

CS 4410
Operating Systems
Storage Devices

Summer 2016
Cornell University

Today

- Devices that can persistently store data.
- Magnetic disk
- Flash storage


Magnetic Disks: Then



Magnetic Disks: Now

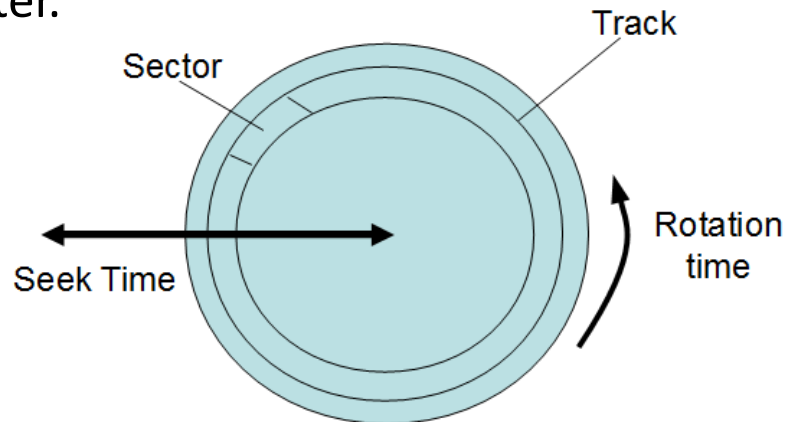


Magnetic Disk: Internal



Disk Speed

- To read from disk, we must specify:
 - cylinder number, track number, sector number.
- disk access time = seek time + rotation time + transfer time
 - seek time: the time to move the disk arm to the desired track
 - rotation time: the time for the desired sector to rotate to the disk head.
 - transfer time: the time for the desired data to flow between the drive and the computer.




Disk Scheduling

- Disk access involves moving mechanical parts, such as arm and platter.
- So, disk access is expensive (milliseconds),
 - comparing to memory access (nanoseconds).
- The disk should be used efficiently.
 - as few movements as possible when serving a sequence of requests.
- Goal : small access time and large bandwidth.
- The OS may employ a disk scheduling algorithm to achieve the above target.


FCFS

- Consider a disk queue with requests for I/O to blocks on cylinders:
 - 98, 183, 37, 122, 14, 124, 65, 67
- The disk head is initially at cylinder 53.
- Total head movement of 640 cylinders




SSTF

- Selects request with minimum seek time from current head position.
- Total head movement of 236 cylinders.




SCAN

- The disk arm starts at one end of the disk.
- Moves toward the other end, servicing requests.
- Head movement is reversed when it gets to the other end of disk.




C-SCAN

- The head moves from one end of the disk to the other.
- Servicing requests as it goes.
- When it reaches the other end it immediately returns to the beginning of the disk.
- Provides a more uniform wait time than SCAN.




C-LOOK

- Arm only goes as far as last request in each direction.
- Then reverses direction immediately.



Flash storage (memory)



- Solid state storage.
- No moving parts.
- Better random IO performance than disks.
- Less power consumption.
- Less vulnerable to physical damage.
- Used in:
 - thumb drives (USBs),
 - smart phones,
 - machine room servers,
 - laptop computers ...
- But, more expensive (per byte) than disks.

Flash storage

- NOR flash storage allows individual words (e.g., 4 bytes) to be written and read.
 - Used for storing device firmware, since it can be executed in place.
- NAND flash storage allows reads and writes of a page (2-4 KB).

Accessing a NAND flash storage

- Erase blocks
 - Before writing, flash memory should be erased in blocks.
 - Slow process
- Write page
- Read page

Today

- Devices that can persistently store data.
- Disk
 - Scheduling requests
- Flash storage

Coming up...

- Next lecture: RAID
- HW4:
 - ex 1, 2, 3, 4
 - Due on Tuesday
- Next exam on Thursday