

The Memory Hierarchy

Kai Shen

1

Random-Access Memory (RAM)

- Key features
 - RAM is traditionally packaged as a chip.
 - Basic storage unit is normally a **cell** (one bit per cell).
 - Multiple RAM chips form a memory.
- 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.
 - Relatively insensitive to electrical noise (EMI), radiation, etc.
- Dynamic RAM (DRAM)
 - Each cell stores bit with a capacitor. One transistor is used for access
 - Value must be refreshed every 10-100 ms.
 - More sensitive to disturbances (EMI, radiation,...) than SRAM.
 - Slower and cheaper than SRAM.

2

Conventional DRAM Organization

- $d \times w$ DRAM:
 - dw total bits organized as d **supercells** of size w bits

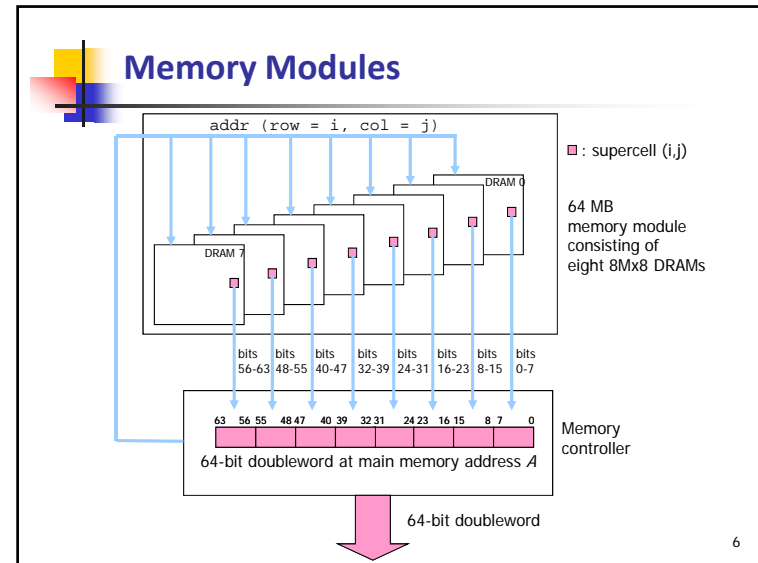
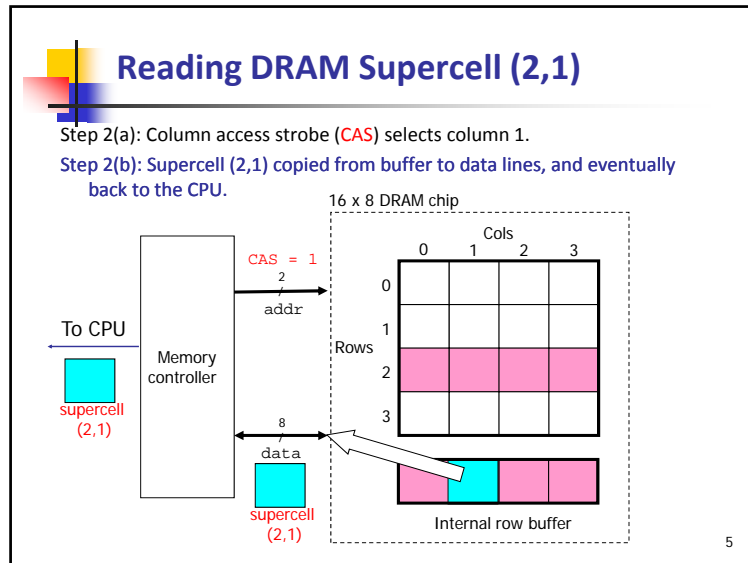
3

Reading DRAM Supercell (2,1)

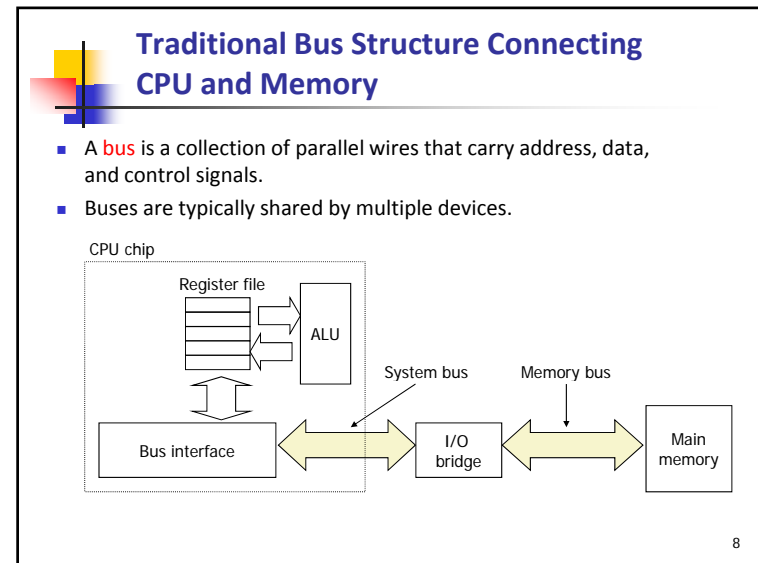
Step 1(a): Row access strobe (**RAS**) selects row 2.

