

CACHE MEMORY

ABEBE A. APRIL, 2019

REMEMBER THIS GUY



Von Neumann Architecture:

- CPU ;
- Memory;
- I/O Module



INTRODUCTION

- **Today's focus: Memory Module** of von Neumann's architecture

INTRODUCTION

- Although simple in concept computer memory exhibits wide range of:
 - Type: internal/external, permanent/temporary
 - Technology: optical/magnetic/semiconductor
 - Organization: words/bytes/blocks
 - Performance: fast/slow
 - and cost: Expensive/cheap
- No single technology is optimal in satisfying all of these

INTRODUCTION

Typically:

- Higher performance → higher cost;
- Lower performance → lower cost

INTRODUCTION

A computer has a hierarchy of memory subsystems:

- some internal to the system
 - i.e. directly accessible by the processor;
- some external
 - accessible via an I/O module;

CLASSIFICATION OF MEMORY SYSTEMS ACCORDING TO THEIR KEY CHARACTERISTICS:

Location

Internal (e.g. processor registers, main memory, cache)
External (e.g. optical disks, magnetic disks, tapes)

Capacity

Number of words
Number of bytes

Unit of Transfer

Word
Block

Access Method

Sequential
Direct
Random
Associative

Performance

Access time
Cycle time
Transfer rate

Physical Type

Semiconductor
Magnetic
Optical
Magneto-optical

Physical Characteristics

Volatile/nonvolatile
Erasable/nonerasable

CHARACTERISTICS OF MEMORY SYSTEMS

Location: either **internal** or **external** to the processor.

- Forms of internal memory:
 - **registers;**
 - **cache;**
 - **RAM**
- Forms of external memory:
 - **disk;**
 - **magnetic tape**
 - devices that are accessible to the processor via **I/O** controllers.

CHARACTERISTICS OF MEMORY SYSTEMS

- **Capacity: amount** of information the memory is capable of **holding**.
 - Typically expressed in terms of **bytes** (1 byte = 8 bits) or **words**;
 - A **word** represents each **addressable** block of the memory
 - common word lengths are **8, 16, and 32 bits**;
 - External memory capacity is typically expressed in terms of **bytes**;

CHARACTERISTICS OF MEMORY SYSTEMS

- **Unity of transfer:**
 - For **internal** memory, the unit of transfer is equal to the number of **electrical lines** into and out of the memory module.
 - This may be equal to the **word length**, but is often larger, such as 64,128,or 256 bits.
 - To clarify this point, consider three related concepts for internal memory:

CHARACTERISTICS OF MEMORY SYSTEMS

- **Word:**
 - The “**natural**” unit of **organization** of memory.
 - The size of the word is typically equal to the number of bits used to represent an **integer** and to the **instruction length**.
 - there are many exceptions.
- **Addressable units:**
 - the addressable unit is the **word**.
 - However, many systems allow addressing at the **byte** level.
 - In any case, the relationship between the length in bits A of an address and the number N of addressable units is N .
- **Unit of transfer:**
 - this is the **number of bits** read out of or written into memory at a time.
 - The unit of transfer need not equal a **word** or an **addressable unit**.
 - For external memory,
 - data are often transferred in much larger units than a word, and these are referred to as **blocks**.

CHARACTERISTICS OF MEMORY SYSTEMS

- **Access Method:** How are the units of memory accessed?
 - **Sequential Method:** Memory is organized into units of data, called records.
 - Access must be made in a specific **linear** sequence;
 - Stored addressing information is used to assist in the retrieval process.
 - A shared read-write head is used;
 - The head must be moved from its one location to the another;
 - Passing and rejecting each intermediate record;
 - **Highly variable** times.



