COT 4600 Operating Systems Fall 2009

Dan C. Marinescu

Office: HEC 439 B

Office hours: Tu-Th 3:00-4:00 PM

Lecture 5

- Last time:
  - Project.
- Today:
  - Names and the basic abstractions
  - Storage
- Next Time
  - Naming in computing systems

Names and the three basic abstractions

- Memory ➔ stores named objects
  - write(name, value) value ← READ(name)
  - file system: /dcm/classes/Fall09/Lectures/Lecture5.ppt
- Interpreters ➔ manipulates named objects
  - machine instructions ADD R1,R2
  - modules ➔ Variables call sort(table)
- Communication Links ➔connect named objects
  - HTTP protocol used by the Web and file systems

    Host: boticelli.cs.ucf.edu

    put /dcm/classes/Fall09/Lectures/Lecture5.ppt

    get /dcm/classes/Fall09/Lectures/Lecture5.ppt

Latency and Bandwidth

- Important concepts for physical characterization.
- Applies to all three abstractions.
- Informal
  - Bandwidth ➔ number of operations per second!
  - Latency ➔ to get there
- The bandwidth of the CPU, Memory, and I/O sbsystems must be balanced.

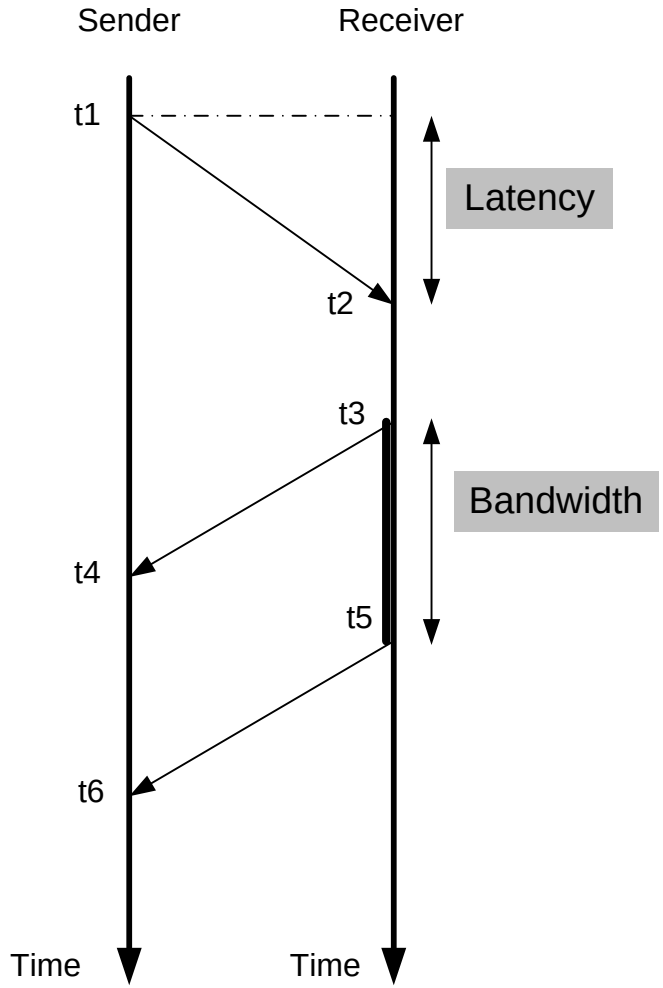

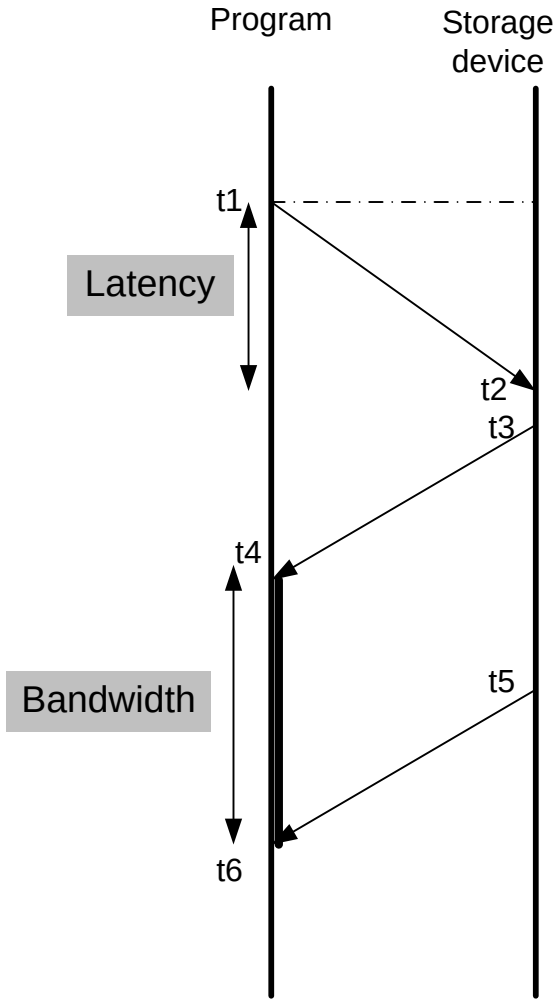

Communication latency- time it takes the first bit sent to reach the receiver

Bandwidth- number of bits/bytes transmitted per unit of time

Operation latency- time it takes the command to read the device

