

CSE 120

Principles of Operating Systems

Fall 2000

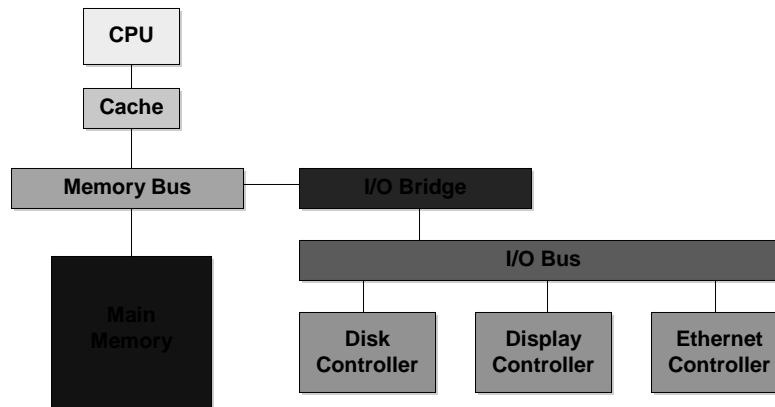
Lecture 12: Disk I/O

Geoffrey M. Voelker

Secondary Storage

- The next topic that we will cover is secondary storage
 - Disk I/O (today)
 - File systems (next time)

I/O Diagram



November 13, 2000

CSE 120 -- Lecture 12 -- Disk I/O

3

Secondary Storage

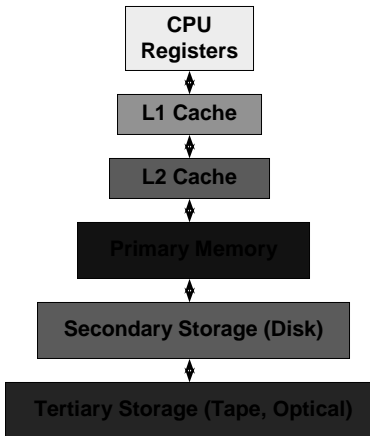
- Secondary storage is usually:
 - ◊ Anything outside of "primary memory"
 - ◊ Anything that does not permit direct instruction execution or data fetch via machine load/store instructions
- Characteristics
 - ◊ It's large – hundreds of megabytes, gigabytes, terabytes
 - ◊ It's cheap – 70 GB disks cost \$1000
 - ◊ It's persistent – data survives loss of power
 - ◊ It's slow – milliseconds to access (why is a millisecond slow?)

November 13, 2000

CSE 120 -- Lecture 12 -- Disk I/O

4

Memory Hierarchy



- Each level acts as a cache of data for the level below

November 13, 2000

CSE 120 -- Lecture 12 -- Disk I/O

5

Secondary Storage Devices

- Drums
 - ◊ Ancient history
- Magnetic disks
 - ◊ Fixed
 - ◊ Removable (floppy)
- Optical disks
 - ◊ Write-once, read-many (CDR)
 - ◊ Write-many, ready-many

- We're going to focus on the use of fixed (hard) magnetic disks for implementing secondary storage

November 13, 2000

CSE 120 -- Lecture 12 -- Disk I/O

6

Disks and the OS

- Disks are messy physical devices:
 - Errors, bad blocks, missed seeks, etc.
- The job of the OS is to hide this mess from higher level software
 - Low-level device control (initiate a disk read, etc.)
 - Higher-level abstractions (files, databases, etc.)
- The OS may provide different levels of disk access to different clients
 - Physical disk block (surface, cylinder, sector)
 - Disk logical block (disk block #)
 - File logical (file block, record, or byte #)


November 13, 2000

CSE 120 -- Lecture 12 -- Disk I/O

7

Physical Disk Structure

- Disk components
 - Platters
 - Surfaces
 - Tracks
 - Sectors
 - Cylinders
 - Arm
 - Heads



November 13, 2000

CSE 120 -- Lecture 12 -- Disk I/O

8

Disk Interaction

- Specifying disk requests requires a lot of info:
 - Cylinder #, surface #, track #, sector #, transfer size...
- Older disks required the OS to specify all of this
 - The OS needed to know all disk parameters
- Modern disks are more complicated
 - Not all sectors are the same size, sectors are remapped, etc.
- The disk provides a higher-level interface (SCSI)
 - The disk exports its data as a logical array of blocks [0...N]
 - » It maps logical blocks to cylinder/surface/track/sector
 - Only need to specify the logical block # to read/write
 - But now the disk parameters are hidden from the OS

November 13, 2000

CSE 120 -- Lecture 12 -- Disk I/O

9

Disk Parameters (1999)

- IBM Ultrastar 36XP drive
 - Form factor: 3.5"
 - Capacity: 36.4 GB
 - Rotation rate: 7,200 RPM
 - Platters: 10
 - Surfaces: 20
 - Sector size: 512-732 bytes
 - Cylinders: 11,494
 - Cache: 4MB
 - Transfer rate: 17.9 MB/s (inner) – 28.9 MB/s (outer)
 - Full seek: 14.5 ms
 - Head switch: 0.3 ms

November 13, 2000

CSE 120 -- Lecture 12 -- Disk I/O

10

Disk Performance

- Disk request performance depends upon a number of steps
 - ♦ Seek – moving the disk arm to the correct cylinder
 - » Depends on how fast disk arm can move (changing very slowly)
 - ♦ Rotation – waiting for the sector to rotate under the head
 - » Depends on rotation rate of disk (increasing, but slowly)
 - ♦ Transfer – transferring data from surface into disk controller electronics, sending it back to the host
 - » Depends on density (increasing fast)
- When the OS uses the disk, it tries to minimize the cost of all of these steps
 - ♦ Particularly seeks and rotation

November 13, 2000

CSE 120 -- Lecture 12 -- Disk I/O

11

Disk Scheduling

- Because seeks are so expensive (milliseconds!), the OS tries to schedule disk requests that are queued waiting for the disk
 - ♦ FCFS (do nothing)
 - » Reasonable when load is low
 - » Long waiting times for long request queues
 - ♦ SSTF (shortest seek time first)
 - » Minimize arm movement (seek time), maximize request rate
 - » Favors middle blocks
 - ♦ SCAN (elevator)
 - » Service requests in one direction until done, then reverse
 - ♦ C-SCAN
 - » Like SCAN, but only go in one direction (typewriter)

November 13, 2000

CSE 120 -- Lecture 12 -- Disk I/O

12

Disk Scheduling (2)

- In general, unless there are request queues, disk scheduling does not have much impact
 - ◆ Important for servers, less so for PCs
- Modern disks often do the disk scheduling themselves
 - ◆ Disks know their layout better than OS, can optimize better
 - ◆ Ignores, undoes any scheduling done by OS

Summary

- I/O overview
- Memory hierarchy
- Secondary storage
 - ◆ Large, persistent, but slow
- Disks
 - ◆ Physical structure
 - ◆ Interface
 - ◆ Performance
 - ◆ Scheduling

Next time...

- Read Chapter 17 (skip 17.5, 17.8)