Systems I

The Memory Hierarchy

Topics

- Storage technologies
- Capacity and latency trends
- The hierarchy

Random-Access Memory (RAM)

Key features

- RAM is packaged as a chip.
- Basic storage unit is a cell (one bit per cell).
- Multiple RAM chips form a memory.

Static RAM (SRAM)

- Each cell stores bit with a six-transistor circuit.
- Retains value indefinitely, as long as it is kept powered.
- Relatively insensitive to disturbances such as electrical noise.
- Faster and more expensive than DRAM.

Dynamic RAM (DRAM)

- Each cell stores bit with a capacitor and transistor.
- Value must be refreshed every 10-100 ms.
- Sensitive to disturbances, slower and cheaper than SRAM.

Flash RAM - it's in your ipod and cell phone

- Each cell stores 1 or more bits on a "floating-gate" capacitor
- Keeps state even when power is off
- As cheap as DRAM, but much slower

RAM Summary

	Tran. per bit	Access time	Persist?	Sensitive?	Cost	Applications
SRAM	6	1X	Yes	No	100x	cache memories
DRAM	1	10X	No	Yes	1X	Main memories, frame buffers
Flash	1/2-1	10000X	Yes	No	1X	Disk substitute

Conventional DRAM Organization

d x w DRAM:

- dw total bits organized as d supercells of size w bits

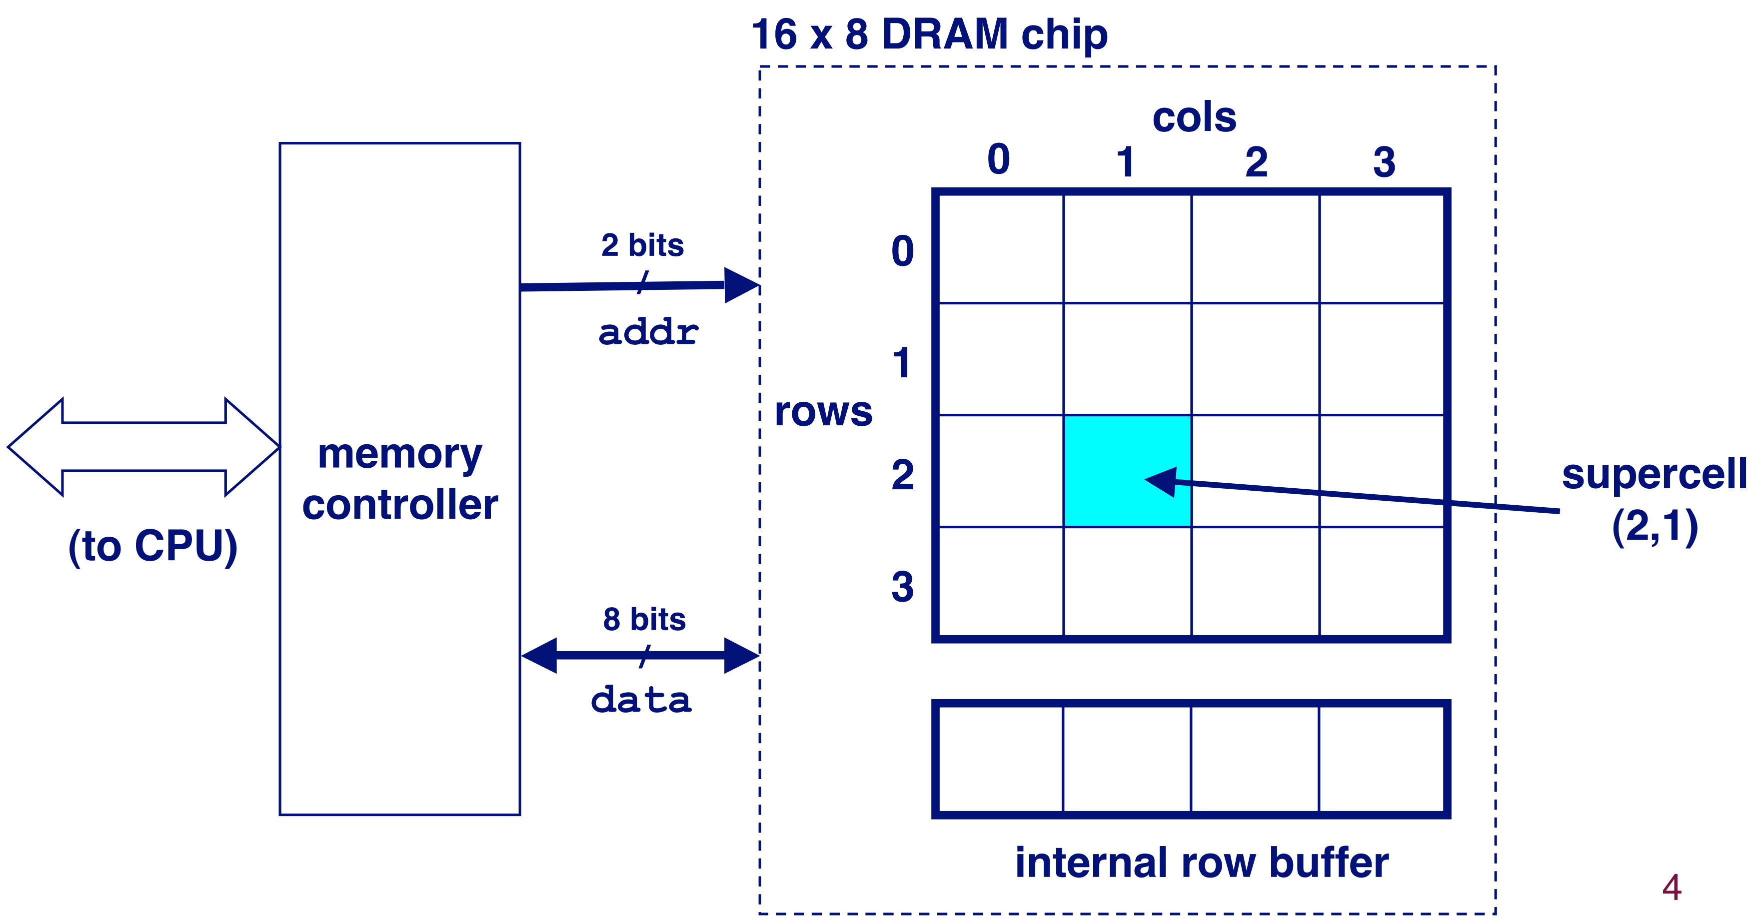

Reading DRAM Supercell (2,1)

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

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

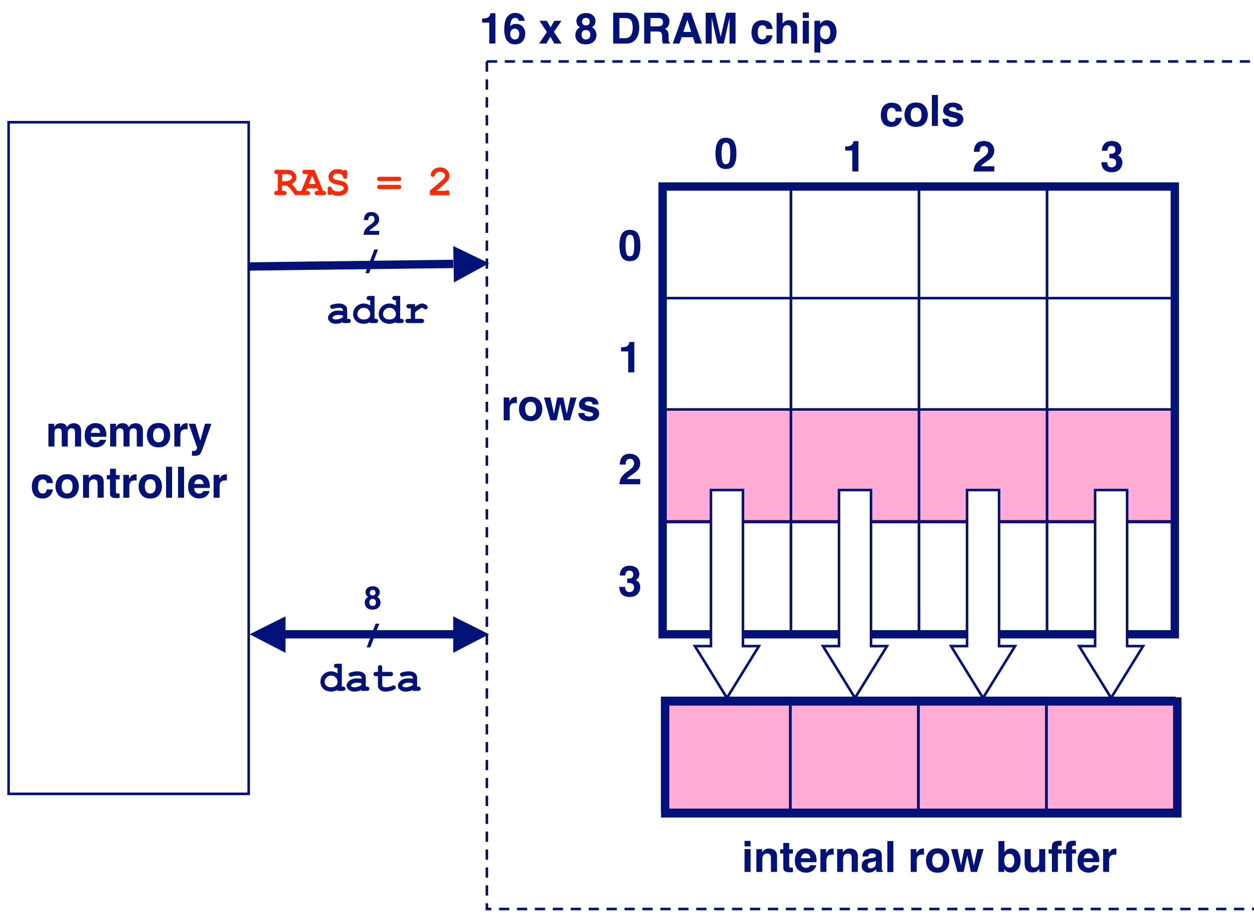

Reading DRAM Supercell (2,1)

Step 2(a): Column access strobe (CAS) selects column 1.

Step 2(b): Supercell (2,1) copied from buffer to data lines, and eventually back to the CPU.