Bandwidth- number of bits/bytes transmitted per unit of time

Memory

- Hardware memory:
  - Devices
    - RAM (Random Access Memory) chip
    - Flash memory ➔ non-volatile memory that can be erased and reprogrammed
    - Magnetic tape
    - Magnetic Disk
    - CD and DVD
  - Systems
    - RAID
    - File systems
    - DBMS (Data Base management Systems)

Attributes of the storage medium/system

- Durability ➔ the time it remembers
- Stability ➔ whether or not the data is changed during the storage
- Persistence ➔ property of data storage system, it keeps trying to preserve the data

Critical properties of a storage medium/system

- Read/Write Coherence ➔ the result of a READ of a memory cell should be the same as the most recent WRITE to that cell.
- Before-or-after atomicity ➔the result of every READ or WRITE is as if that READ or WRITE occurred either completely before or completely after any other READ or WRITE

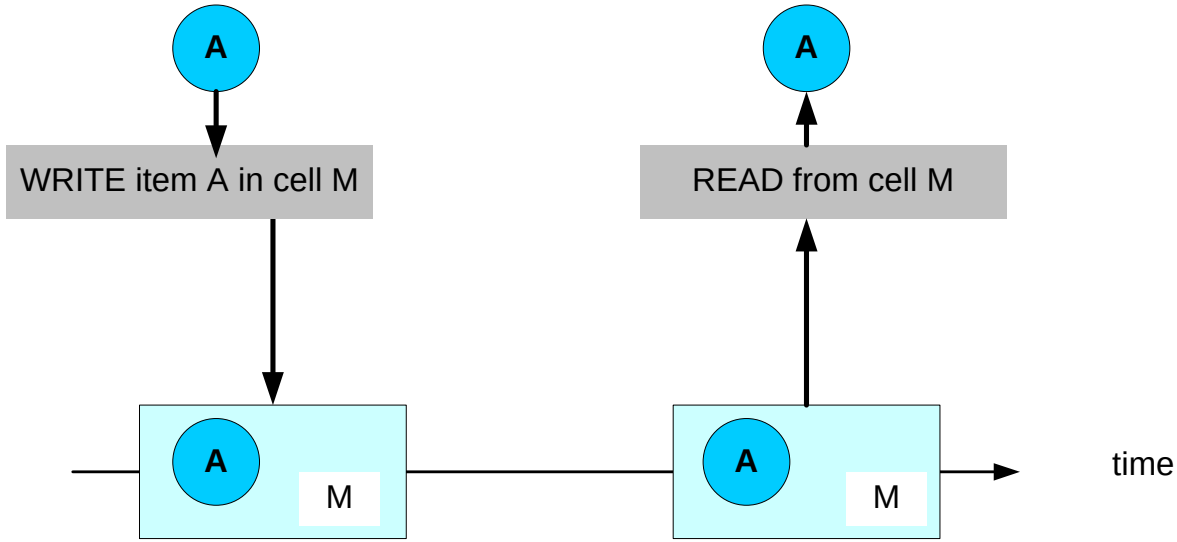

Read/Write Coherence ➔ the result of a READ of a memory cell should be the same as the most recent WRITE to that cell.

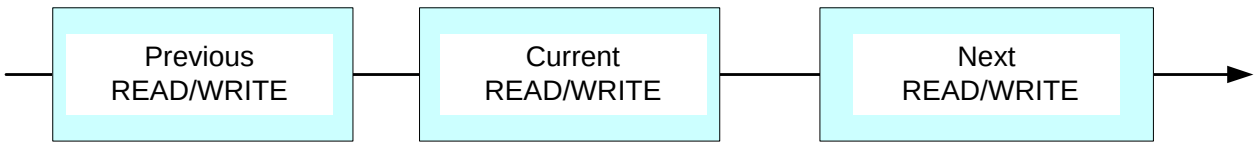

Before-or-after atomicity ➔ the result of every READ or WRITE is as if that READ or WRITE occurred either completely before or completely after any other READ or WRITE

Why it is hard to guarantee the critical properties?

