

Last Class

- Signals
 - Types and interface
 - Signal handlers and concurrency issues

1

Today


- Memory and storage technologies
- Virtual memory

Quiz 8 due Monday

2

Random-Access Memory (RAM)

- Key features
 - RAM is packaged in a chip (potentially as part of a processor)
 - Basic storage unit is a **cell** (one bit per cell)
 - Multiple RAM chips may be used to 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



3

SRAM vs DRAM 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

4

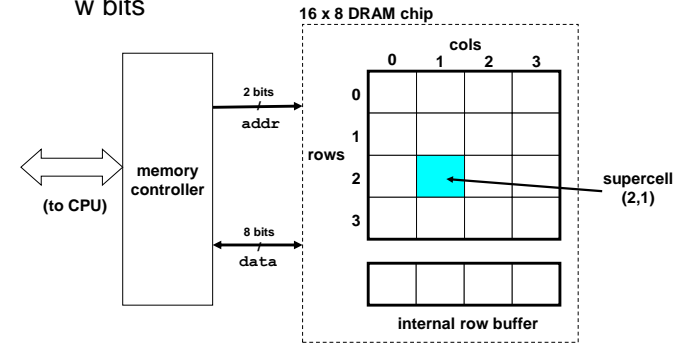
Nonvolatile Memories

- 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 (block-level) 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
 - More to come ...

5

Conventional DRAM Organization

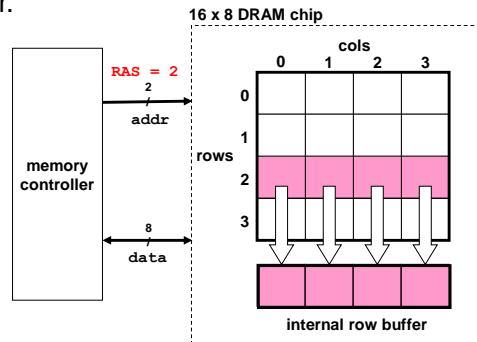
- $d \times w$ DRAM:
 - dw total bits organized as d **supercells** of size w bits



6

Reading DRAM Supercell (2,1)

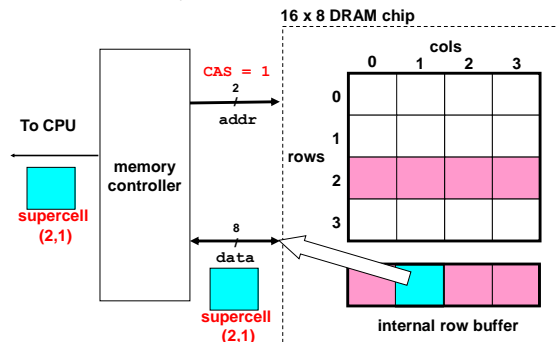
- Step 1(a): Row access strobe (**RAS**) selects row 2.
- Step 1(b): Row 2 copied from DRAM array to row buffer.



7

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.



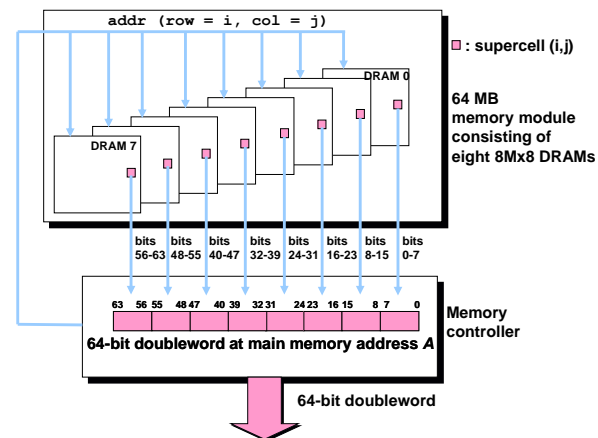
8

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)

9

Memory Modules



10


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
 - Phase Change Memory (PCM)
- Firmware**
 - Program stored in a ROM
 - Boot time code, BIOS (basic input/output system)
 - graphics cards, disk controllers
- Replacing volatile memory?