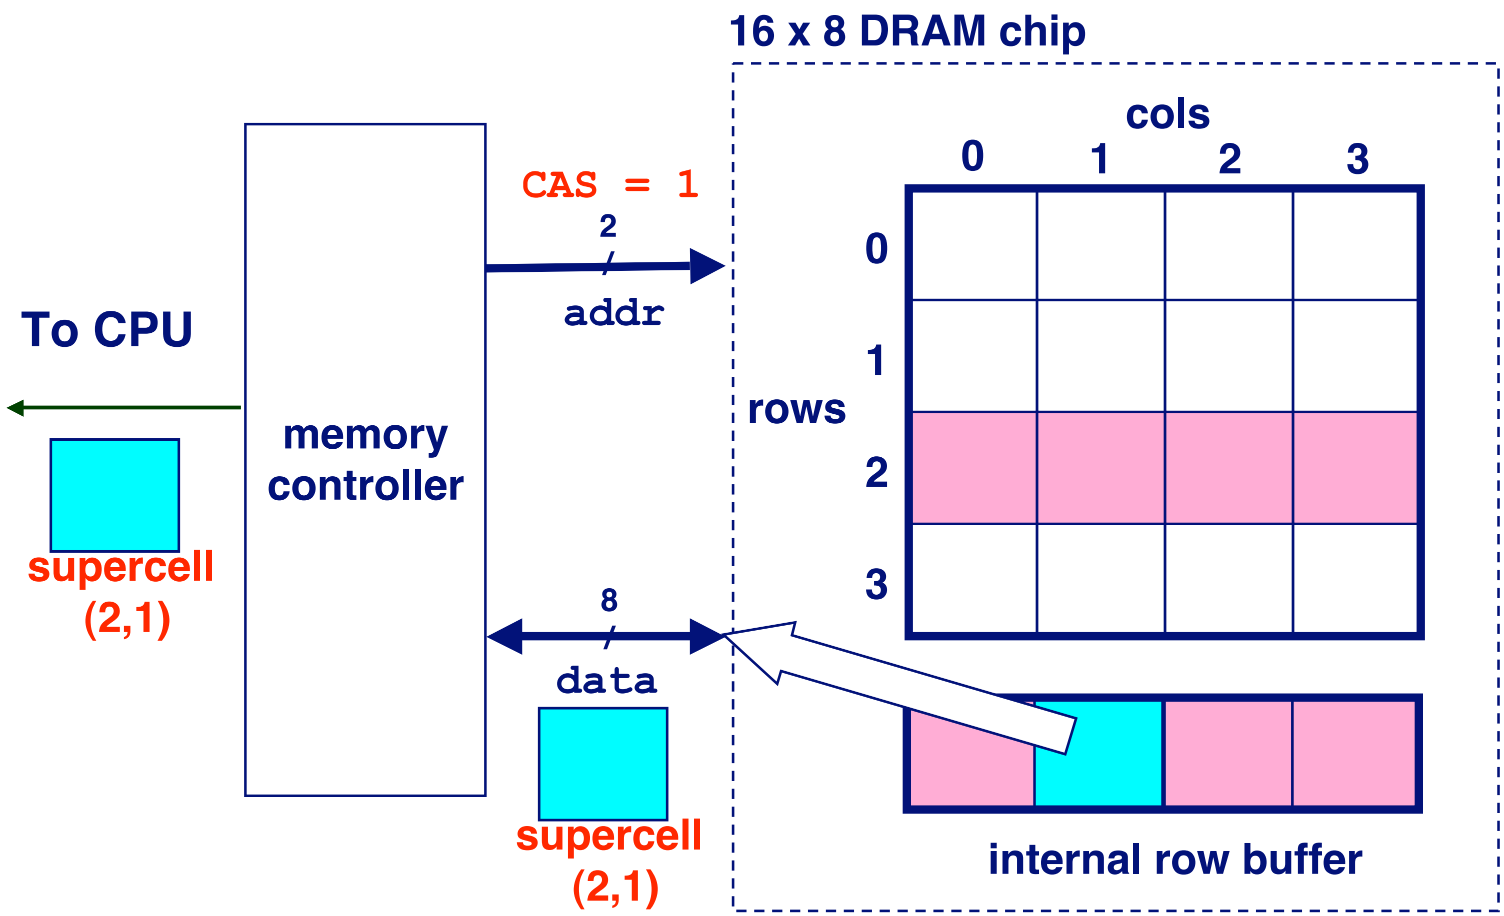

Memory Modules

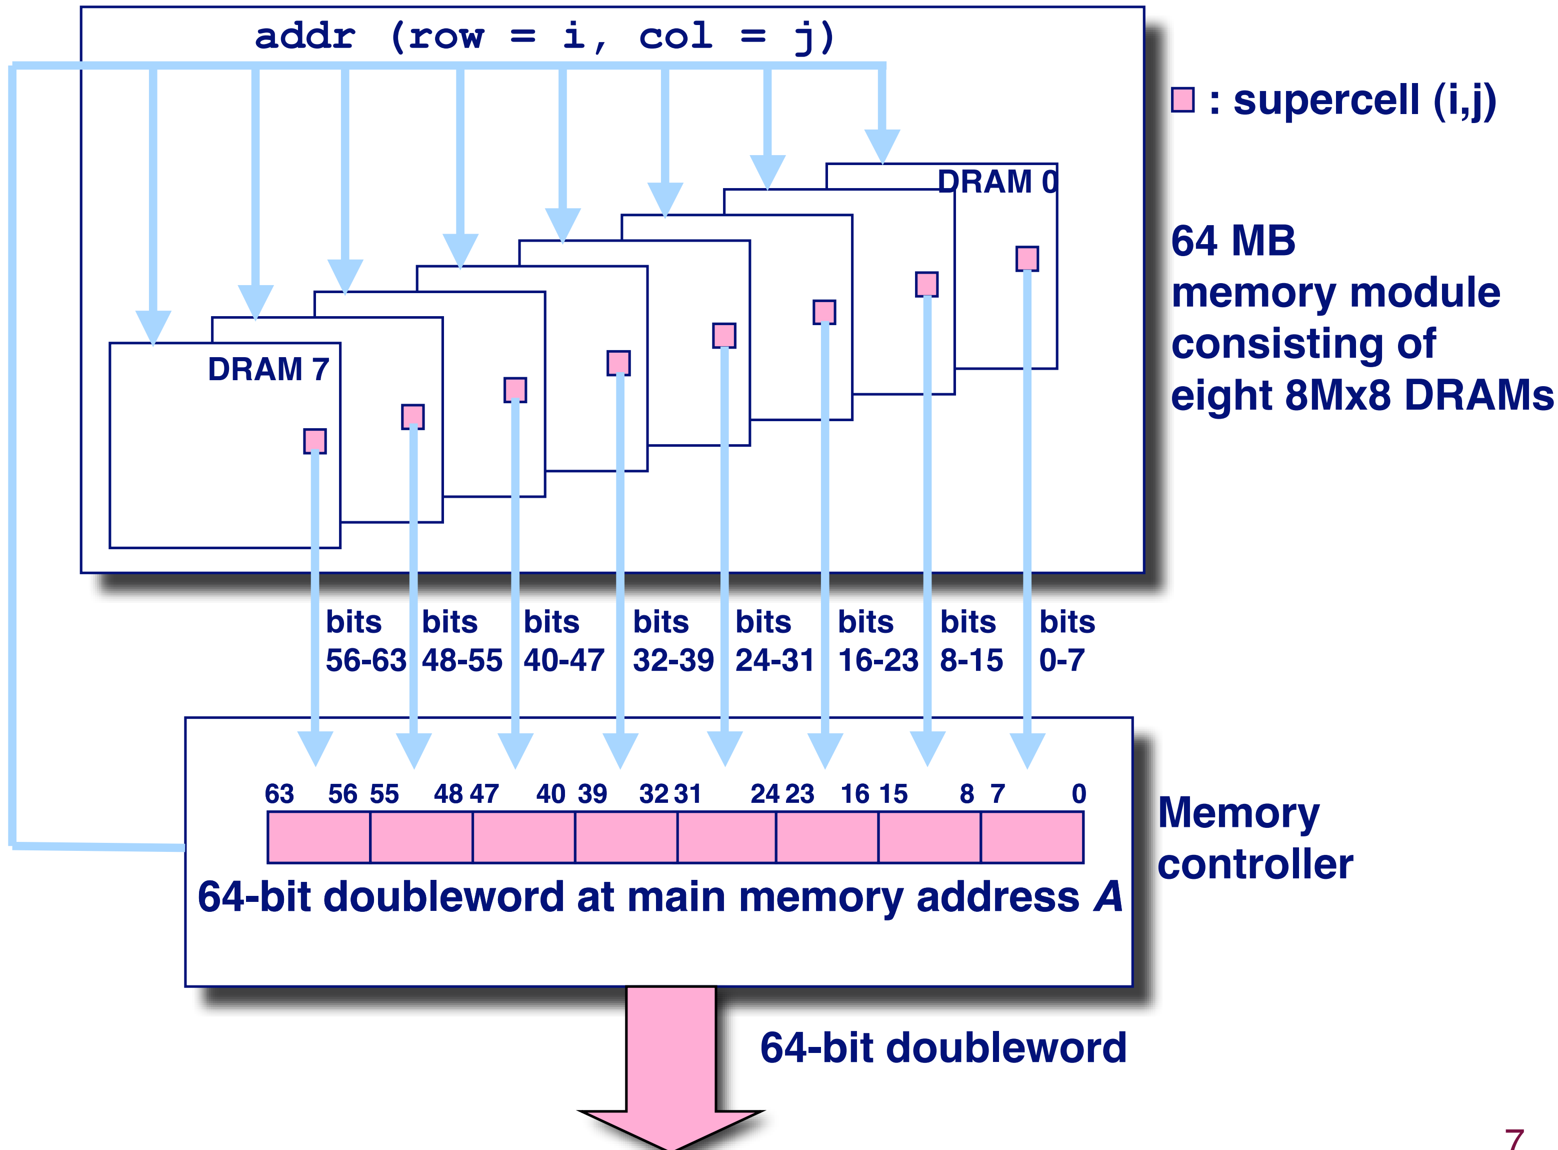

Enhanced DRAMs

All enhanced DRAMs are built around the conventional DRAM core.

- **Fast page mode DRAM (FPM DRAM)**
  - **Access contents of row with [RAS, CAS, CAS, CAS, CAS] instead of [(RAS,CAS), (RAS,CAS), (RAS,CAS), (RAS,CAS)].**
- **Extended data out DRAM (EDO DRAM)**
  - **Enhanced FPM DRAM with more closely spaced CAS signals.**
- **Synchronous DRAM (SDRAM)**
  - **Driven with rising clock edge instead of asynchronous control signals.**
- **Double data-rate synchronous DRAM (DDR SDRAM)**
  - **Enhancement of SDRAM that uses both clock edges as control signals.**
- **Video RAM (VRAM)**
  - **Like FPM DRAM, but output is produced by shifting row buffer**
  - **Dual ported (allows concurrent reads and writes)**

Nonvolatile Memories

DRAM and SRAM are volatile memories

- **Lose information if powered off.**

Nonvolatile memories retain value even if powered off.

- **Generic name is read-only memory (ROM).**
- **Misleading because some ROMs can be read and modified.**

Types of ROMs

- **Programmable ROM (PROM)**
- **Eraseable programmable ROM (EPROM)**
- **Electrically eraseable PROM (EEPROM)**
- **Flash memory**

Firmware

- **Program stored in a ROM**
  - **Boot time code, BIOS (basic input/ouput system)**
  - **graphics cards, disk controllers.**

Typical Bus Structure Connecting CPU and Memory

A bus is a collection of parallel wires that carry address, data, and control signals.

Buses are typically shared by multiple devices.

