

# Data Storage

CS157B

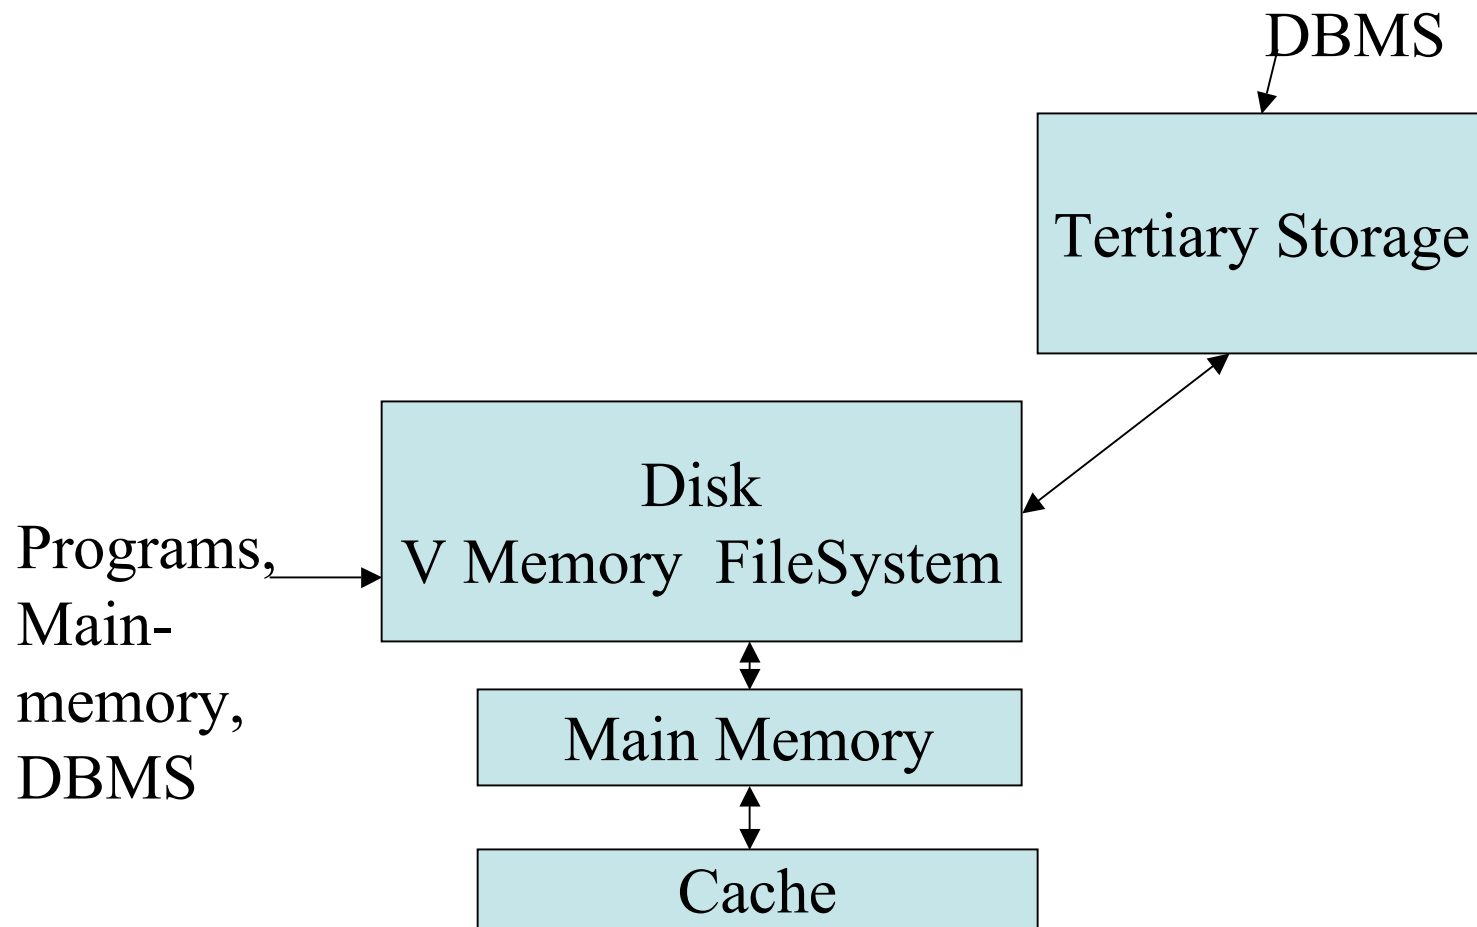
Chris Pollett

Jan 26, 2005.

