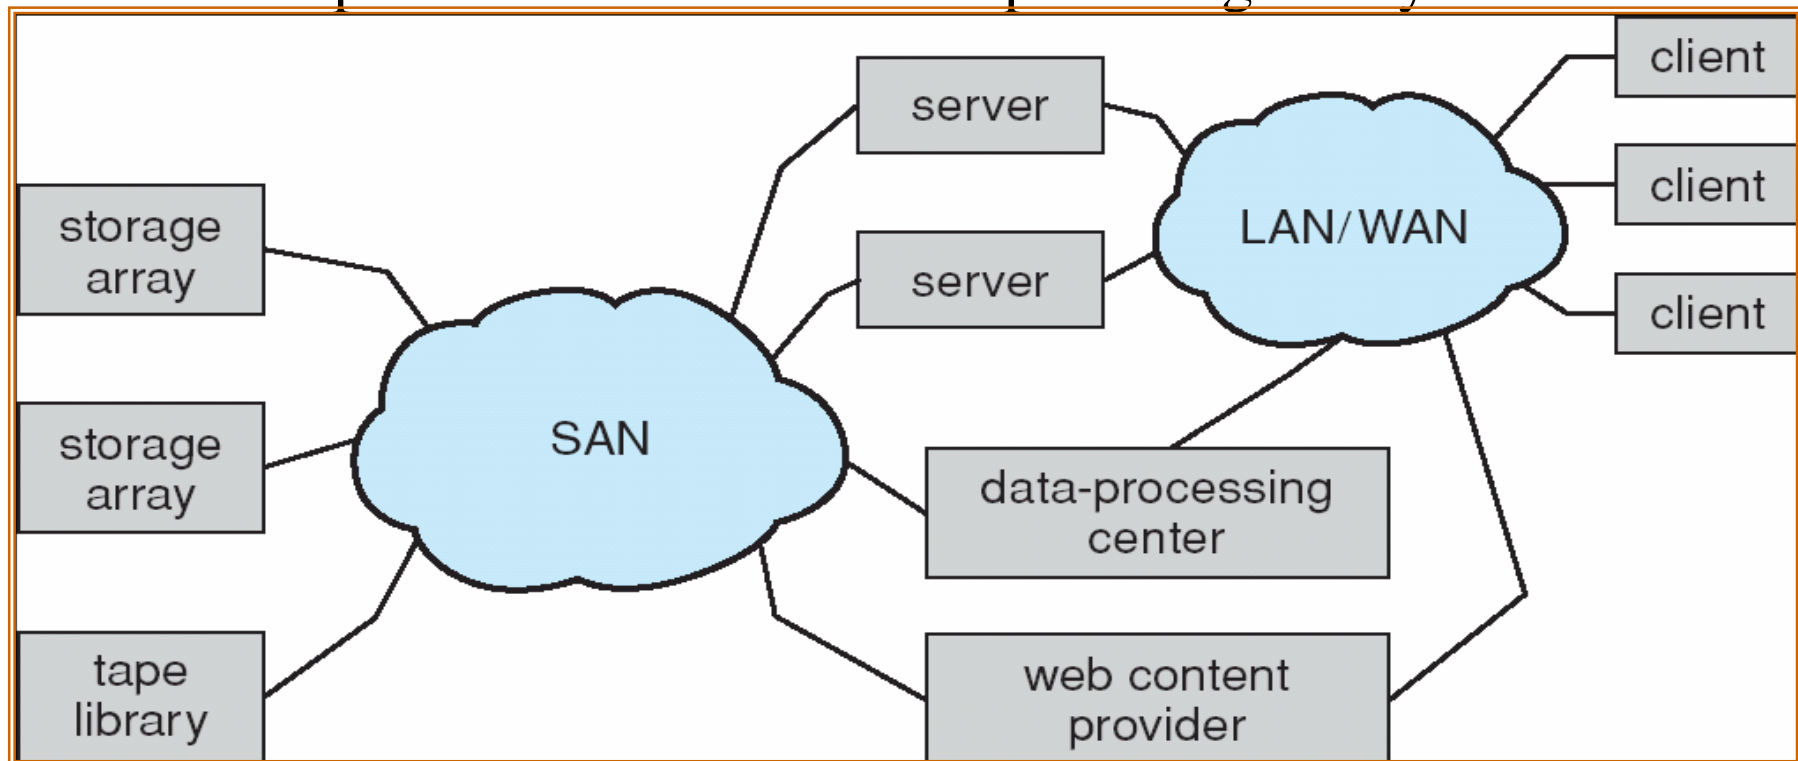
Performance problems?