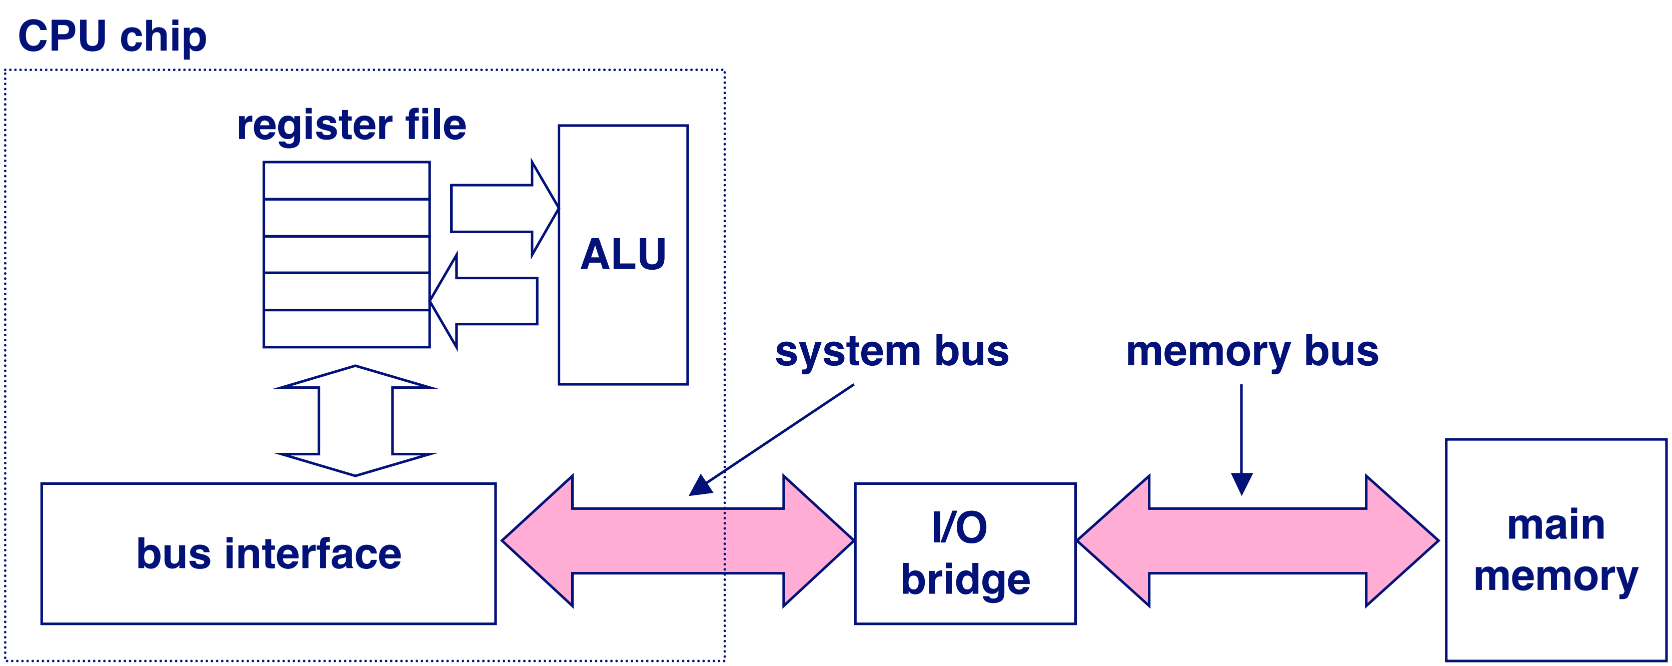

Memory Read Transaction (1)

CPU places address A on the memory bus.

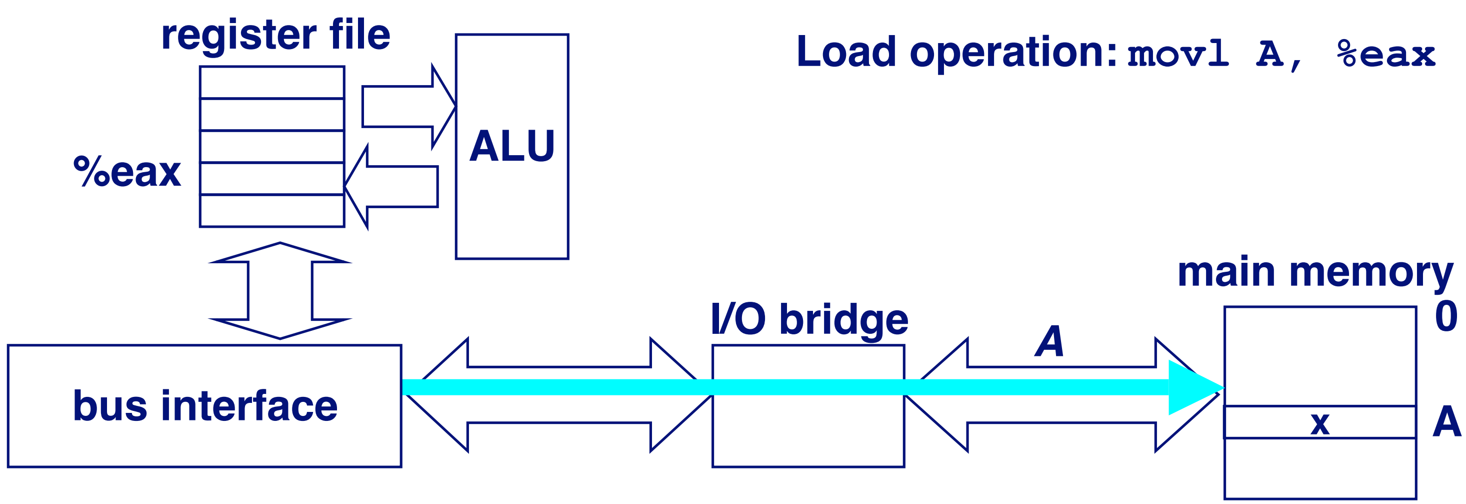

Memory Read Transaction (2)

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

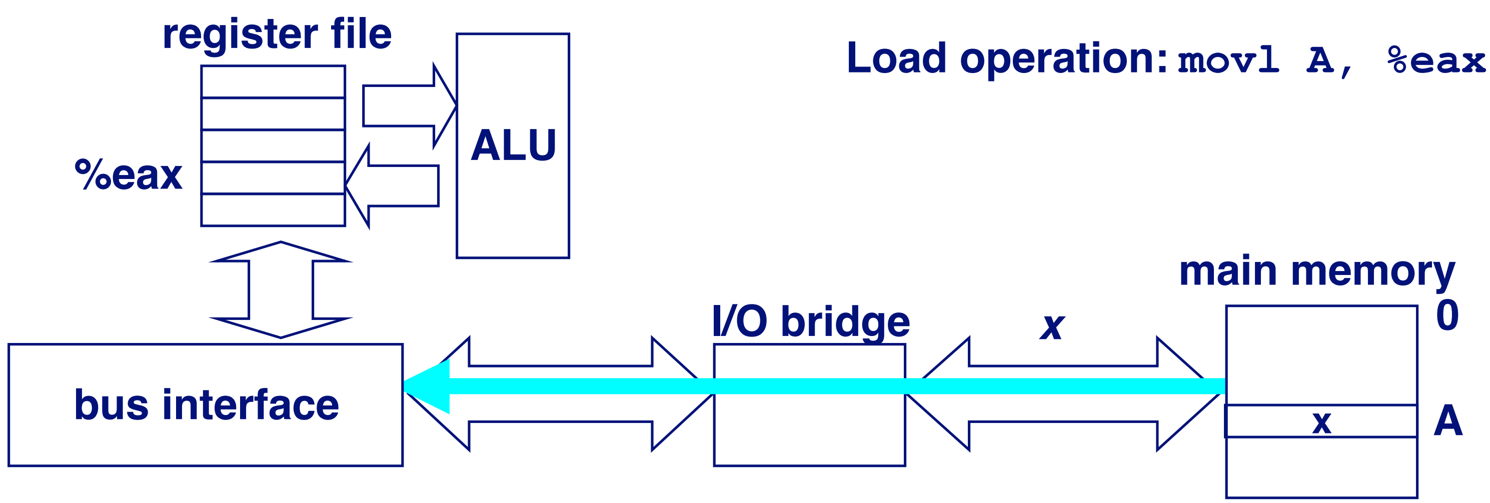

Memory Read Transaction (3)

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

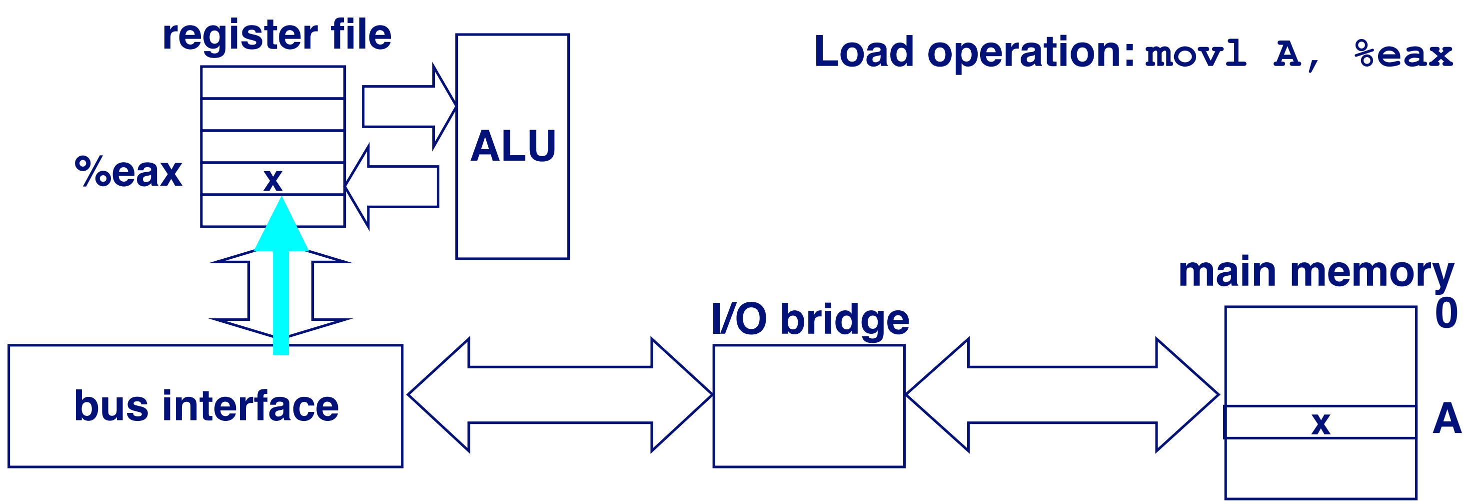

Memory Write Transaction (1)