- Concurrency ➔ multiple threads could READ/WRITE to the same cell
- Remote storage ➔ The delay to reach the physical storage may not guarantee FIFO operation
- Optimizations ➔ data may be buffered to increase I/O efficiency
- Cell size may be different than data size ➔ data may be written to multiple cells.
- Replicated storage ➔ difficult to maintain consistency.

Access type; access time

- Sequential access
  - Tapes
  - CD/DVD
- Random access devices
  - Disk
    - Seek
    - Search time
    - Read/Write time
  - RAM

Physical memory organization

- RAM ➔ two dimensional array. To select a flip-flop provide the x and y coordinates.
- Tapes ➔ blocks of a given length and gaps (special combination of bits.
- Disk:
  - Multiple platters
  - Cylinders correspond to a particular position of the moving arm
  - Track ➔ circular pattern of bits on a given platter and cylinder
  - Record ➔ multiple records on a track

Names and physical addresses

- Location addressed memory ➔ the hardware maps the physical coordinates to consecutive integers, addresses
- Associative memory ➔ unrestricted mapping; the hardware does not impose any constraints in mapping the physical coordinates

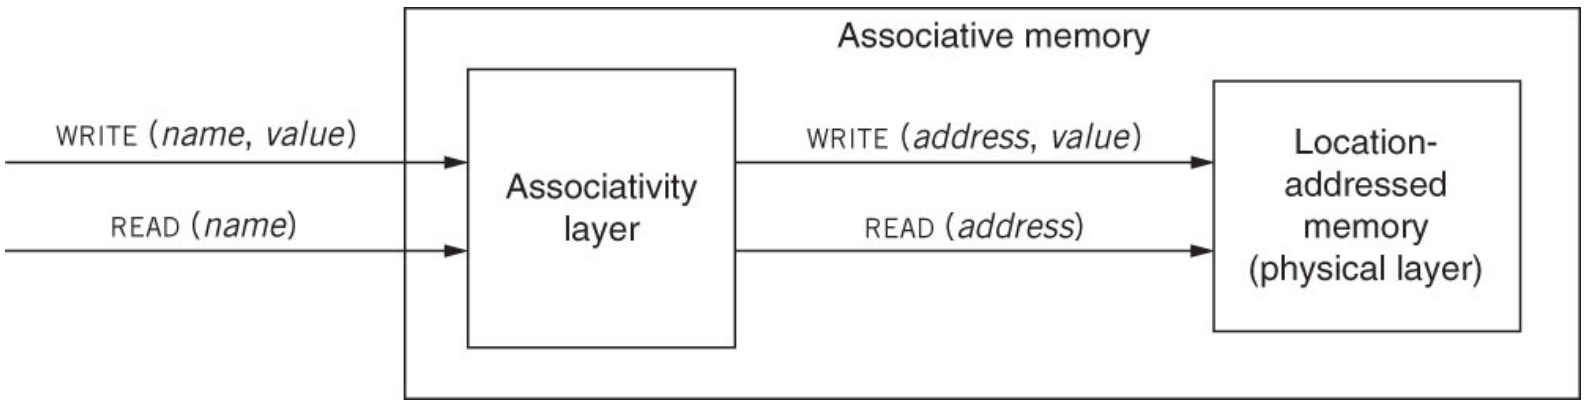

Figure 2.2 from textrbook

RAID – Redundant Array of Inexpensive Disks

- The abstraction put to work to increase performance and durability.
- Raid 0 ➔ allows concurrent reading and writing. Increases performance but does not improve reliability.
- Raid 1 ➔ increases durability by replication the block of data on multiple disks (mirroring)
- Raid 2 ➔ Disks are synchronized and striped in very small stripes, often in single bytes/words.
  - Error correction calculated across corresponding bits on disks, and is stored on multiple parity disks (Hamming codes).