Step 1(b): Row 2 copied from DRAM array to row buffer.

4



- ### Nonvolatile Memories
- DRAM and SRAM are volatile memories
 - Lose information if powered off.
 - 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 memory: EEPROMs with partial (sector) erase capability
 - Wears out after about 100,000 erasings.
 - Uses for nonvolatile memories
 - Firmware programs stored in a ROM (BIOS, controllers for disks, network cards, graphics accelerators, security subsystems,...)
 - Solid state disks (replace rotating disks in thumb drives, smart phones, mp3 players, tablets, laptops,...)
 - Disk caches
- 7



Memory Read Transaction (1)

- CPU places address A on the memory bus.

Register file
%eax
ALU
Bus interface
I/O bridge
Main memory
0
A
x

Load operation: `movl A, %eax`

9

Memory Read Transaction (2)

- Main memory reads A from the memory bus, retrieves word x, and places it on the bus.

Register file
%eax
ALU
Bus interface
I/O bridge
Main memory
0
A
x

Load operation: `movl A, %eax`

10

Memory Read Transaction (3)

- CPU read word x from the bus and copies it into register %eax.

Register file
%eax
ALU
Bus interface
I/O bridge
Main memory
0
A
x

Load operation: `movl A, %eax`

11

Memory Write Transaction (1)

- CPU places address A on bus. Main memory reads it and waits for the corresponding data word to arrive.

Register file
%eax
ALU
Bus interface
I/O bridge
Main memory
0
A

Store operation: `movl %eax, A`

12

Memory Write Transaction (2)

- CPU places data word y on the bus.

Register file
%eax
ALU
Bus interface
I/O bridge
Main memory
0
A
y

Store operation: `movl %eax, A`

13

Memory Write Transaction (3)

- Main memory reads data word y from the bus and stores it at address A .

register file
%eax
ALU
bus interface
I/O bridge
main memory
0
A
y

Store operation: `movl %eax, A`

14

What's Inside A Disk Drive?

Arm
Spindle
Platters
Actuator
Electronics (including a processor and memory!)
SCSI connector

Image courtesy of Seagate Technology

15

Disk Geometry

- Disks consist of **platters**, each with two **surfaces**.
- Each surface consists of concentric rings called **tracks**.
- Each track consists of **sectors**.

Tracks
Surface
Spindle
Track k
Sectors

16

Disk Operation

The disk surface spins at a fixed rotational rate.

The read/write head is attached to the end of the arm and flies over the disk surface on a thin cushion of air.

By moving radially, the arm can position the read/write head over any track.

spindle

17

Disk Geometry (Multiple-Platter View)

- Aligned tracks form a cylinder.

Surface 0
Surface 1
Surface 2
Surface 3
Surface 4
Surface 5

Platter 0
Platter 1
Platter 2

Cylinder k

Spindle

18

Disk Operation (Multi-Platter View)

Read/write heads move in unison from cylinder to cylinder

Arm

Spindle

19

Disk Seek and Rotation

Access to BLUE sector and then access to RED sector.