COMP 7970, Auburn University

Slide 04b-9

Storage Area Network

- Common in large storage environments (and becoming more common)
- Multiple hosts attached to multiple storage arrays - flexible



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 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.
- Minimize seek time
- Seek time \approx 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 (Cont.)

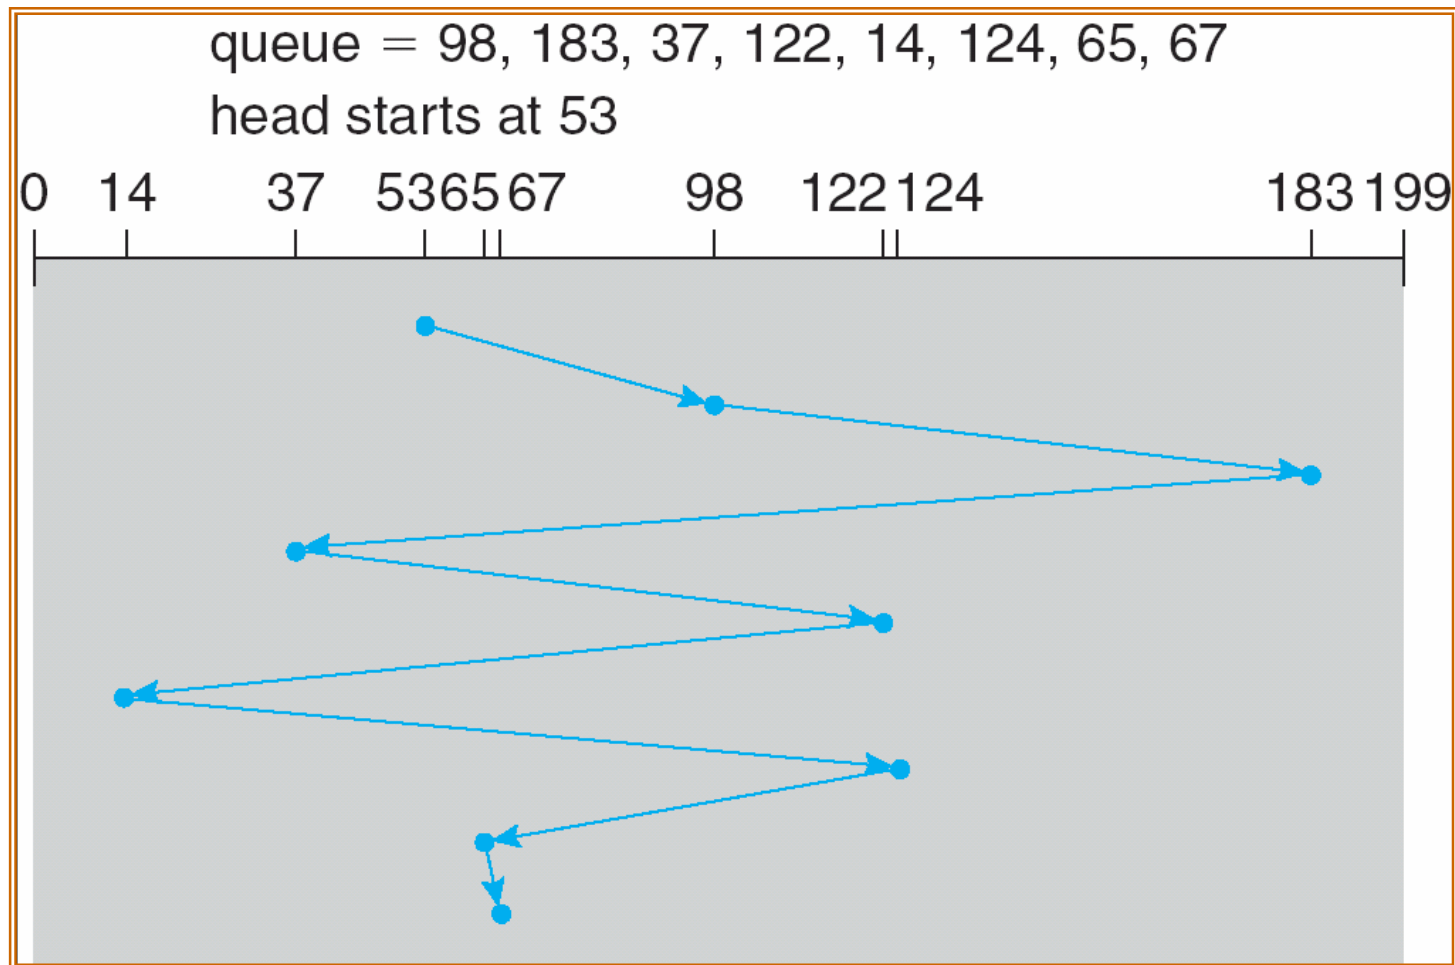
- Several algorithms exist to schedule the servicing of disk I/O requests.
- We illustrate them with a request queue (0-199).