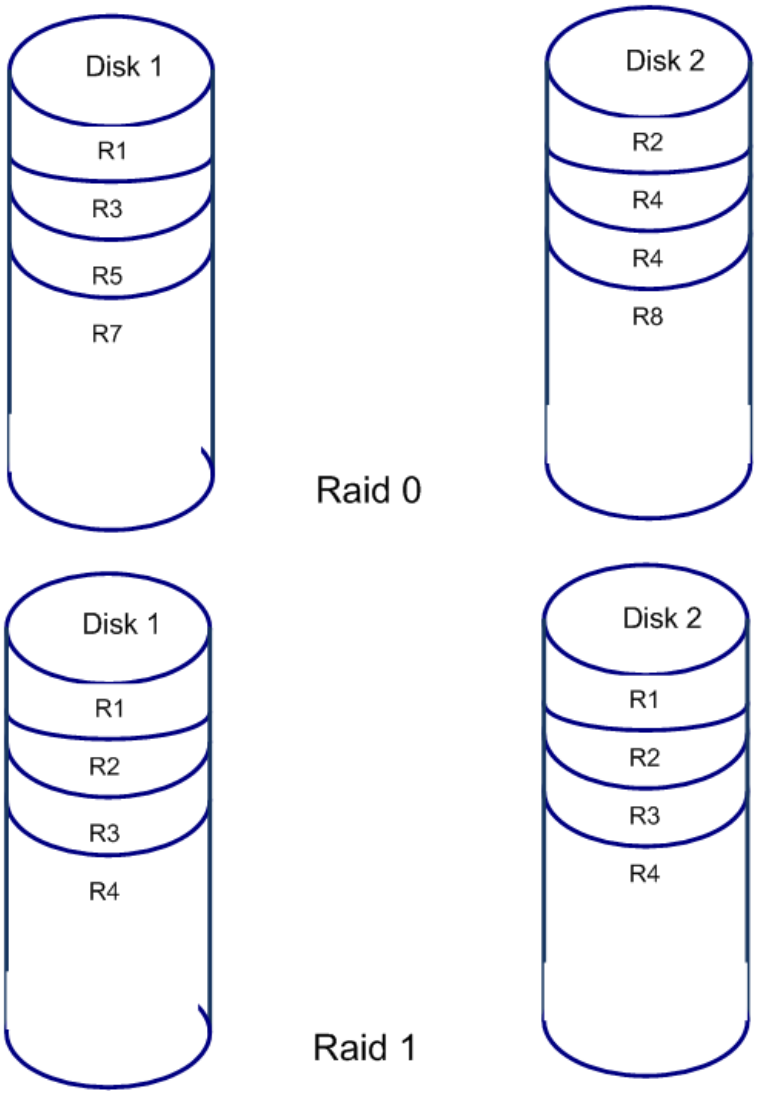
Disk 1
R1
R3
R5
R7
Disk 2
R2
R4
R4
R8
Raid 0
Disk 1
R1
R2
R3
R4
Disk 2
R1
R2
R3
R4
Raid 1

RAID (cont'd)

- Raid 3 ➔**Striped set with dedicated parity** or **bit interleaved parity** or **byte level parity**.
- Raid 4 ➔ improves reliability, it adds error correction

RAID (cont'd)

- Raid 5 ➔ striped disks with parity combines three or more disks to protect data against loss of any one disk.
  - The storage capacity of the array is reduced by one disk
- Raid 6 ➔ striped disks with dual parity combines four or more disks to protect against loss of any two disks.
  - Makes larger RAID groups more practical.
  - Large-capacity drives lengthen the time needed to recover from the failure of a single drive. Single parity RAID levels are vulnerable to data loss until the failed drive is rebuilt: the larger the drive, the longer the rebuild will take. Dual parity gives time to rebuild the array without the data being at risk if a (single) additional drive fails before the rebuild is complete.

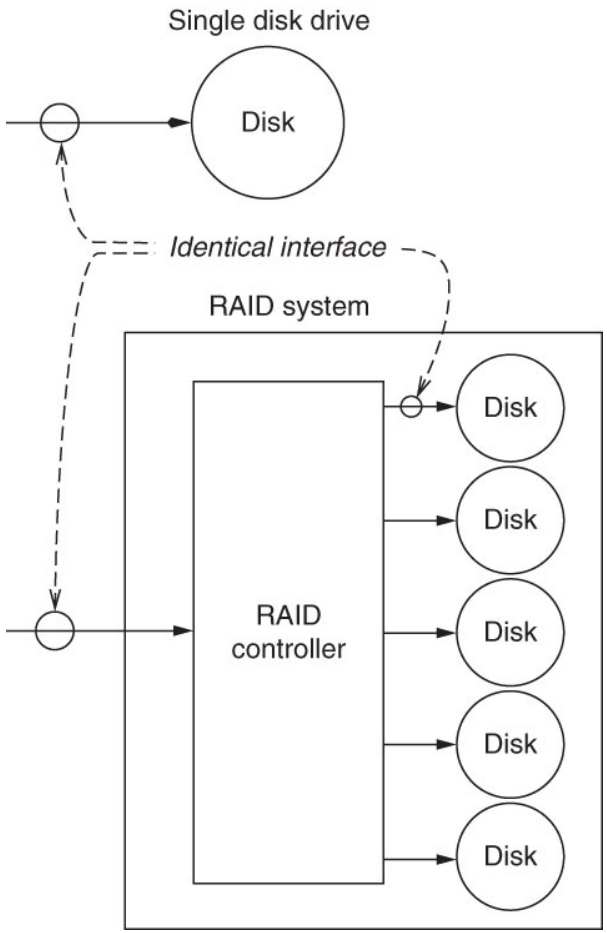

Figure 2.3 from textbook