


# Announcements

- Project #5 is due on April 27<sup>th</sup> at 5:00 pm

# Magnetic Disks



Total capacity: up to 6TB

Collection of platters (1-20)

Rotate at 3600-15000 RPM

Size - usually 2.5-3.5 inch

1,000-50,000 tracks per platter

Track consists of ~100-700 sectors

zones: vary number of tracks/sector based on distance from center

# Access Times

- **Seek: Move disk arm over appropriate track**
  - Seek times vary depending on locality
  - Times are order of milliseconds
- **Rotational delay: Wait until desired information is under disk arm**
  - A disk that rotates at 10,000 RPM will take 6.0 ms to complete a full rotation
  - Improving only a few percent per year
- **Transfer time: time taken to transfer a block of bits**
  - Minimum transfer is one sector
  - Depends on recording density of track, rotation speed, block size
  - Achieved transfer rate for many blocks can also be influenced by other system bottlenecks (software, hardware)
  - Rates range from 2 to 40 MB per second

# Solid State Disks (SSD)

- Random Access nearly as fast as sequential
- Limited number of writes to a sector possible
  - Controller needs to move things around
- Implemented to provide same HW interface as disks
  - IDE and SCSI attached
- Long Term reliability of media still unknown
  - Will they be readable if idle for 5-10 years?

# Disk Scheduling Exercise

- After talking about the goals of disk head scheduling
  - The class split into 7 groups
  - Each group created a scheduling algorithm
  - Groups presented their solutions
- Algorithms presented
  - Shortest request first
  - Scan scheduling
  - Circular Scan
  - Sorted Batch Scheduling
    - Batch up to  $n$  requests
    - Service batch with minimum seek time ordering

# Disk Scheduling

- **First come, first served**
  - ordering may lead to lots of disk head movement
  - i.e. 1, 190, 3, 170, 4, 160 etc.
  - total number of tracks traversed : 863
- **Shortest seek time first: select request with the minimum seek time from current head position**
  - move head to closest track
  - i.e. 1,3,4,160,170,190
  - total number of tracks traversed: 189
  - potential problem with distant tracks not getting service for an indefinite period

# Disk Scheduling

- Scan scheduling - read-write head starts at one end of the disk, moves to the other, servicing requests as it reaches each track
  - Consider example: 1, 190, 3, 170, 4, 160
  - If head starts at track 64 and moves towards 0, the ordering would be 4,3,1,160,170,190
  - Total distance 265
- C-Scan (circular scan)
  - disk head sweeps in only one direction
  - when the disk head reaches one end, it returns to the other
  - Consider example: 1, 190, 3, 170, 4, 160
  - If head starts at track 64 and moves towards 0, the ordering would be 4,3,1,190,170,160
  - Total distance 282

# Disk Cache

- Buffer in main memory for disk sectors
- Cache contains copy of some of the sectors on a disk. When I/O request is made for a sector, a check is made to find out if sector is in the disk cache
- Replacement strategy:
  - Least recently used: block that has been in the cache longest with no reference gets replaced
  - Least frequently used: block that experiences fewest references gets replaced