11

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



12

Memory Read Transaction (1)


- CPU places address A on the memory bus.



13

Memory Read Transaction (2)


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



14

Memory Read Transaction (3)


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



15

Memory Write Transaction (1)


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



16

Memory Write Transaction (2)


- CPU places data word y on the bus.



17

Memory Write Transaction (3)

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



18

What's Inside A Disk Drive?





Image courtesy of Seagate Technology

19

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**.



20

Disk Geometry (Multiple-Platter View)

- Aligned tracks form a cylinder.



21


Disk Capacity

- Capacity:** maximum number of bits that can be stored.
 - Vendors express capacity in units of gigabytes (GB), where $1 \text{ GB} = 10^9 \text{ Bytes}$.
- 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.

22

Recording zones

- 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, outer zones have more sectors/track than inner zones.
 - So we use **average** number of sectors/track when computing capacity.



23

Computing Disk Capacity

$$\text{Capacity} = (\# \text{ bytes/sector}) \times (\text{avg. } \# \text{ sectors/track}) \times (\# \text{ tracks/surface}) \times (\# \text{ surfaces/platter}) \times (\# \text{ platters/disk})$$

Example:


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

$$\begin{aligned} \text{Capacity} &= 512 \times 300 \times 20000 \times 2 \times 5 \\ &= 30,720,000,000 \\ &= 30.72 \text{ GB} \end{aligned}$$

24

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.


25

Disk Operation (Multi-Platter View)



26

Disk Structure - top view of single platter




Surface organized into tracks

Tracks divided into sectors

27


Disk Access



Head in position above a track

28


Disk Access



Rotation is counter-clockwise

29


Disk Access – Read



About to read blue sector

30

Disk Access – Read




After BLUE
read

After reading blue sector

31

Disk Access – Read




After BLUE
read

Red request scheduled next

32

Disk Access – Seek




After BLUE read Seek for RED

Seek to red's track

33

Disk Access – Rotational Latency




After BLUE read Seek for RED Rotational latency

Wait for red sector to rotate around

34

Disk Access – Read




After BLUE read Seek for RED Rotational latency After RED read

Complete read of red

35

Disk Access – Service Time Components



After BLUE read Seek for RED Rotational latency After RED read

Data transfer Seek Rotational latency Data transfer

36

Disk Access Time

- Average time to access some target sector approximated by :
 - $T_{\text{access}} = T_{\text{avg seek}} + T_{\text{avg rotation}} + T_{\text{avg transfer}}$
- **Seek time** ($T_{\text{avg seek}}$)
 - Time to position heads over cylinder containing target sector.
 - Typical $T_{\text{avg seek}}$ is 3—9 ms
- **Rotational latency** ($T_{\text{avg rotation}}$)
 - Time waiting for first bit of target sector to pass under r/w head.
 - $T_{\text{avg rotation}} = 1/2 \times 1/\text{RPMs} \times 60 \text{ sec}/1 \text{ min}$
 - Typical $T_{\text{avg rotation}} = 7200 \text{ RPMs}$
- **Transfer time** ($T_{\text{avg transfer}}$)
 - Time to read the bits in the target sector.
 - $T_{\text{avg transfer}} = 1/\text{RPM} \times 1/(\text{avg \# sectors/track}) \times 60 \text{ secs}/1 \text{ min.}$

37

Disk Access Time Example

- Given:
 - Rotational rate = 7,200 RPM
 - Average seek time = 9 ms.
 - Avg # sectors/track = 400.
- Derived:
 - $T_{\text{avg rotation}} = 1/2 \times (60 \text{ secs}/7200 \text{ RPM}) \times 1000 \text{ ms/sec} = 4 \text{ ms.}$
 - $T_{\text{avg transfer}} = 60/7200 \text{ RPM} \times 1/400 \text{ secs/track} \times 1000 \text{ ms/sec} = 0.02 \text{ ms}$
 - $T_{\text{access}} = 9 \text{ ms} + 4 \text{ ms} + 0.02 \text{ ms}$
- 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 than DRAM.

38

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”.

