

Jin-Soo Kim
(jinsoo.kim@snu.ac.kr)

Systems Software &
Architecture Lab.
Seoul National University

Fall 2020

The Memory Hierarchy

Chap. 5.1 – 5.2



Memory Technologies

- **Static RAM (SRAM)**
 - Each cell stores a bit with a four or six-transistor circuit
 - Retains value indefinitely, as long as it is kept powered
 - Faster and more expensive than DRAM
- **Dynamic RAM (DRAM)**
 - Each cell stores a bit with a capacitor. One transistor is used for access
 - Value must be refreshed every 10 – 100 ms
 - Slower and cheaper than SRAM


Technology	Typical access time	\$ Per GiB in 2012
SRAM	0.5 – 2.5ns	\$500 -- \$1000
DRAM	50 – 70ns	\$10 -- \$20
Flash	5 μ s – 50 μ s	\$0.75 -- \$1.00
Disk	5ms – 20ms	\$0.05 -- \$0.10

Non-volatile Memories

- Nonvolatile memories retain value even if powered off
 - Read-only memory (**ROM**): programmed during production
 - Programmable ROM (**PROM**): can be programmed once
 - Erasable PROM (**EPROM**): can be bulk erased (UV, X-ray)
 - Electrically erasable PROM (**EEPROM**): electronic erase capability
 - **Flash** memories: EEPROMs with partial (block-level) erase capability (NOR vs. NAND)
 - Intel **Optane** memory: slower than DRAM, denser and better cost/GiB than DRAM
- Uses for nonvolatile memories
 - Firmware programs stored in a ROM (BIOS, Disk/network/graphics controllers, ...)
 - USB drives, smartphones, tablets, SSDs (Solid-State Drives), disk caches, ...
 - Main memory?

The CPU-Memory Gap

- The gap widens between DRAM, disk, and CPU speeds





Locality to the Rescue!

- Question:
How can we make a memory as fast as SRAM and as cheap as DRAM (or even disk)?
- The key to bridging this CPU-Memory gap is a fundamental property of computer programs known as **locality**

Principle of Locality

- Programs tend to use data and instructions with addresses near or equal to those they have used recently
- Programs access a small portion of their address space at any time
- Temporal locality
 - Recently referenced items are likely to be accessed again soon
e.g., instructions in a loop, induction variables
- Spatial locality
 - Items near those accessed recently are likely to be accessed soon
 - e.g., sequential instruction access, array data



Principle of Locality: Example

```
sum = 0;
for (i = 0; i < n; i++)
    sum += a[i];
return sum;
```

■ Data

- Reference array elements in succession (stride-1 reference pattern)
- Reference `sum` each iteration

Spatial locality

Temporal locality

■ Instructions

- Reference instructions in sequence
- Cycle through loop repeatedly

Spatial locality

Temporal locality

Memory Hierarchy (I)

- Some fundamental and enduring properties of hardware and software
 - Fast storage technologies cost more per byte, have less capacity, and require more power (heat!)
 - The gap between CPU and main memory speed is widening
 - Well-written programs tend to exhibit good locality
- These fundamental properties complement each other beautifully
- They suggest an approach for organizing memory and storage systems known as a **memory hierarchy**

Memory Hierarchy (2)

“We are therefore forced to recognize the possibility of constructing a hierarchy of memories, each of which has greater capacity than the preceding but which is less quickly accessible.”

-- A. W. Burks, H. H. Goldstein, J. von Neumann, Preliminary Discussion of the Logical Design of Electronic Computing Instrument, June 1946.

- Taking advantage of locality
 - Store everything on disk
 - Copy recently accessed (and nearby) items from disk to smaller DRAM memory (main memory)
 - Copy more recently accessed (and nearby) items from DRAM to smaller SRAM memory (cache memory)