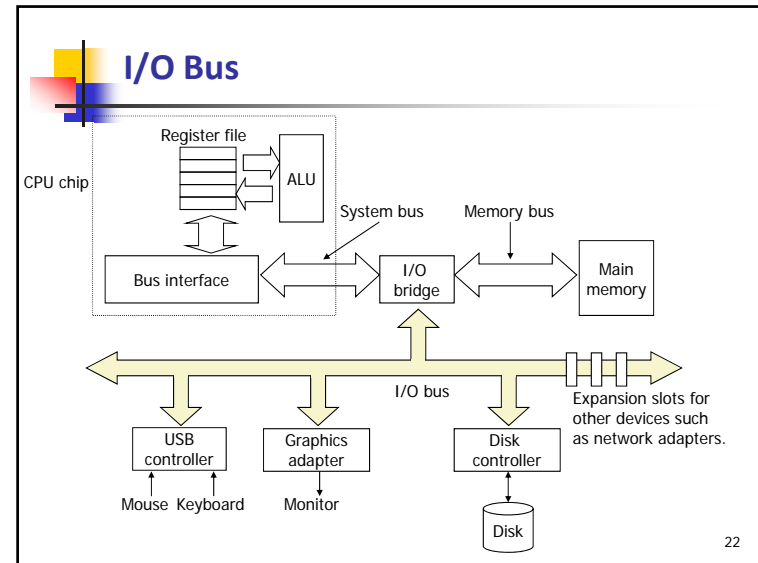
After BLUE read Seek for RED Rotational latency After RED read

20

Disk Access Time

- Time to access some target sector approximated by:
 - Seek_time + Rotation_time + Transfer_time
- Seek time**
 - Time to position heads over cylinder containing target sector.
 - Typically 3–9 ms
- Rotational latency**
 - Time waiting for first bit of target sector to pass under r/w head.
 - Average rotational latency?
 - For a 10KRPM disk?
- Transfer time**
 - Time to read the bits in the target sector.
 - 80MB/s transfer speed. Time for transferring 4KB?
- Access time dominated by seek time and rotational latency!

21



Flash-based Solid State Disks (SSDs)

The diagram shows the internal structure of a Solid State Disk (SSD). It is connected to an I/O bus. The I/O bus sends "Requests to read and write logical disk blocks" to a Flash translation layer. This layer manages the mapping between logical blocks and physical flash memory. The Flash memory is organized into blocks, such as Block 0 and Block B-1. Each block contains multiple pages (Page 0, Page 1, ..., Page P-1).

- Pages: 2KB to 8KB, Blocks: 64 to 256 pages
- Data read/written in units of pages.
- Page can be rewritten only after its block has been erased
- A block wears out after 100,000 repeated writes.

23

SSD Performance Characteristics


- Performance characteristics**
 - Sequential reads/writes comparable (a bit faster) than disks
 - Random reads are much faster than disks (0.1ms vs. 10ms)
 - Random writes are somewhat slower than reads (still much faster than disks)
- Why are random writes slow?**
 - Erasing a block is slow (around 1 ms) and it forces a copy of all useful pages in the block

24

SSD Tradeoffs vs Rotating Disks

- Advantages
 - No moving parts → faster, less power, more rugged
- Disadvantages
 - Have the potential to wear out
 - Mitigated by “wear leveling logic” in flash translation layer
 - In 2010, about 100 times more expensive per byte
- Applications
 - MP3 players, smart phones, laptops
 - Beginning to appear in desktops and servers

25



Locality to the Rescue!

The key to bridging this CPU-Memory gap is a fundamental property of computer programs known as **locality**

27

Locality

- Principle of Locality:** Programs tend to use data and instructions with addresses near or equal to those they have used recently
- Temporal locality:**
 - Recently referenced items are likely to be referenced again in the near future
- Spatial locality:**
 - Items with nearby addresses tend to be referenced close together in time

28

Locality Example

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

- Data references
 - Reference array elements in succession (stride-1 reference pattern). Spatial locality
 - Reference variable `sum` each iteration. Temporal locality
- Instruction references
 - Reference instructions in sequence. Spatial locality
 - Cycle through loop repeatedly. Temporal locality

29

Qualitative Estimates of Locality

- **Claim:** Being able to look at code and get a qualitative sense of its locality is a key skill for a professional programmer.
- **Question:** Does this function have good locality with respect to array `a`?

```
int sum_array_rows(int a[M][N])
{
    int i, j, sum = 0;

    for (i = 0; i < M; i++)
        for (j = 0; j < N; j++)
            sum += a[i][j];
    return sum;
}
```