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

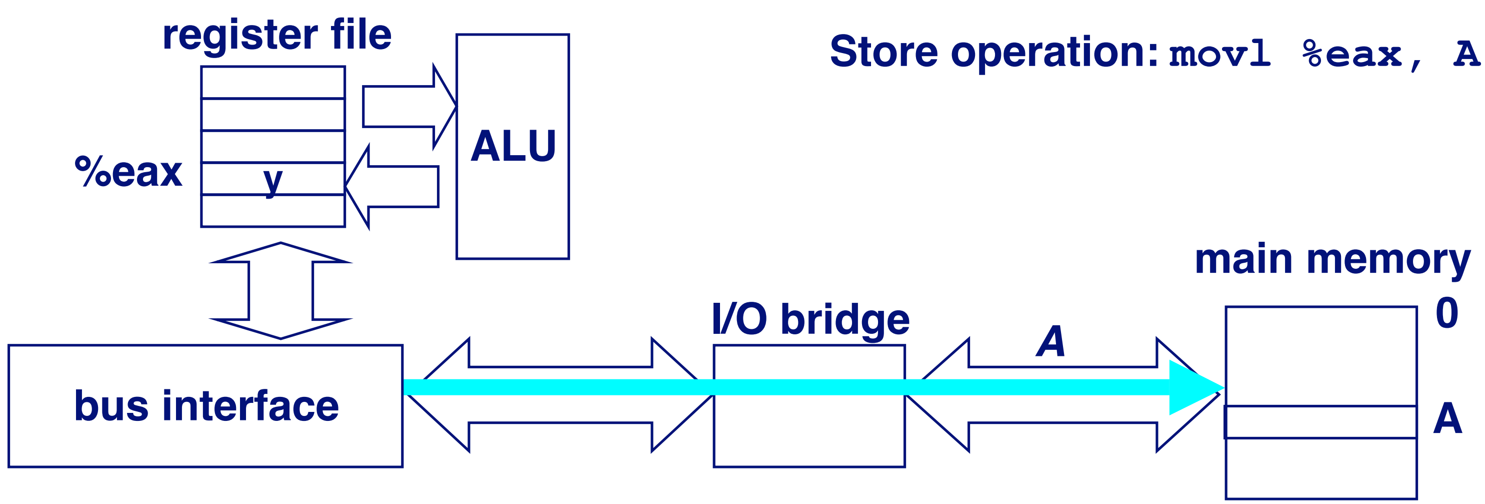

Memory Write Transaction (2)

CPU places data word y on the bus.

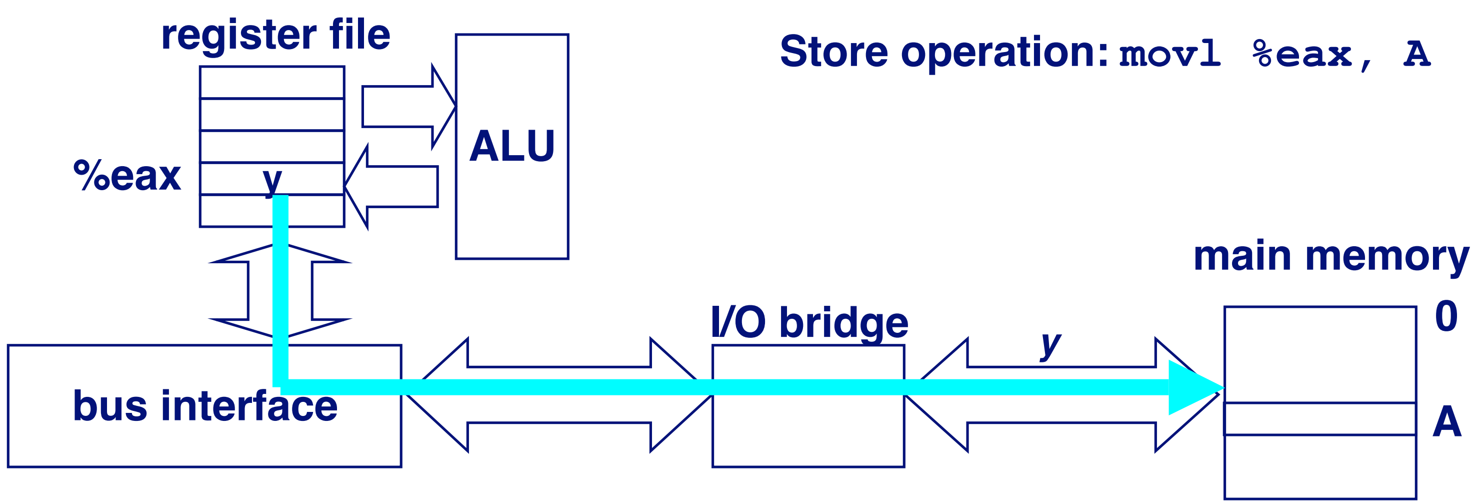

Memory Write Transaction (3)

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

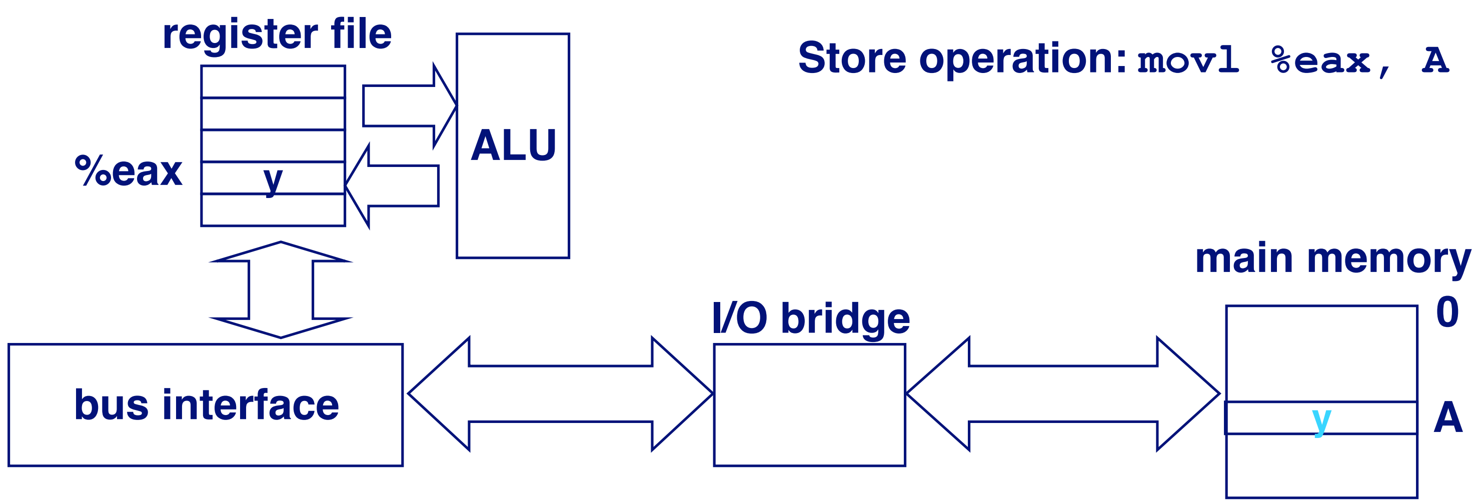

Disk Geometry

Disks consist of platters, each with two surfaces.

Each surface consists of concentric rings called tracks.

Each track consists of sectors separated by gaps.

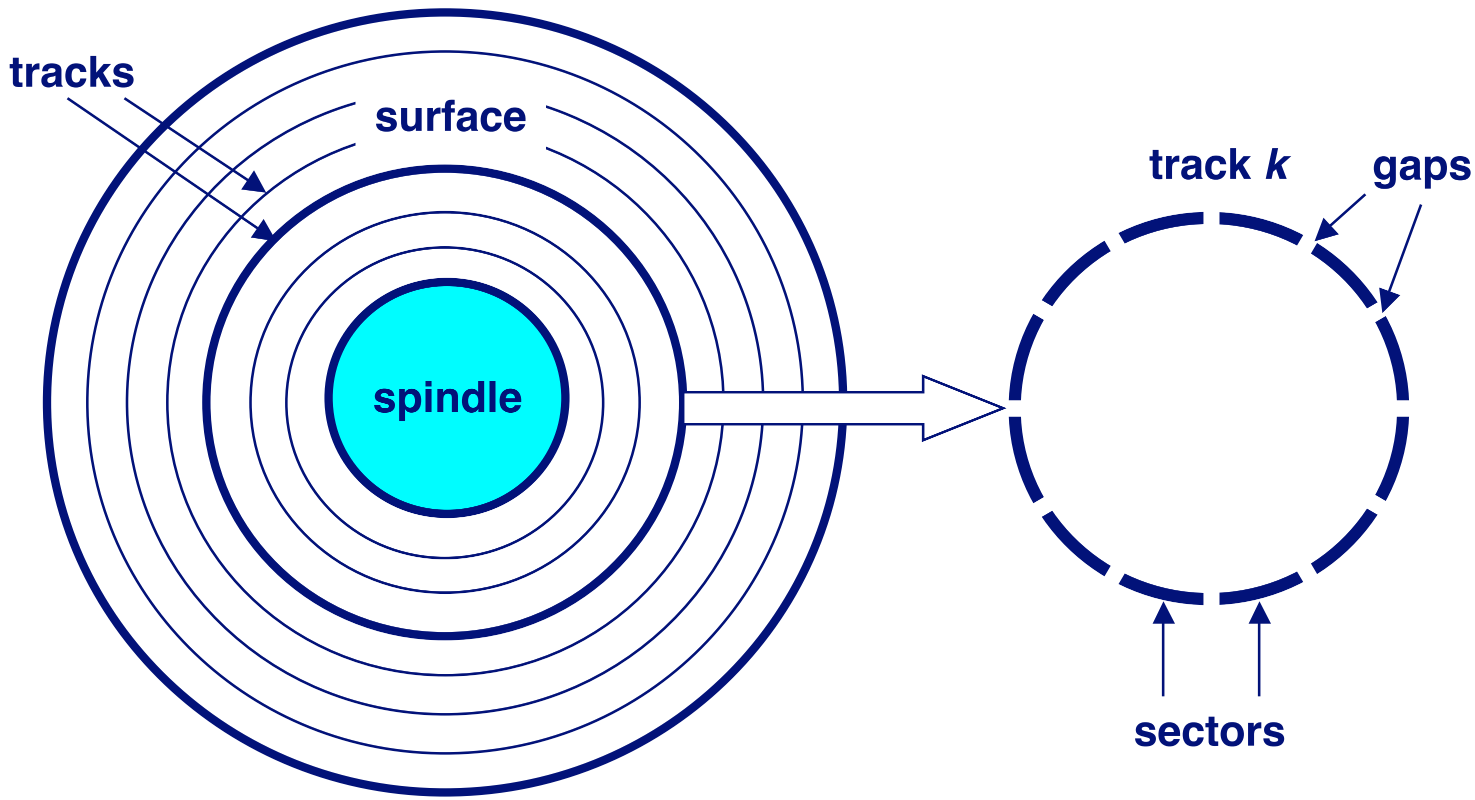

Disk Geometry (Multiple-Platter View)

Aligned tracks form a cylinder.