# Outline

- The Memory Hierarchy
- Disks

# The Memory Hierarchy: How Computer Memory Is Organized



# Cache

- Usually part of the computer's microprocessor
- Is capable of holding both data and instructions
- Only limited memory
- Sometime cache changes not immediately reflected in main memory
- Might have two levels: on-board cache on microprocessor and L2 on another chip
- Cache memory takes on the order of nanosecond to tens of nanoseconds to access.
- Main memory look up order of magnitude slower

# Main Memory

- Nowadays have around 1/2GB to 1 GB of main memory in most new PCs.
- Memory notation is in powers of 2.
  - 1 KB (one kilobyte) is  $2^{10}$  bytes = 1024 bytes
  - 1 MB (one megabyte) is  $2^{20}$  bytes = 1048576
  - 1GB (one gigabyte) is  $2^{30}$  bytes
  - 1 TB (one terabyte) is  $2^{40}$  bytes
  - 1 PB (one petabyte) is  $2^{50}$  bytes
- Moore's Law: ICs are improving at an exponential rate.
- Dubious is can continue. Some things do not follow. Example network communication.

# Virtual Memory

- Programs make use of a virtual addressing scheme for storing data during program execution.
- Goal is to try to fit as much data in memory as can at one time. If need to look up some data that is not in memory then go to disk to read/store.
- For 32 bit processors one word can store up to  $2^{32}$  different numbers.
- So pretend have 4GB of main memory
- Swapping between memory and disk done using page sized blocks (4KB to 56KB).

# Secondary Storage

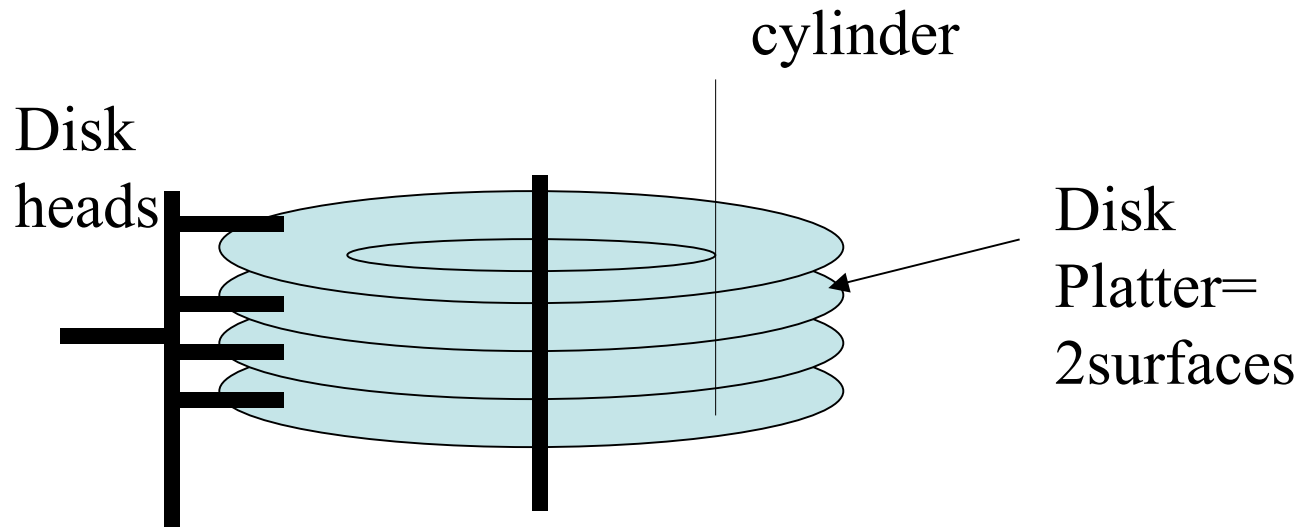
- Usually some kind of slower memory device than main memory which is more permanent and capacious.
- Usually, secondary storage means a disk.
- Disks are about 100,000 times slower than memory
- 300GB drives reasonably common
- Data on disk organized in blocks.
- Disk reads/Disk writes (Disk I/Os) operate on the block level
- Memory read off disk, is usually temporarily stored in a buffer

# Tertiary Storage

- Slower but yet more capacious memory.
- Tape drives, Optical-Disk Juke Boxes, Tape



# How Disks Work



- Given platter is broken into track which are further broken into arcs called sectors

# Disk Controller

- Small processor in the disk drive capable of:
  - Controlling the actuator arm that move the head assembly
  - Selecting a surface to read/write from
  - Transferring bits to/from the desired sector to/from the computer's memory

# Disk Storage Characteristics

- Rotation speed: say 5400 RPMs
- Number of Platters: say 30 surfaces
- Number of tracks per surface: say 10000
- Number of bytes per track: say 100000

# Disk Access Characteristics

- Several components makes up the time from a read/write request until the block appears in main memory (called latency):
  - processor/controller time
  - time to position head to correct cylinder (seek time).  
10-40 mS
  - time to rotate head over sector (rotational latency) 10 mS
  - transfer time - time for whole sector/gaps between adjacent sectors to rotate under head.