39

Disk Access Time Example

- Given:
 - Rotational rate = 7,200 RPM
 - Average seek time = 9 ms.
 - Avg # sectors/track = 400.
- Derived:
 - $T_{\text{avg rotation}} = 1/2 \times (60 \text{ secs}/7200 \text{ RPM}) \times 1000 \text{ ms/sec} = 4 \text{ ms.}$
 - $T_{\text{avg transfer}} = 60/7200 \text{ RPM} \times 1/400 \text{ secs/track} \times 1000 \text{ ms/sec} = 0.02 \text{ ms}$
 - $T_{\text{access}} = 9 \text{ ms} + 4 \text{ ms} + 0.02 \text{ ms}$
- 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 than DRAM.

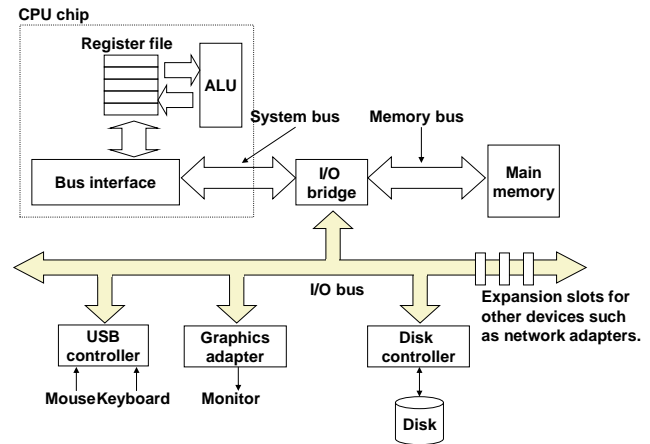
40

Quiz 8/Breakout

- Consider a magnetic disk with 1000 cylinders rotating at a speed of 7200 revolutions per minute. The sector size on the disk is 1024 bytes. Requests for sectors arriving from the file system can be located anywhere on the disk with equal probability and the locations of the requested sectors are independent of each other. Under these assumptions, the average seek time is 7 milliseconds. The transfer time per sector is 0.01 milliseconds (i.e., the raw data rate is 100 Megabytes per second). What is the effective data rate?

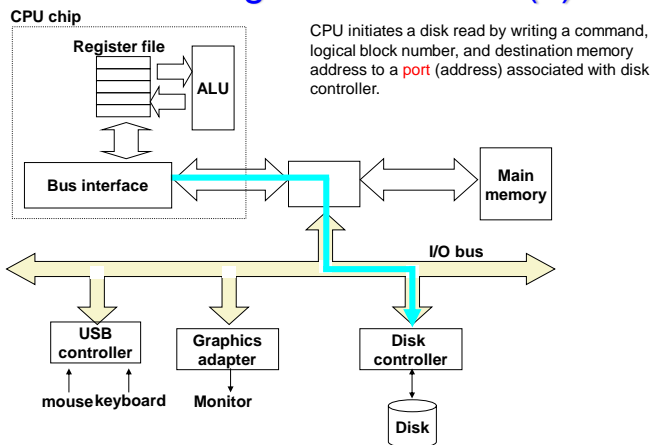
41

I/O Bus



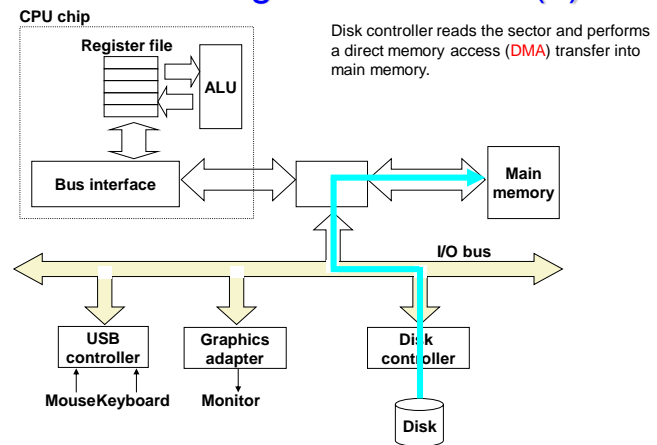
42

Reading a Disk Sector (1)