CHARACTERISTICS OF MEMORY SYSTEMS

- **Direct Access Memory:**


- Involves a shared read-write mechanism;
- Individual records have a unique address;
- Requires accessing general record vicinity plus sequential searching, counting, or waiting to reach the final location;
- Access time is also **variable**;



Magnetic Disk

CHARACTERISTICS OF MEMORY SYSTEMS

- **Random Access:** Each addressable location in memory has a unique, **physically wired-in** addressing mechanism.
 - Constant time;
 - independent of the sequence of prior accesses;
 - Any location can be selected at random and directly accessed;
 - Main memory and some cache systems are random access



CHARACTERISTICS OF MEMORY SYSTEMS

- **Associative:** a memory whose storage locations are identified by their contents, or by a part of their contents, rather than by their names or positions
 - Word is retrieved based on a portion of its contents rather than its address;
 - Retrieval time is constant independent of location or prior access patterns

CHARACTERISTICS OF MEMORY SYSTEMS

- **Performance:**
 - **Access time (latency):**
 - For RAM: time to perform a read or write operation;
 - For Non-RAM: time to position the read-write head at desired location;
 - **Memory cycle time:** Primarily applied to RAM:
 - Access time + additional time required before a second access;
 - Required for electrical signals to be terminated/regenerated;
 - Concerns the system bus.

CHARACTERISTICS OF MEMORY SYSTEMS

- **Transfer time:** Rate at which data can be transferred in / out of memory;
 - For RAM:
 - For Non-RAM: $T_n = T_A + \frac{n}{R}$, where:
 - T_n : Average time to read or write n bits;
 - T_A : Average access time;
 - n : Number of bits
 - R : Transfer rate, in bits per second (bps.)

CHARACTERISTICS OF MEMORY SYSTEMS

- **Physical characteristics:**

- **Volatile:** information decays naturally or is lost when powered off;
- **Nonvolatile:** information remains without deterioration until changed:
 - no electrical power is needed to retain information.;
 - E.g.: Magnetic-surface memories are nonvolatile;
- Semiconductor memory (memory on integrated circuits) may be either volatile or nonvolatile.

MEMORY HIERARCHY

Design constraints on memory can be summed up by three questions:

- **How much?**
 - If memory exists, applications will likely be developed to use it.
- **How fast?**
 - Best performance achieved when memory keeps up with the processor;
 - I.e. as the processor execute instructions, memory should minimize pausing / waiting for instructions or operands.
- **How expensive?**
 - Cost of memory must be reasonable in relationship to other components

MEMORY HIERARCHY

Memory tradeoffs are a sad part of reality

- Faster access time, greater cost per bit;
- Greater capacity:
 - Smaller cost per bit;
 - Slower access time;

MEMORY HIERARCHY

These tradeoff imply a **dilemma**:


- Large capacity memories are desired:
 - low cost and because the capacity is needed;
- However, to meet performance requirements, the designer needs:
 - to use expensive, relatively lower-capacity memories with short access times

How can we solve this issue? Or at least mitigate the problem? Any ideas?

MEMORY HIERARCHY

The way out of this **dilemma**:

- Don't rely on a **single memory**;
- Instead employ a **memory hierarchy**;
- **Supplement**:
 - smaller, more expensive, faster memories with...
 - ...larger, cheaper, slower memories;




MEMORY HIERARCHY

As one goes down the hierarchy:

- Decreasing cost per bit;
- Increasing capacity;
- Increasing access time;
- Decreasing frequency of access memory by processor

Decreasing frequency of memory access by processor.

But why is this key to success?



MEMORY HIERARCHY

But why is this key to success?

- As we go down the hierarchy we gain in size but lose in speed;
- Therefore: not efficient for the processor to access these memories;
- Requires having specific strategies to minimize such accesses;

How can we develop strategies to minimize these accesses?

MEMORY HIERARCHY

How can we develop strategies to minimize these accesses?

Space and Time locality of reference principle:

- **Space:**

- if we access a memory location, close by addresses will very likely be accessed;

- **Time:**

- if we access a memory location, we will very likely access it again;
But why does this happen?

MEMORY HIERARCHY

Suppose that the processor has access to two levels of memory:

- **Level 1 - L1:**

- contains 1000 words and has an access time of $0.01\mu\text{s}$;

- **Level 2 - L2:**

- contains 100,000 words and has an access time of $0.1\mu\text{s}$.

- Assume that:

- if word $\in$ L1, then the processor accesses it directly;
 - If word $\in$ L2, then word is transferred to L1 and then accessed by the processor.

MEMORY HIERARCHY

Strategy to minimize accesses should be:

- Organize data across the hierarchy such that
 - % of accesses to lower levels is substantially less than that of upper levels
- I.e. L2 memory contains all program instructions and data:
 - Data that is currently being used should be in L1;
 - Eventually:
 - Data $\in$ L1 will be swapped to L2 to make room for new data;
 - On average, most references will be to data contained in L1.