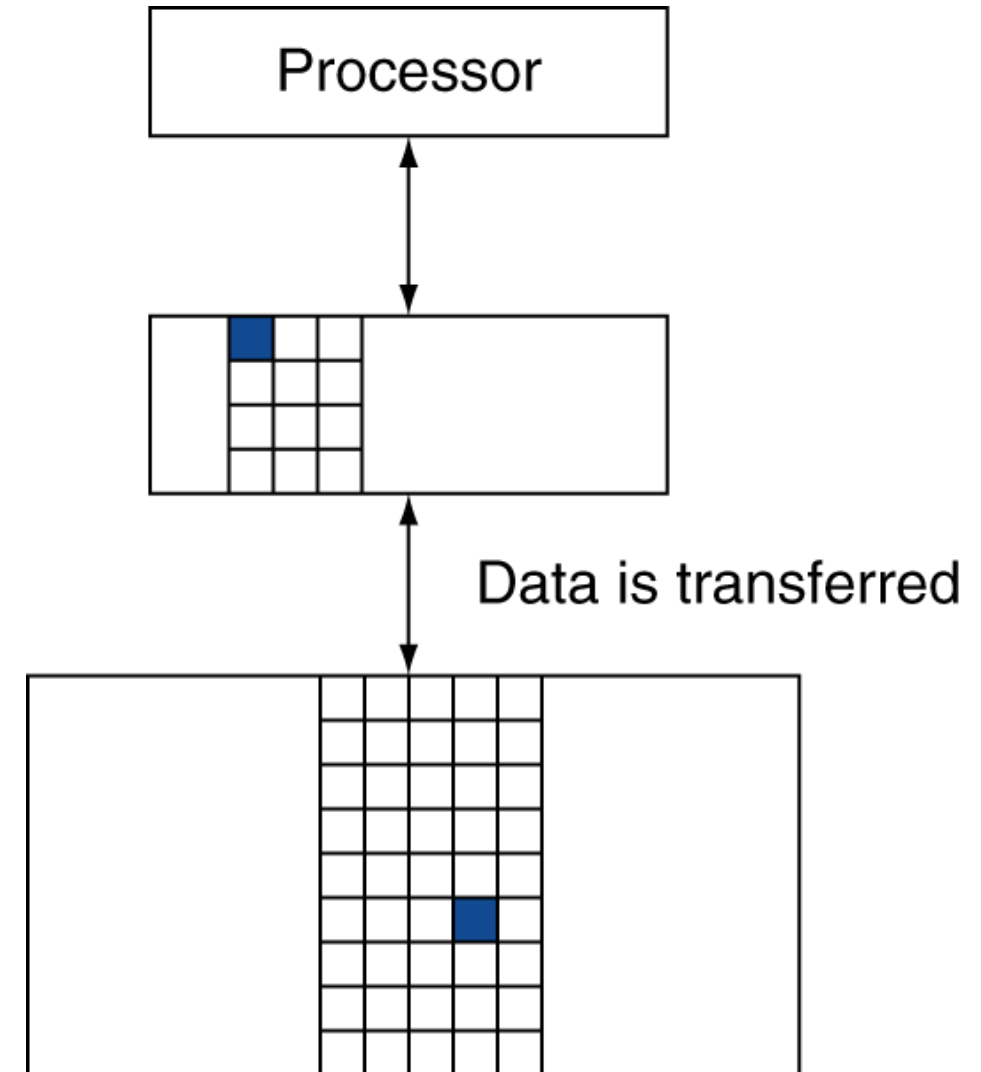
Caches

- A smaller, faster storage device that acts as a staging area for a subset of the data in a larger, slower device
- Fundamental idea of a memory hierarchy
 - For each k , the faster, smaller device at level k serves a cache for the larger, slower device at level $k+1$
- Why do memory hierarchies work?
 - Because of locality, programs tend to access the data at level k more often than they access the data at level $k+1$

Big Idea: The memory hierarchy creates a large pool of storage that costs as much as the cheap storage near the bottom, but that serves data to programs at the rate of the fast storage near the top

Memory Hierarchy Levels

- **Block (or line):** a unit of copying
 - May be multiple words
- **If accessed data is present in upper level**
 - **Hit:** access satisfied by upper level
 - **Hit ratio:** hits / accesses
- **If accessed data is absent**
 - **Miss:** block copied from lower level
 - Then accessed data supplied from upper level
 - Time taken: **miss penalty**
 - **Miss ratio:** misses / accesses = $1 - \text{hit ratio}$



Memory Hierarchy