30

Locality Example

- **Question:** Does this function have good locality with respect to array `a`?

```
int sum_array_cols(int a[M][N])
{
    int i, j, sum = 0;

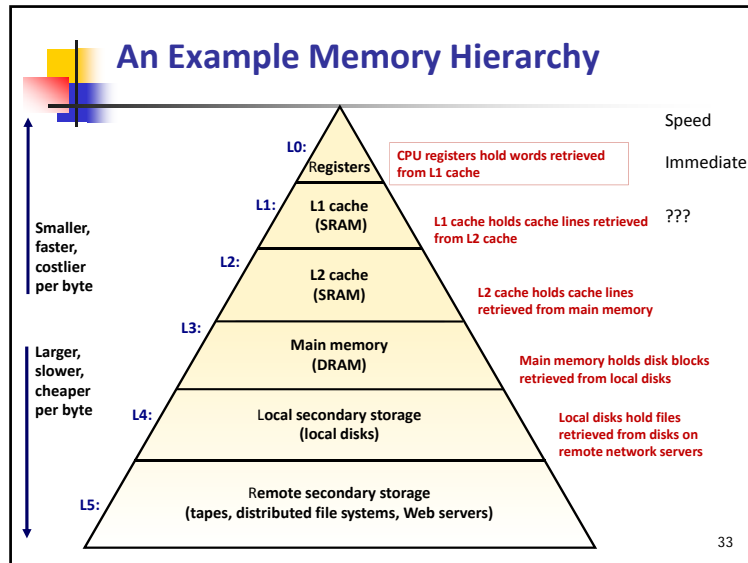
    for (j = 0; j < N; j++)
        for (i = 0; i < M; i++)
            sum += a[i][j];
    return sum;
}
```

31

Memory Hierarchies

- 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**.

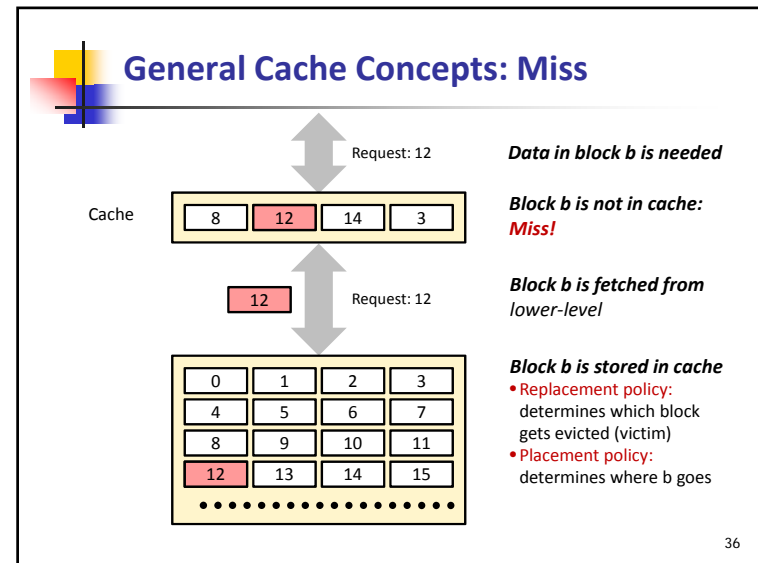
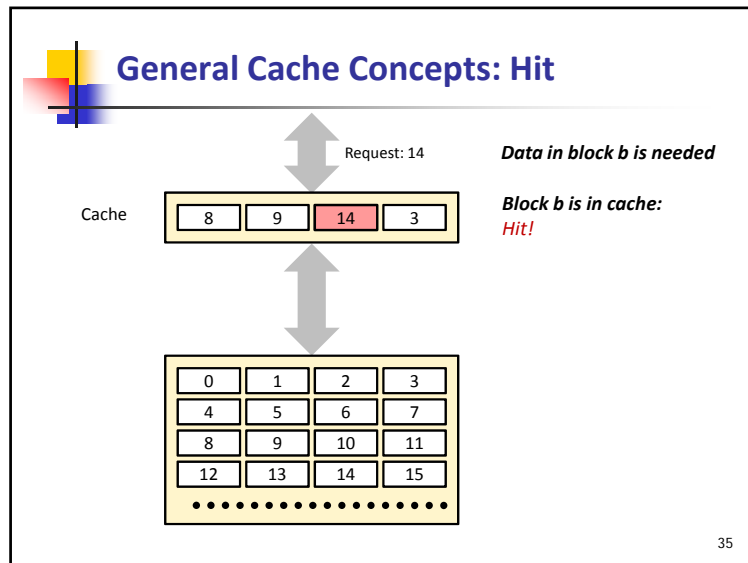
32




Caches

- Cache: 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 as a cache for the larger, slower device at level $k+1$.
- Why do memory hierarchies work?
 - Because of locality, most accesses are satisfied by cache.
 - So you get the storage space of lower-level storage but enjoy the speed of higher level cache.

34





Disclaimer

These slides were adapted from the CMU course slides provided along with the textbook of "Computer Systems: A programmer's Perspective" by Bryant and O'Hallaron. The slides are intended for the sole purpose of teaching the computer organization course at the University of Rochester.

37