MEMORY HIERARCHY


This principle can be applied across more than two levels of memory:

- **Processor registers:**
 - Fastest, smallest, and most expensive type of memory
- Followed immediately by the **cache**:
 - Stages data movement between registers and main memory;
 - Improves performance;
 - Is not usually visible to the processor;
 - Is not usually visible to the programmer.
- Followed by **main memory**:
 - Principal internal memory system of the computer;
 - Each location has a unique address

CACHE MEMORY PRINCIPLES

Cache memory is designed to combine :

- Memory access time of expensive, high-speed memory combined with...
- ...the large memory size of less expensive, lower-speed memory.



CACHE MEMORY

Typically:

- Higher performance → higher cost;
- Lower performance → lower cost

INTERNAL MEMORY

ABEBE A. JUNE, 2019

INTRODUCTION

- Previous chapter discussed memory. This chapter presents:
 - semiconductor main memory subsystems.
 - ROM;
 - DRAM;
 - SRAM memories
 - But why do we need these types of memories? Any ideas?

INTRODUCTION

- Remember these guys?



- Serial memory takes different lengths of time to access information:
 - depends on where the desired location is relative to the current position;

SEMICONDUCTOR MAIN MEMORY

- Memory is a collection of cells with the following properties
 - They exhibit two stable states, used to represent binary 1 and 0;
 - They are capable of being written into (at least once), to set the state;
 - They are capable of being read to sense the state;
 - Access time is the same regardless of the location;

SEMICONDUCTOR MAIN MEMORY

- Cell has three terminals capable of carrying an electrical signal:
 - **select**: selects a memory cell for a read / write operation;
 - **control**: indicates whether a read / write operation is being performed;
 - **read/write**:
 - For writing, terminal sets the state of the cell to 1 or 0
 - For reading, terminal is used for output of the cell's




Figure: Memory cell write operation (Source: Stallings, 2015)



Figure: Memory cell read operation (Source: Stallings, 2015)

SEMICONDUCTOR MAIN MEMORY

- Consider, a memory with a capacity of 1K words of 16 bits each:
 - Capacity:
 - $1K \times 16 = 16Kbits = 2Kbytes$;
 - Each word has an address:
 - 0 to 1023
 - When a word is read or written:
 - memory acts on all 16 bits;
- Based on these concepts how do you think a write operation works? Any

Memory Address		Memory Contents
Binary	Decimal	
00000000	0	10110101 01011100
00000001	1	10101011 10001001
00000010	2	00001101 01000110
.	.	.
.	.	.
.	.	.
.	.	.
.	.	.
111111101	1021	10011101 00010101
111111110	1022	00001101 00011110
111111111	1023	11011110 00100100

SEMICONDUCTOR MAIN MEMORY

- Write operation steps:
 - Apply the binary address of the desired word to the address lines.
 - Apply the data bits that must be stored in memory to the data input lines.
 - Activate the Write control line.
- Memory unit will then transfer the bits to that address.

SEMICONDUCTOR MAIN MEMORY

- Read operation steps:
 - Apply the binary address of the desired word to the address lines.
 - Activate the Read control line.
 - Memory unit will then transfer the bits to the data output lines.

TIMING WAVEFORMS

- Memory unit operation is controlled by an external device (1/2):
 - CPU is synchronized by its own clock pulses;
 - Control signals are employed for memory read / write.
 - Access time - time required for specifying an address and obtaining the word;
 - Write time - time required for specifying an address and storing the word;
- CPU must provide the control signals

TIMING WAVEFORMS

- Memory unit operation is controlled by an external device (2/2):
 - CPU provides the control signals synchronized with its clock:
 - This implies that:
 - read / write operations will take a certain number of clock periods;
 - What does this mean? Lets see with an example =)

EXAMPLE (1/7)

- Assume:
 - CPU with a clock frequency of 50 Mhz;
 - Memory access time: 65ns;
 - Memory write time: 75ns
- How many clock pulses do we need for read / write operations?