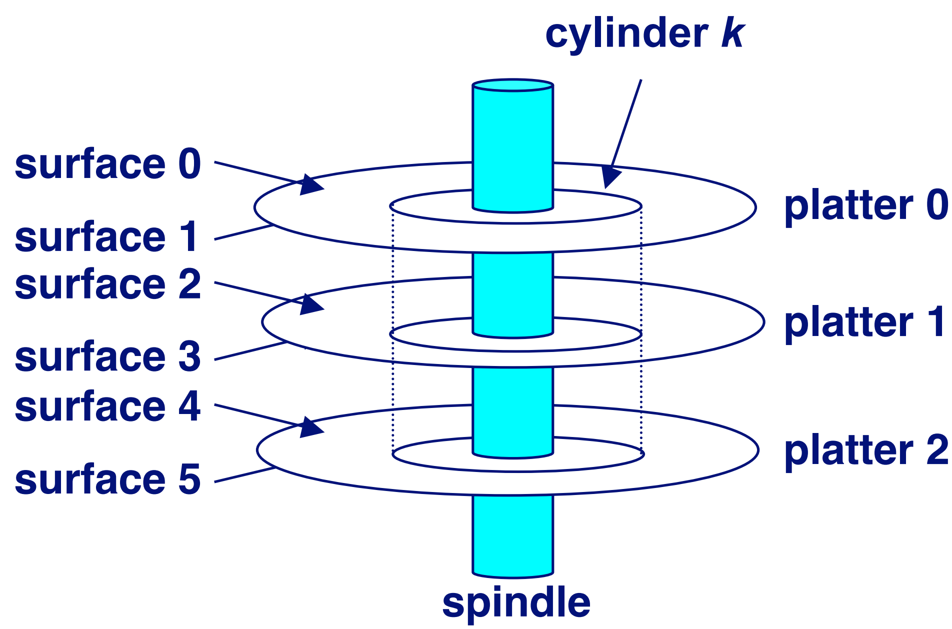

Disk Capacity

Capacity: maximum number of bits that can be stored.

- Vendors express capacity in units of gigabytes (GB), where 1 GB = 10^9.

Capacity is determined by these technology factors:

- Recording density (bits/in): number of bits that can be squeezed into a 1 inch segment of a track.
- Track density (tracks/in): number of tracks that can be squeezed into a 1 inch radial segment.
- Areal density (bits/in2): product of recording and track density.

Modern disks partition tracks into disjoint subsets called recording zones

- Each track in a zone has the same number of sectors, determined by the circumference of innermost track.
- Each zone has a different number of sectors/track

Computing Disk Capacity

Capacity = (# bytes/sector) x (avg. # sectors/track) x (# tracks/surface) x (# surfaces/platter) x (# platters/disk)

Example:

- 512 bytes/sector
- 300 sectors/track (on average)
- 20,000 tracks/surface
- 2 surfaces/platter
- 5 platters/disk

Capacity = 512 x 300 x 20000 x 2 x 5

= 30,720,000,000

= 30.72 GB

Disk Operation (Single-Platter View)

The disk surface spins at a fixed rotational rate

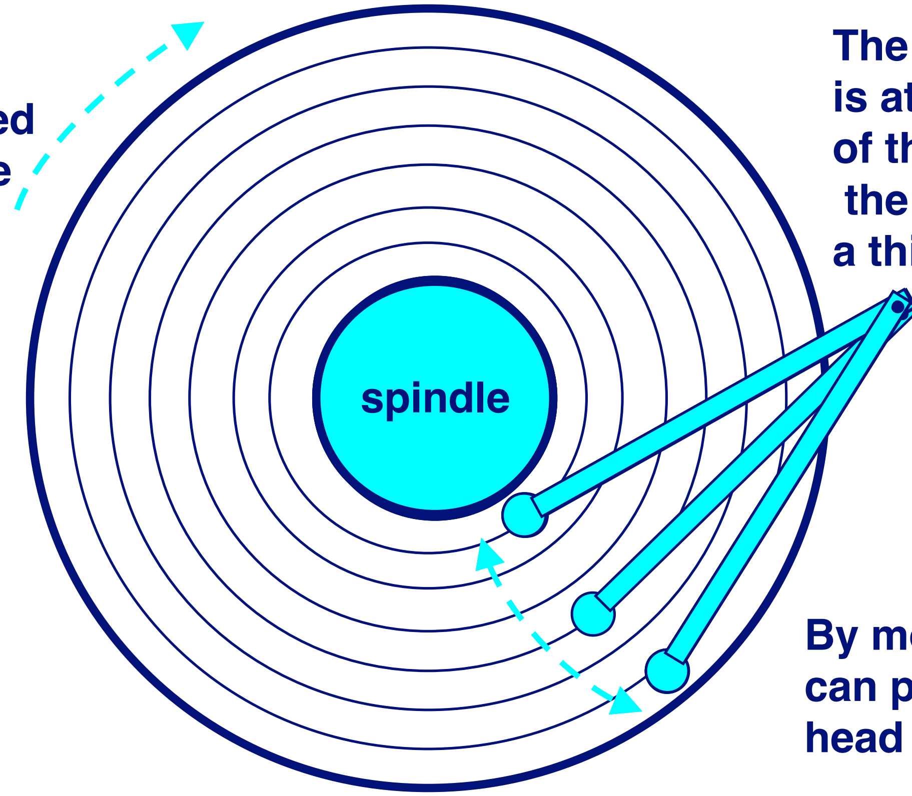

The read/write *head* is attached to the end of the *arm* and flies over the disk surface on a thin cushion of air.

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

Disk Operation (Multi-Platter View)

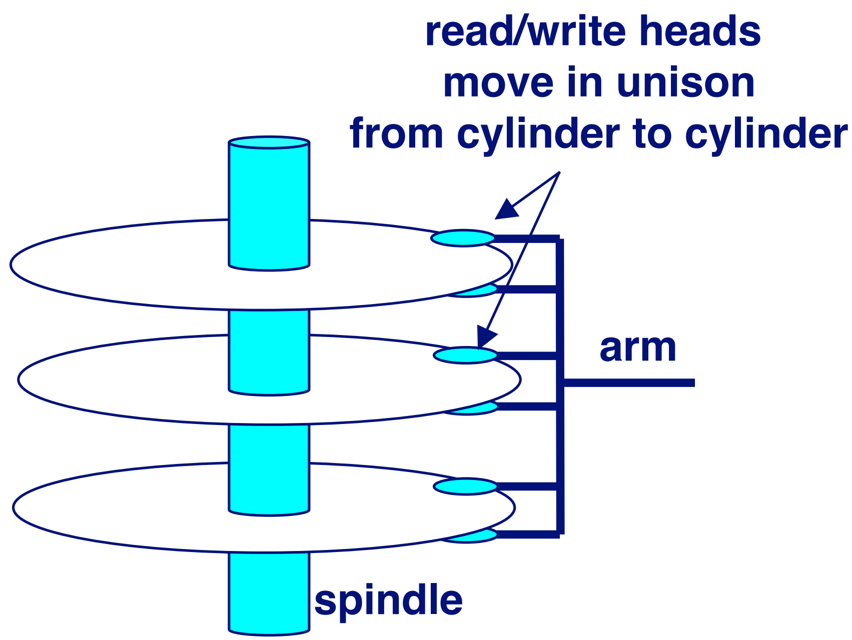

Disk Access Time

Average time to access some target sector approximated by :

- Taccess = Tavg seek + Tavg rotation + Tavg transfer

Seek time (Tavg seek)

- Time to position heads over cylinder containing target sector.
- Typical Tavg seek = 9 ms

Rotational latency (Tavg rotation)

- Time waiting for first bit of target sector to pass under r/w head.
- Tavg rotation = 1/2 x 1/RPMs x 60 sec/1 min

Transfer time (Tavg transfer)

- Time to read the bits in the target sector.
- Tavg transfer = 1/RPM x 1/(avg # sectors/track) x 60 secs/1 min.

Disk Access Time Example

Given:

- Rotational rate = 7,200 RPM
- Average seek time = 9 ms.
- Avg # sectors/track = 400.

Derived:

- Tavg rotation = 1/2 x (60 secs/7200 RPM) x 1000 ms/sec = 4 ms.
- Tavg transfer = 60/7200 RPM x 1/400 secs/track x 1000 ms/sec = 0.02 ms
- Taccess = 9 ms + 4 ms + 0.02 ms

Important points:

- Access time dominated by seek time and rotational latency.
- First bit in a sector is the most expensive, the rest are free.
- SRAM access time is about 4 ns/doubleword, DRAM about 60 ns
  - Disk is about 40,000 times slower than SRAM,
  - 2,500 times slower then DRAM.

Logical Disk Blocks

Modern disks present a simpler abstract view of the complex sector geometry:

- The set of available sectors is modeled as a sequence of b-sized logical blocks (0, 1, 2, ...)

Mapping between logical blocks and actual (physical) sectors

- Maintained by hardware/firmware device called disk controller.
- Converts requests for logical blocks into (surface,track,sector) triples.

Allows controller to set aside spare cylinders for each zone.

- Accounts for the difference in "formatted capacity" and "maximum capacity".

