Announcements

- Reading Chapters 11 (11.6) and 12 (12.1-12.7)
  - suggested problems: 12.6 and 12.8
- Midterm #2
  - it's next week (4/11/96)
  - covers material through and including today's lecture

Magnetic Disks

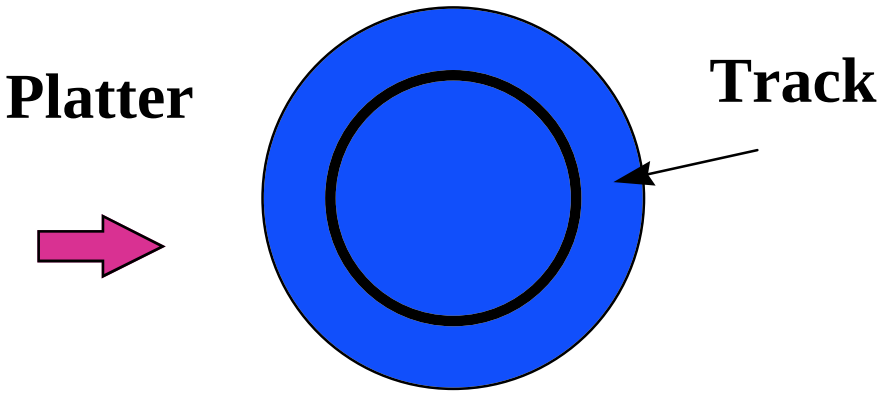

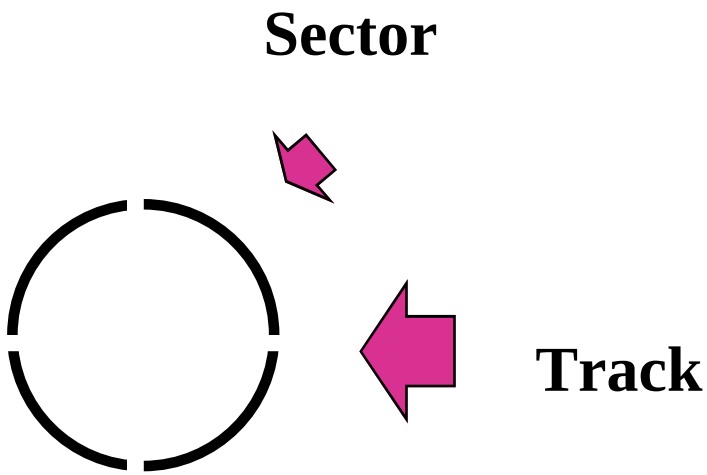

Collection of platters (1-20)
Rotate at 3600-7200 RPM
Size - usually 2.5-3.5 inch
Usually 500-2500 tracks per platter
Track consists of around 64 sectors

Access Times

- Seek: Move disk arm over appropriate track
  - Seek times vary depending on locality - seek times are order of milliseconds
- Rotational delay: Wait time until desired information is under disk arm
  - A disk that rotates at 7200 RPM will take 8.3 ms to complete a full rotation
- Transfer time: time taken to transfer a block of bits (usually a 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 8 MB per second

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. 1,3,4,160,190
  - 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

Bad Blocks

- Some blocks on a disk may not work
  - could be bad from the start (when disk is installed)
  - could go bad during use
- Two options to manage bad blocks
  - disk drive maps the blocks to “replacement” blocks
    - special blocks that are held in reserve for this purpose
  - OS keeps track of where the bad blocks are located and avoids them
- Replacement blocks
  - can be located in tracks at one location, or around the disk
  - provide correct behavior, but change disk performance
- Even if the disk re-maps bad blocks
  - OS could loose data stored on disk
  - needs to be able to recover filesystem from partial update

- Multiple lists: 4.3 BSD Unix has an LRU list and an AGE list
  - LRU list contains blocks that have been used
  - AGE list contains unproven blocks such as blocks that were prefetched because they were contiguous with blocks that had been requested (read-ahead)
  - When a new block (called a buffer in UNIX) is needed, the AGE list is searched first, only if the AGE list is empty does the LRU list get searched

Booting the OS

- How does the OS get loaded and started?
- Process is called booting
  - want to use the OS to load itself
  - but what loads the OS?
- ROM monitor
  - knows how to read from a fixed location on disk and jump into it
- Bootstrap program
  - knows how to load a program from the filesystem and jump into it
- Alternative:
  - put more info into ROM about booting
    - MAC OS has most of the info in ROM
    - hard to change OS without changing ROMs

Swap Space

- Where is swap space located?
  - Is it a “normal” file in the filesystem?
  - Is is in a special location on disk?
- “normal” file
  - ✔ simple, just looks like a file
  - ✔ easy to change size
    - use normal tools
  - slow since it requires all of the filesystem overhead
- separate disk partition
  - ✔ faster
  - harder to change size (need a new partition)

Backups

- Disks can fail, so need to provide a way to copy them
- Two types of backups
  - full backup (all of the data on disks)
  - incremental (data that has changed since last backup)
    - can mark changed files with a field
    - can use the data of the file compared to the last backup
      - permits several levels of backup
- Does the system need to be shutdown for backups?
  - what if a file is moved during a backup?
    - it could get copied 0, 1, or 2 times.
  - easiest answer is to shutdown the machine from dumps