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

Head pointer 53

FCFS

Illustration shows total head movement of 640 cylinders.

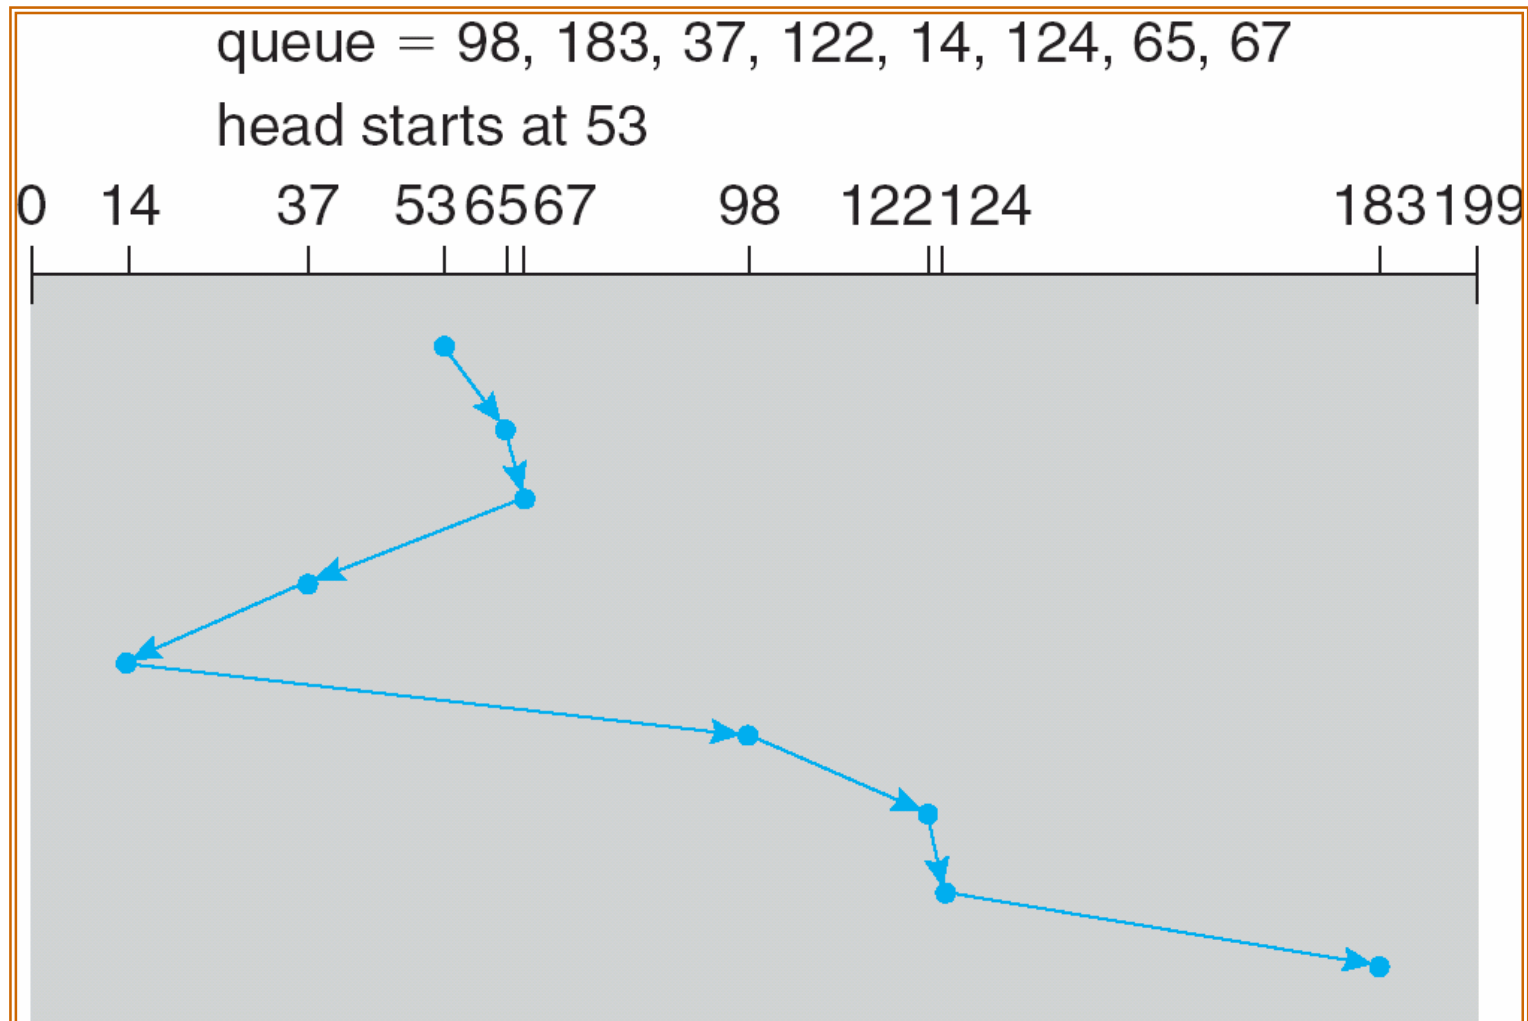


Can you develop a better disk scheduler?

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.)