I/O Bus

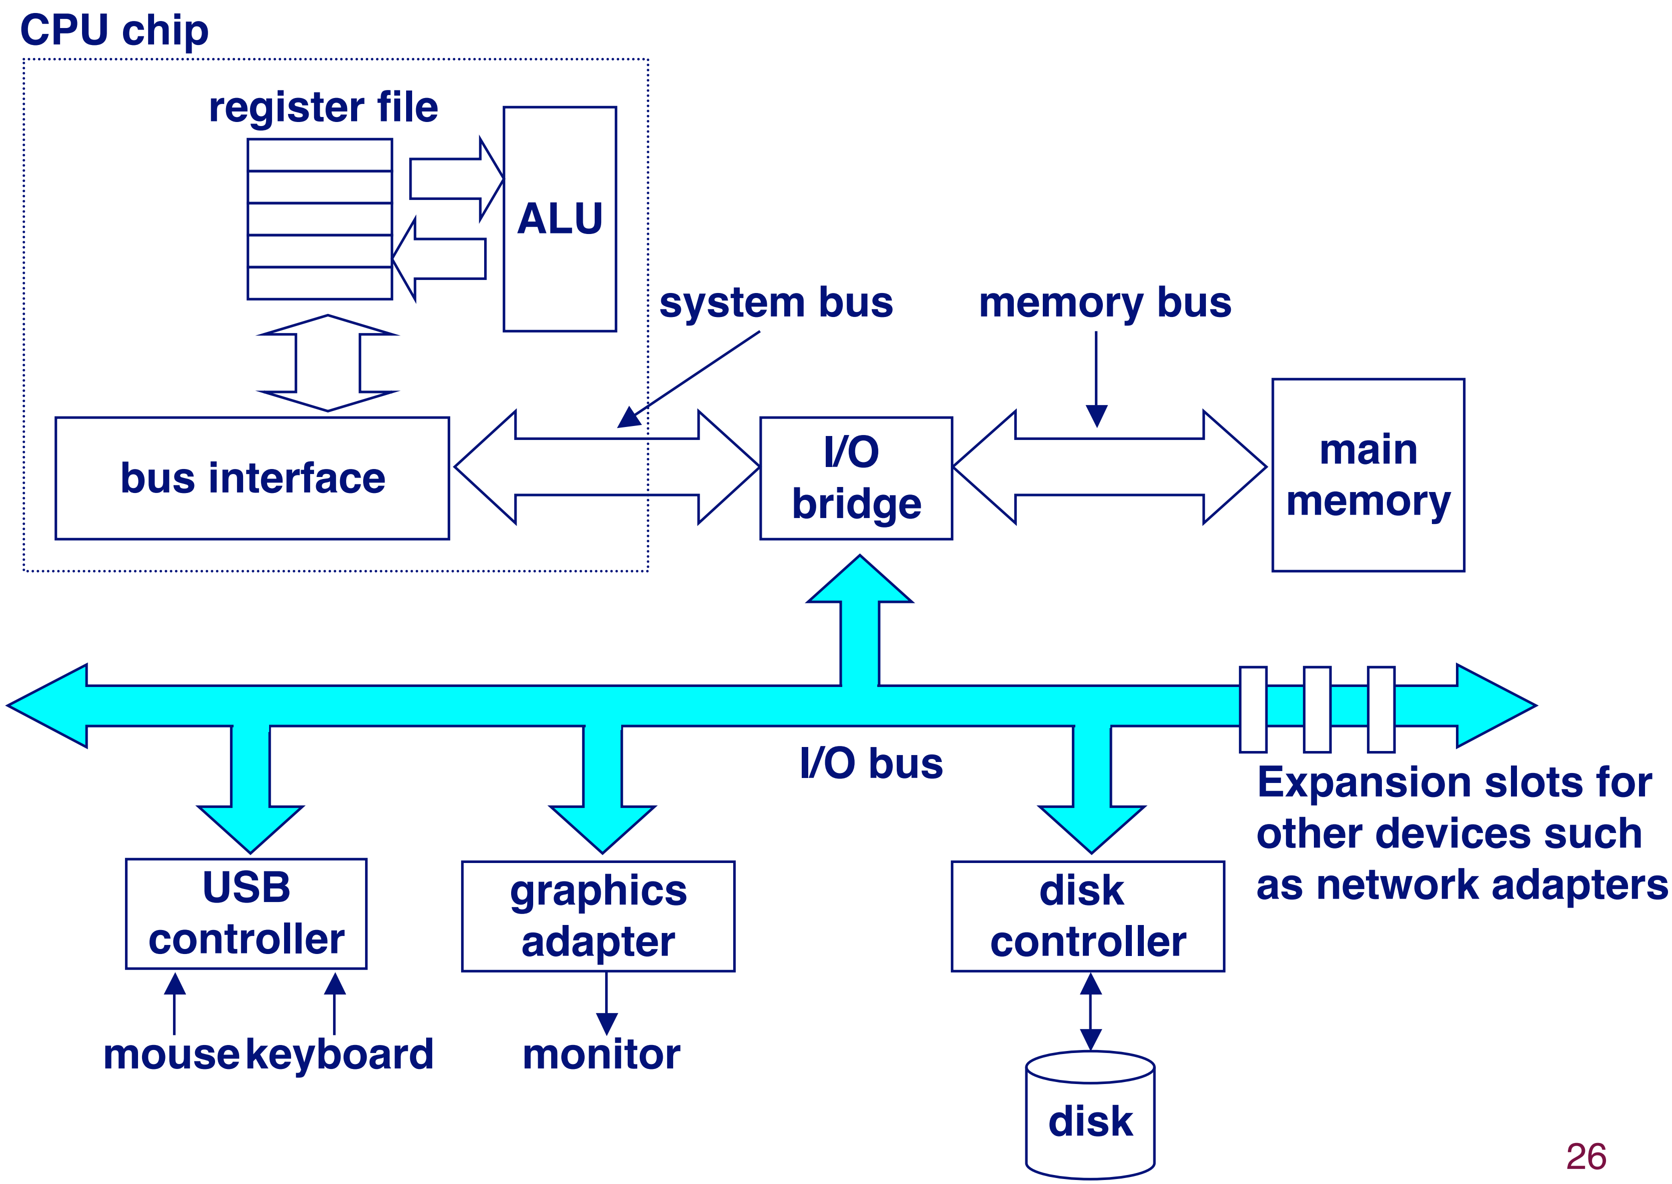

CPU chip
register file
ALU
system bus
memory bus
bus interface
I/O
bridge
main
memory
I/O bus
Expansion slots for
other devices such
as network adapters.
USB
controller
graphics
adapter
disk
controller
mouse
keyboard
monitor
disk

Reading a Disk Sector (1)

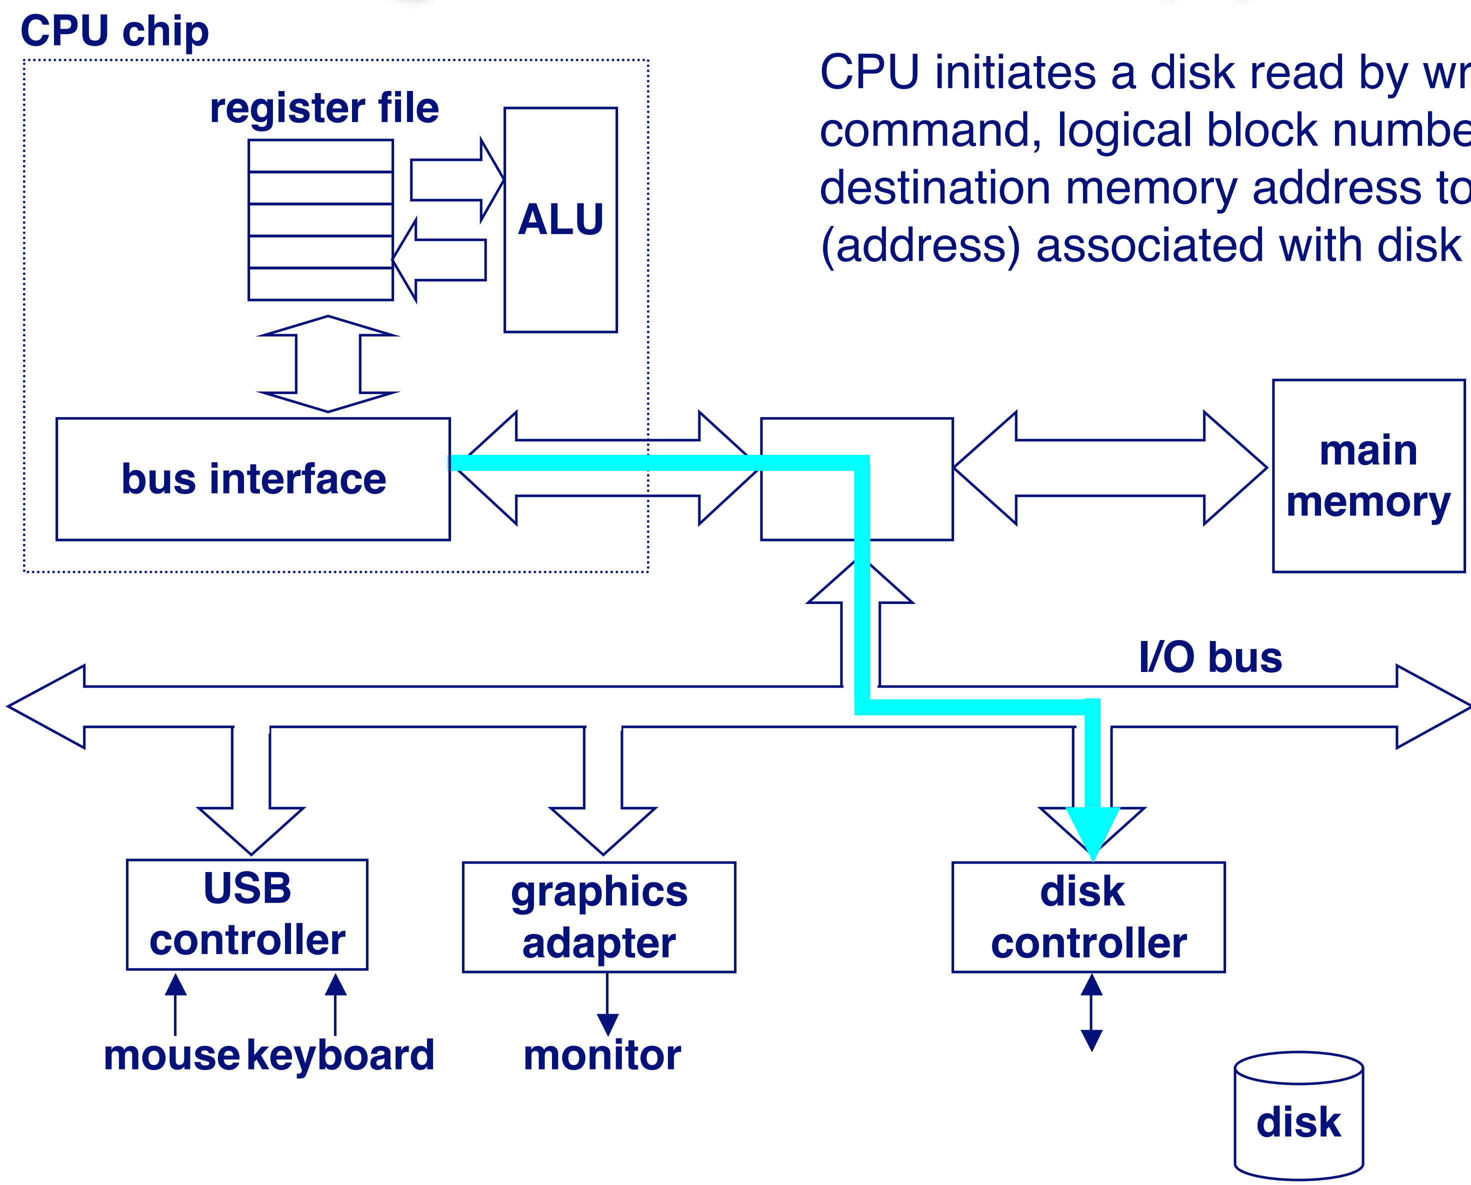

CPU initiates a disk read by writing a command, logical block number, and destination memory address to a port (address) associated with disk controller.