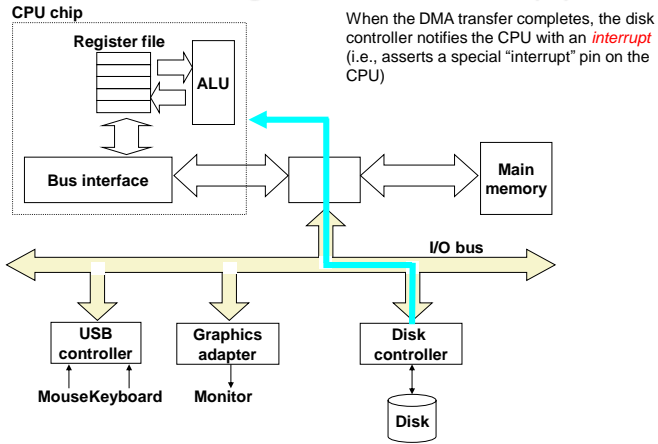
43

Reading a Disk Sector (2)



44

Reading a Disk Sector (3)



45

Solid State Drives

- No mechanical component (moving parts)
- Lower energy requirements
- Speed
 - Reads and writes in the order of 10s of microseconds (reading faster than writing)
 - Erase on the order of a millisecond
- Finite number of erase and write cycles, requiring what is called "wear leveling"

11/18/2020

CSC 2/456


46

Flash Memory (Based on Charge)

- Based on floating-gate transistor

From Computer Desktop Encyclopedia
© 2005 The Computer Language Co., Inc.

EEPROM and Flash Transistor

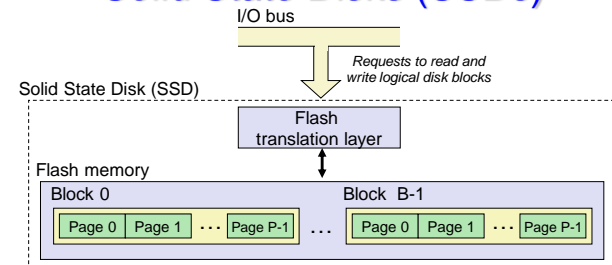


11/18/2021

47

47

Solid State Disks (SSDs)



- Pages: 512KB to 4KB, Blocks: 32 to 128 pages
- Data read/written in units of pages.
- Page can be written only after its block has been erased
- A block wears out after about 100,000 repeated writes.

48

SSD Performance Characteristics

Sequential read thput	550 MB/s	Sequential write thput	470 MB/s
Random read thput	365 MB/s	Random write thput	303 MB/s
Avg seq read time	50 us	Avg seq write time	60 us

- Sequential access faster than random access
 - Common theme in the memory hierarchy
- Random writes are somewhat slower
 - Erasing a block takes a long time (~1 ms)
 - Modifying a block page requires all other pages to be copied to new block
 - In earlier SSDs, the read/write gap was much larger.

Source: Intel SSD 730 product specification.

49


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
 - E.g. Intel SSD 730 guarantees 128 petabyte (128 x 10¹⁵ bytes) of writes before they wear out
- Applications
 - MP3 players, smart phones, laptops
 - Also becoming prevalent in desktops and servers

50

The CPU-Memory Gap

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



51

Locality to the Rescue!


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

52

Coming Up: Virtual Memory

53

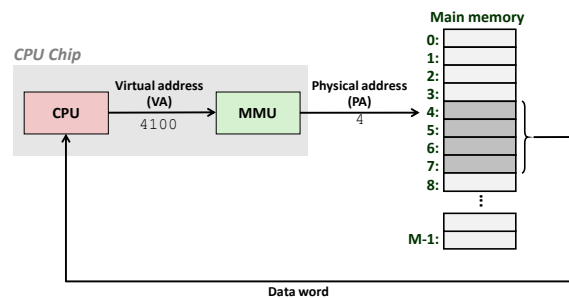
A System Using Physical Addressing



- Used in “simple” systems like embedded microcontrollers in devices like cars, elevators, and digital picture frames

55

A System Using Virtual Addressing



- Used in all modern servers, laptops, and smart phones
- One of the great ideas in computer science

56