EXAMPLE (2/7)

- CPU with a clock frequency of 50 Mhz:
- • How long does one clock pulse take? Any Ideas?
- $f = 1/p \Leftrightarrow$
- $p = 1/f \Leftrightarrow$
- $p = 1/50\text{Mhz} \Leftrightarrow$
- $p = 1 / 50 \times \text{hz} \Leftrightarrow$
- $p = 2 \times \text{s} \Leftrightarrow 20\text{ns}$

EXAMPLE (3/7)

- CPU clock pulse: 20ns;
- Assume:
 - Memory access time: 65ns;
 - Memory write time: 75ns
- How many clock pulses do we need for read / write operations?
 - 4 clock pulses for read / write operations;

EXAMPLE (4/7)

For the **write operation**:




Figure: Write cycle (Source: (Mano and Kime, 2007))

EXAMPLE (5/7)

- For the write operation:
 - 1 CPU must provide the address (T1 pulse);
 - 2 Memory enable is set (T1 pulse);
 - 3 Data is supplied (T2 pulse);
 - 4 Read/Write signal set to 0 for write operation (T2 pulse):
 - • CPU waits for T2 to let the address signals stabilize;
 - • Otherwise: wrong address may be used!
 - • Signal must stay activated long enough for the operation to finish;
 - 5 When T4 completes, write operation has ended with 5ns to spare

EXAMPLE (6/7)

For the **read operation**:




Figure: Read cycle (Source: (Mano and Kime, 2007))

EXAMPLE (7/7)

- For the read operation:
 - 1 CPU must provide the address (T1 pulse);
 - 2 Memory enable is set (T1 pulse);
 - 3 Read/Write signal set to 1 for read operation (T2 pulse):
 - Signal must stay activated long enough for the operation to finish;
 - Selected word is placed onto the data output lines;
 - 4 When T4 completes, write operation has ended with 15ns to spare

MAJOR TYPES OF SEMICONDUCTOR MEMORY

- For the read operation:
 - 1 CPU must provide the address (T1 pulse);
 - 2 Memory enable is set (T1 pulse);
 - 3 Read/Write signal set to 1 for read operation (T2 pulse):
 - Signal must stay activated long enough for the operation to finish;
 - Selected word is placed onto the data output lines;
 - 4 When T4 completes, write operation has ended with 15ns to spare

MAJOR TYPES OF SEMICONDUCTOR MEMORY

LET'S HAVE A LOOK AT THE DIFFERENT TYPES OF MEMORY TECHNOLOGIES:

Memory Type	Category	Erasure	Write Mechanism	Volatility
Random-access memory (RAM)	Read-write memory	Electrically, byte-level	Electrically	Volatile
Read-only memory (ROM)	Read-only memory	Not possible	Masks	Nonvolatile
Programmable ROM (PROM)				
Erasable PROM (EPROM)	Read-mostly memory	UV light, chip-level	Electrically	
Electrically Erasable PROM (EEPROM)		Electrically, byte-level		
Flash memory		Electrically, block-level		

Figure: Semiconductor memory types (Source: (Stallings, 2015))

RANDOM-ACCESS MEMORY

- Randomly access any possible address;
- Possible both to read and write data;
- Volatile: requires a power supply;
- For a chip with m words with n bits per word:
 - Consists of an array with $m \times n$ binary storage cells
- E.g.: DRAM and SRAM.
 - Lets have a look at these

DYNAMIC RAM (DRAM)

- Made with cells that store data as charge on capacitors:
- Capacitor charge presence or absence is interpreted as a binary 1 or 0;
- Capacitors have a natural tendency to discharge;
- Requires periodic charge refreshing to maintain data storage

DRAM CELL COMPONENTS:

- If voltage goes to the address line:
 - Transistor closes:
 - Current flows to capacitor;
- If no voltage goes to the address line:
 - Transistor opens:
 - No current flows to capacitor
- *Based on the DRAM cell components:*
- *How do you think a write operation works?*




Figure: DRAM Cell (Source: (Stallings, 2015))

DYNAMIC RAM (DRAM)

- Write operation:
 - 1 Voltage signal is applied to the bit line
 - • Low voltage = 0;
 - • High voltage = 1;
 - 2 A signal is then applied to the address line:
 - • transistor closes...
 - • ...charge goes to the capacitor
- *How do you think a read operation works?*

DYNAMIC RAM (DRAM)

- Read operation:
 - 1 Address line is activated
 - 2 Transistor turns on
 - 3 Capacitor charge goes to bit line...;
 - ...Low voltage = 0
 - ...High voltage = 1
 - 4 Cell readout discharges the capacitor:
 - state must be restored;
- *Do you have any idea of other type of technology that can be employed?*

STATIC RAM (SRAM)

- Binary values are stored using SR flip-flop configurations:
 - Remember flip-flops?
 - Same logic elements used in the processor registers;
 - will hold its data as long as power is supplied to it.