Reading a Disk Sector (2)

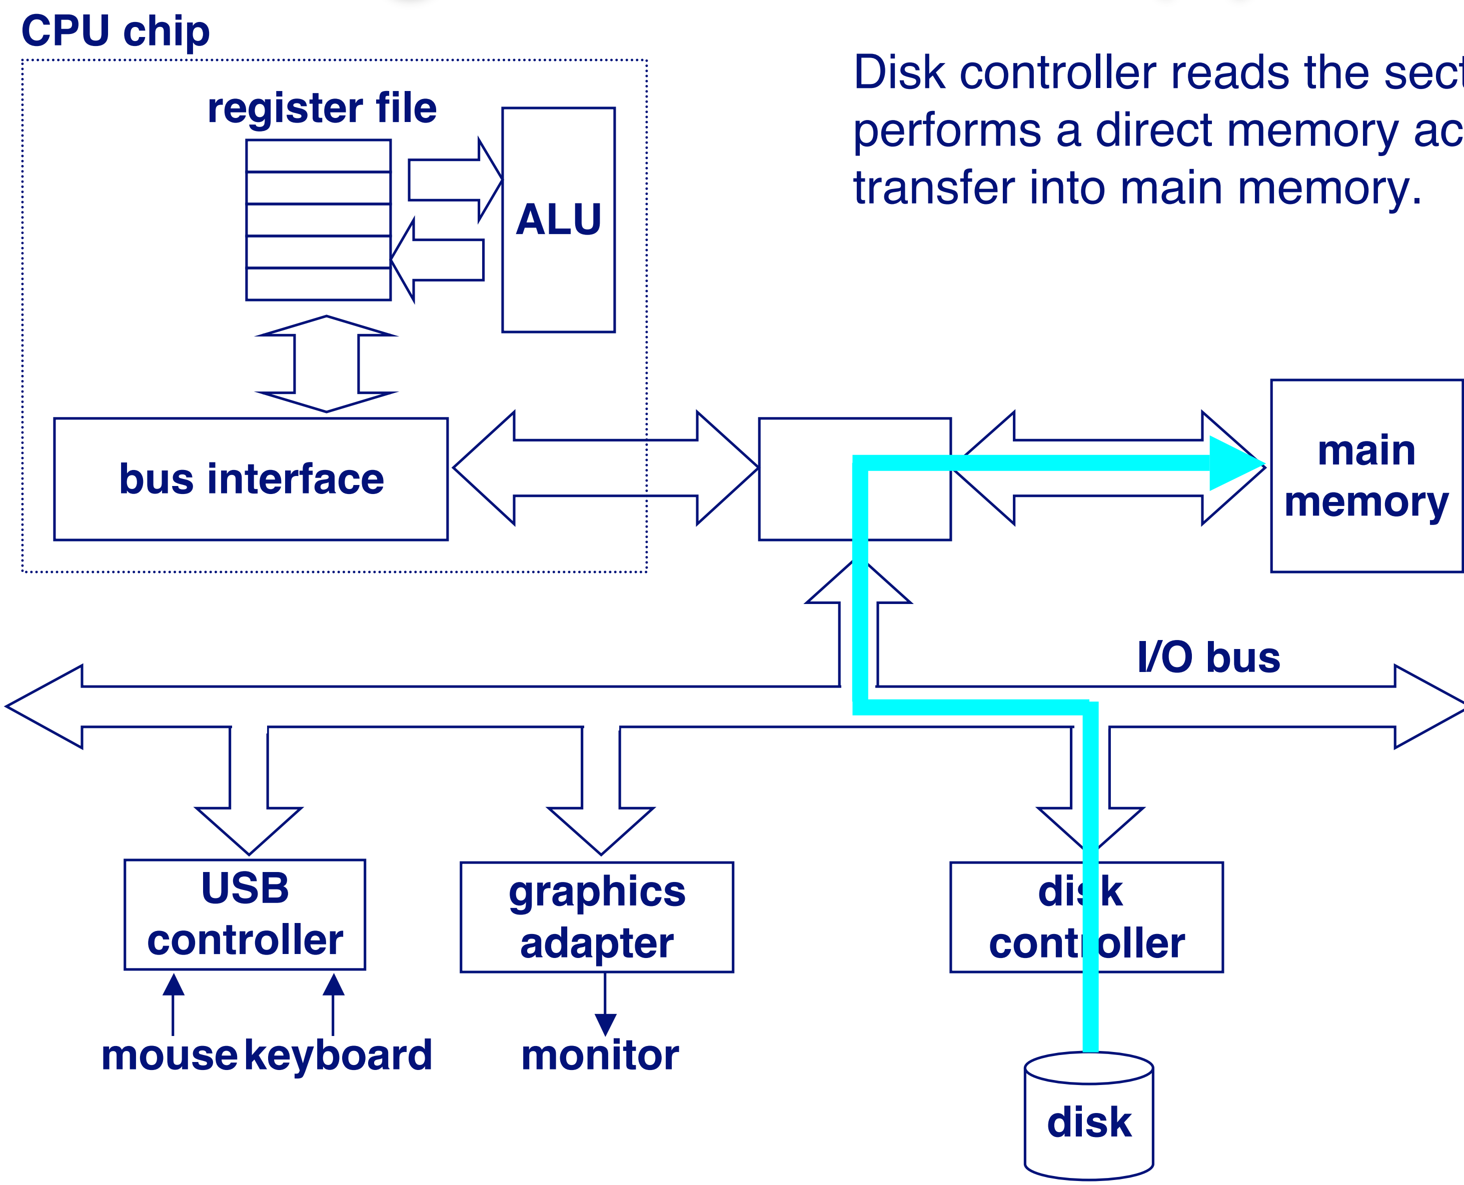

Disk controller reads the sector and performs a direct memory access (DMA) transfer into main memory.

Reading a Disk Sector (3)

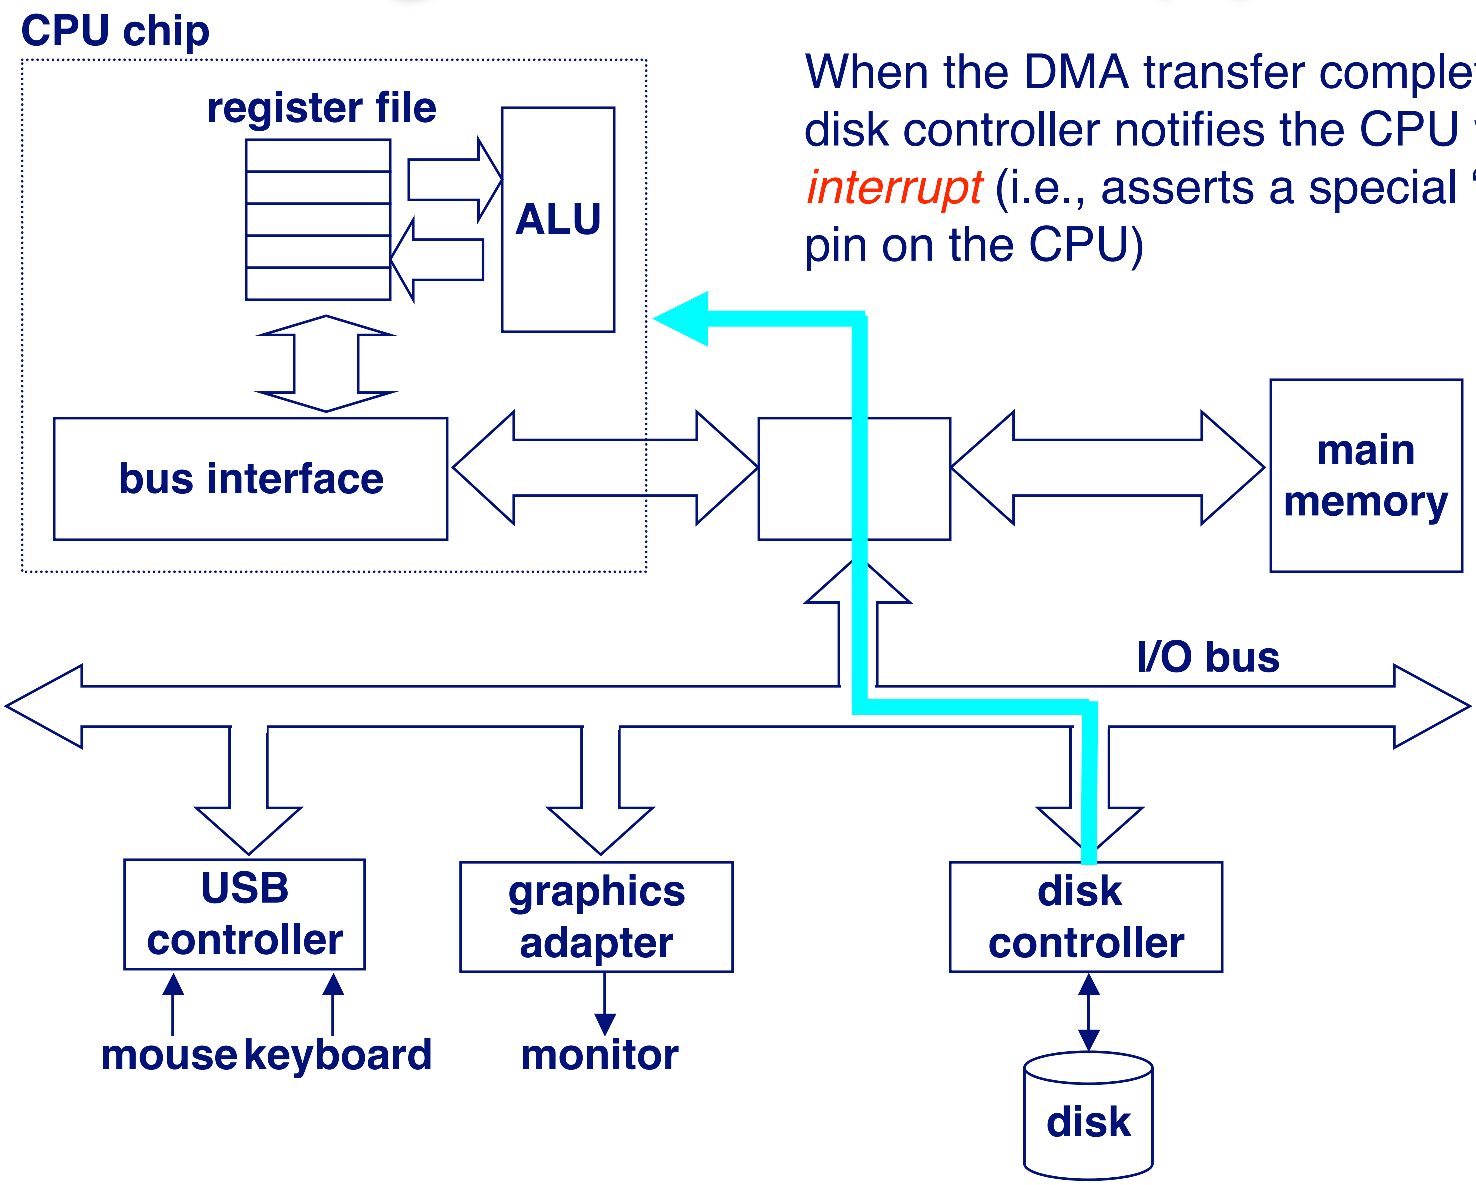

When the DMA transfer completes, the disk controller notifies the CPU with an *interrupt* (i.e., asserts a special “interrupt” pin on the CPU)

Storage Trends

SRAM

metric	1980	1985	1990	1995	2000	*2000:1980*
$/MB	19,200	2,900	320	256	100	*190*
access (ns)	300	150	35	15	2	*100*

DRAM

metric	1980	1985	1990	1995	2000	*2000:1980*
$/MB	8,000	880	100	30	1	*8,000*
access (ns)	375	200	100	70	60	*6*
typical size(MB)	0.064	0.256	4	16	64	*1,000*

Disk

metric	1980	1985	1990	1995	2000	*2000:1980*
$/MB	500	100	8	0.30	0.05	*10,000*
access (ms)	87	75	28	10	8	*11*
typical size(MB)	1	10	160	1,000	9,000	*9,000*

(Culled from back issues of Byte and PC Magazine)

CPU Clock Rates

	1980	1985	1990	1995	2000	*2000:1980*
processor	8080	286	386	Pent	P-III	
clock rate(MHz)	1	6	20	150	750	750
cycle time(ns)	1,000	166	50	6	1.6	750

The CPU-Memory Gap

The increasing gap between DRAM, disk, and CPU speeds.