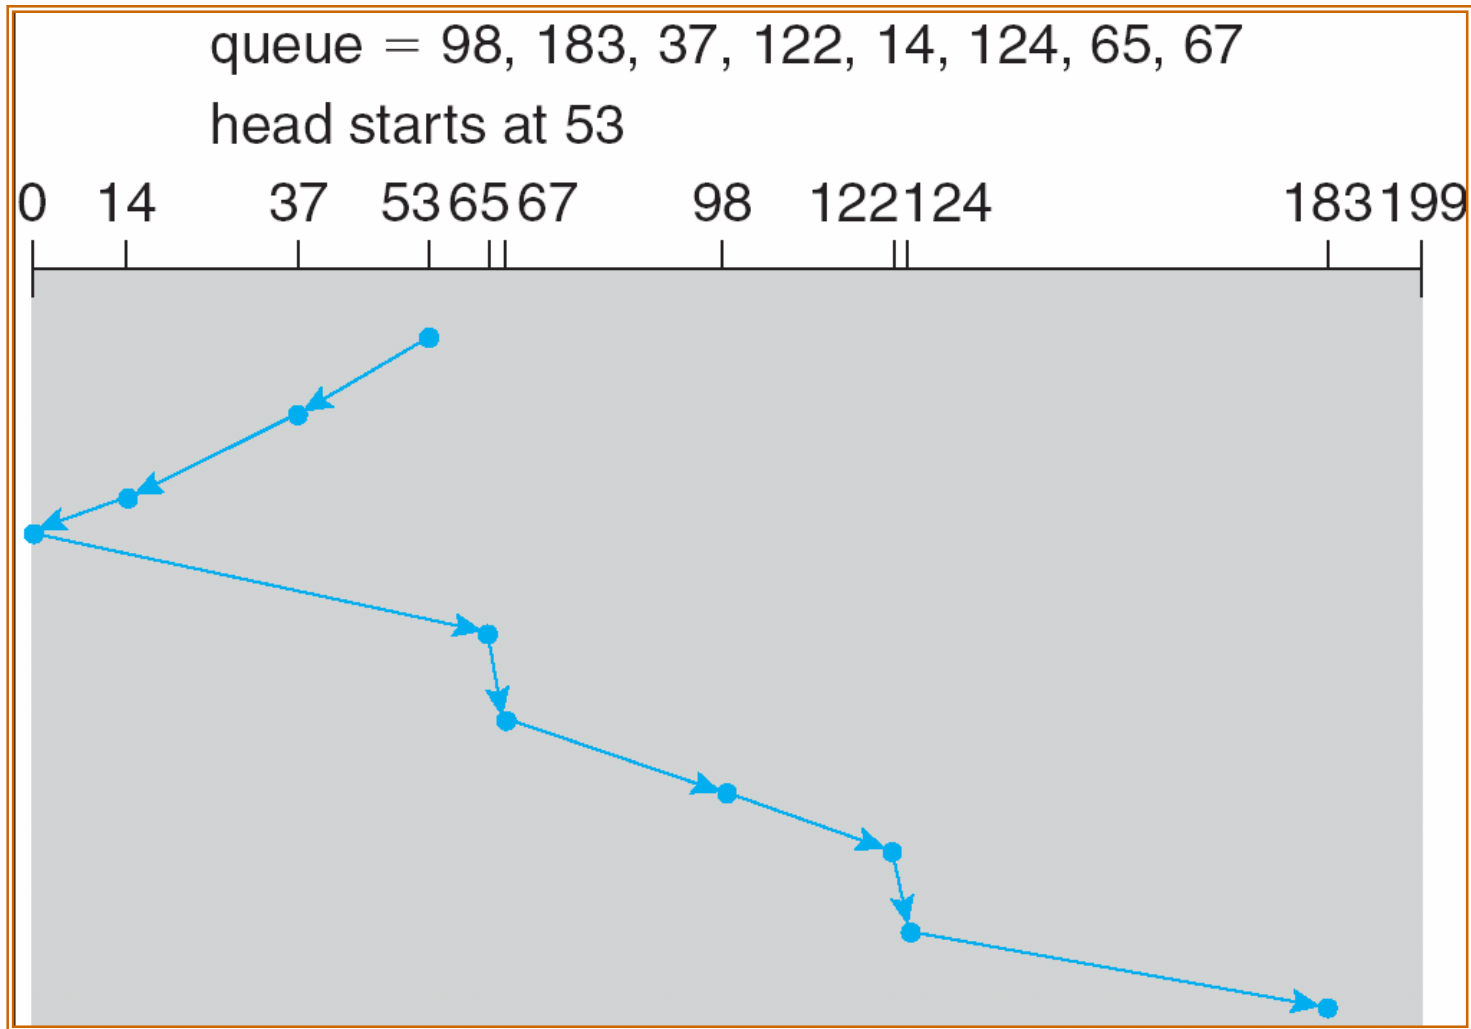


Can you further improve the performance?

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*.
- Illustration shows total head movement of **236** cylinders.

SCAN (Cont.)