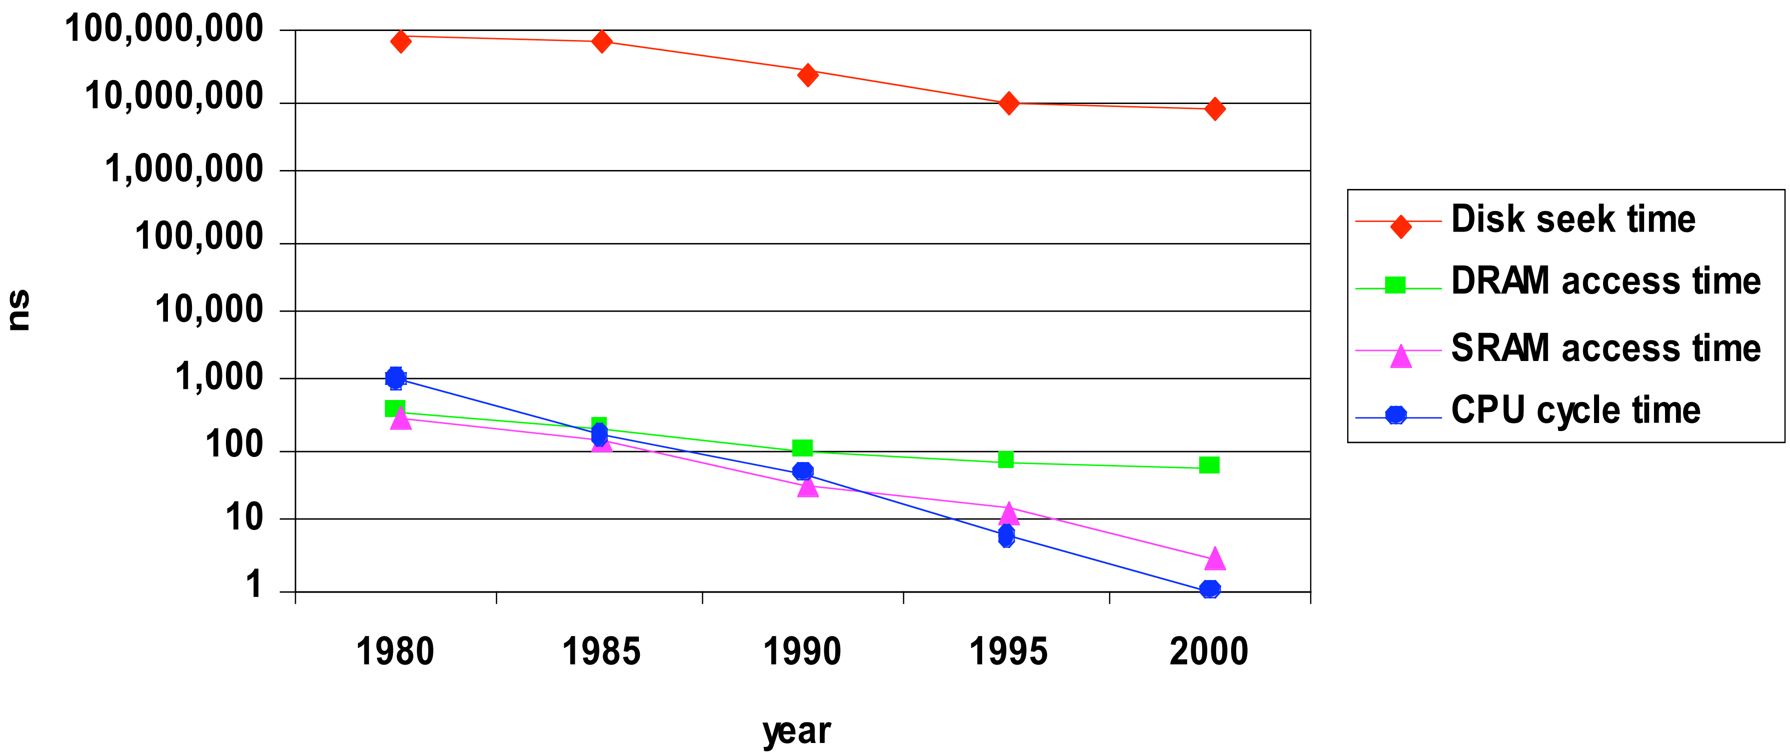

An Example Memory Hierarchy

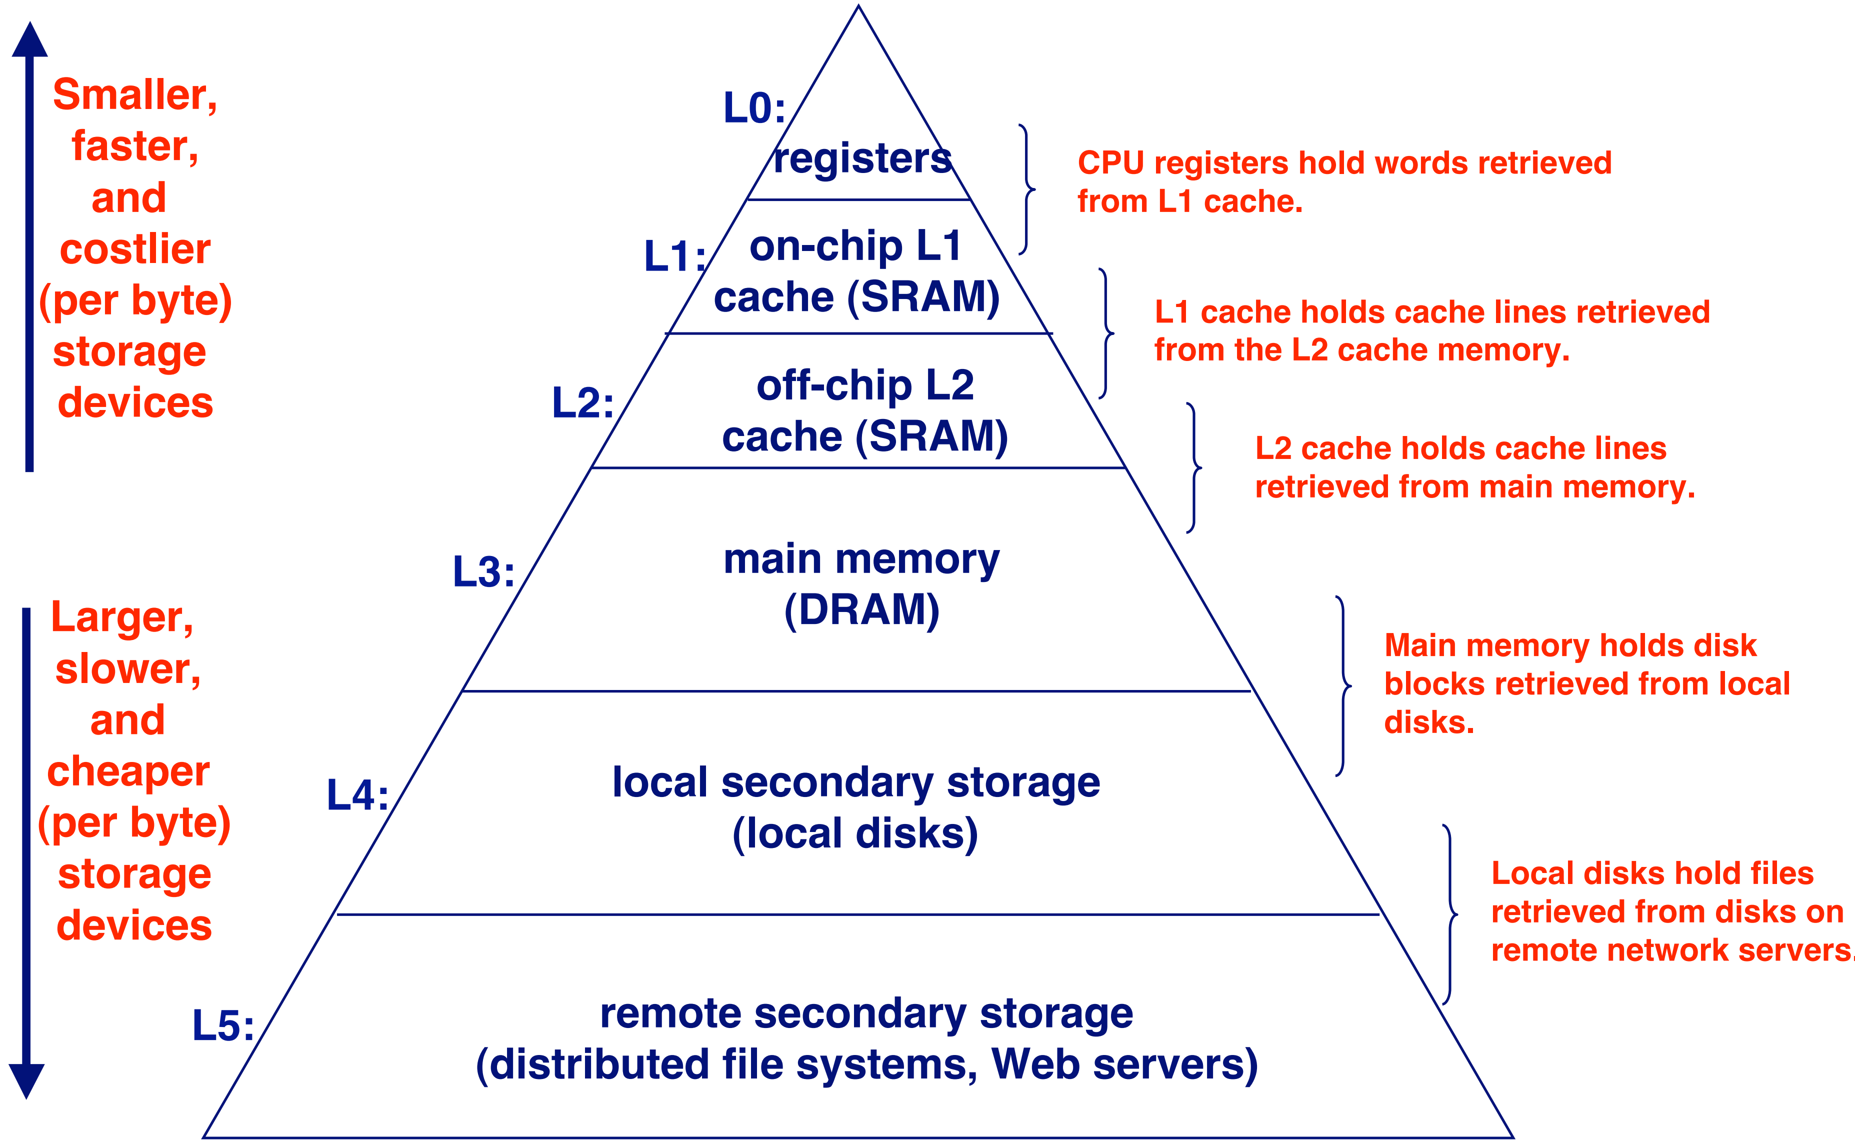

Summary

Today

- Memory and storage technologies
- Trends
- Hierarchy of capacity and latency

Next time

- Principles of locality
- Cache architectures