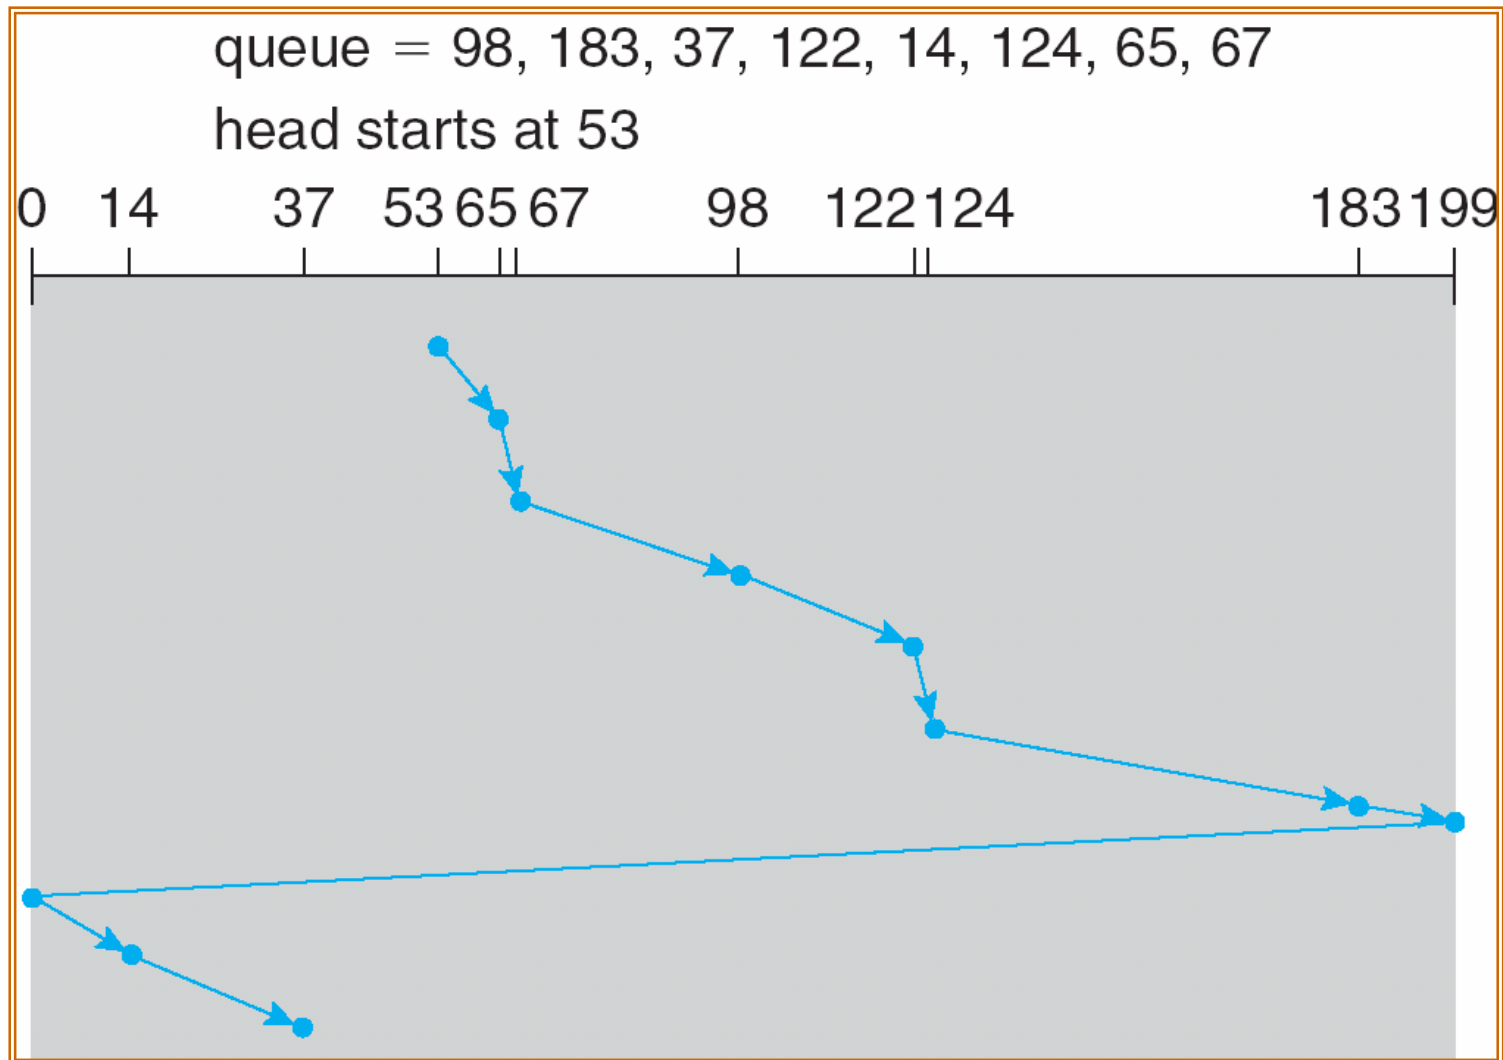


Head movement of **236** cylinders.

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.

C-SCAN (Cont.)



Is there a problem here?

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.)