RAM

- How to choose between DRAM and SRAM? Any ideas?
- • What are the pros / cons of DRAM?
- • What are the pros / cons of SRAM?

DRAM VS SRAM

- Both are volatile:
 - power must be continuously supplied to preserve the bit values;
- DRAMs:
 - Periodically refresh capacitor's charge;
- SRAM:
 - No need to periodically refresh;

DRAM VS SRAM

- DRAM cell is simpler and smaller than a SRAM cell:
 - • (1 transistor, 1 capacitor) vs. SR flip-flop
 - • DRAM is denser (smaller cells = more cells per unit area) ;
 - • Due to its simplicity:
 - • DRAM is cheaper than corresponding SRAM;

DRAM VS SRAM

- DRAM requires the supporting refresh circuitry (readout):
 - This is a fixed cost, does not increase with size
 - This cost is more than compensated by the smaller cost of DRAM cells;
 - Thus, DRAMs tend to be favored for large memory requirements;

DRAM VS SRAM

- SRAMs are faster and more expensive than DRAMs
 - • No need to refresh the circuitry after a readout;
- • Because of these characteristics:
 - • DRAM: used in main memory;
 - • SRAM: used in cache;

DRAM VS SRAM

Memory technology	Typical access time	\$ per GiB in 2012
SRAM semiconductor memory	0.5–2.5 ns	\$500–\$1000
DRAM semiconductor memory	50–70 ns	\$10–\$20
Flash semiconductor memory	5,000–50,000 ns	\$0.75–\$1.00
Magnetic disk	5,000,000–20,000,000 ns	\$0.05–\$0.10

Figure: Access time and price per bit (Source: (Patterson and Hennessy, 2013))

READ-ONLY MEMORY (ROM)

- Contains a permanent pattern of data that cannot be changed;
 - • Therefore:
 - • Possible to read a ROM;
 - • Not possible to write data more than once;
 - • Advantage:
 - • Data / Program is permanently in main memory;
 - • No need to load from a secondary storage device;
 - • Original BIOS were stored in ROM chips.
 - • Nonvolatile: no power source is required;

READ-ONLY MEMORY (ROM)

- Not possible to write data more than once:
 - • Data is wired into the chip as part of the fabrication process
 - • Data insertion is expensive;
 - • No room for error:
 - • If one bit is wrong, the whole batch of ROMs must be thrown out.

READ-ONLY MEMORY (ROM)

- Different types of ROM (1/4):
 - • Programmable ROM (PROM):
 - • nonvolatile and may be written into only once;
 - • Writing may be performed at a later time than the original chip fabrication.
 - • Special equipment is required for the writing process;
 - • More expensive than ROM;

READ-ONLY MEMORY (ROM)

- Different types of ROM (2/4):
 - • Erasable programmable read-only memory (EPROM):
 - • Can be altered multiple times;
 - • Entire memory contents need to be erased;
 - • Special equipment is required for the erasure process:
 - • Ultra-violet radiation;
 - • Slow erasure procedure (>20 minutes);
 - • More expensive than PROM;

READ-ONLY MEMORY (ROM)

- Different types of ROM (3/4):
 - • Electrically erasable programmable read-only memory (EEPROM):
 - • Can be altered multiple times;
 - • Special equipment is required for the erasure process:
 - • No need to erase prior contents...
 - • ...only the byte or bytes addressed are updated.
 - • Ultra-violet radiation;
 - • Faster erasure procedure than EPROM;
 - • More expensive than EPROM;

READ-ONLY MEMORY (ROM)

- Different types of ROM (4/4):
 - • Flash memory:
 - • named because of the speed with which it can be reprogrammed;
 - • between EPROM and EEPROM in both cost and functionality:
 - • Uses an electrical erasing technology;
 - • Possible to erase blocks of memory;
 - • Does not provide byte-level erasure;
 - • Technology used in pendrives;

ERROR CORRECTION

- Semiconductor memory systems are subject to errors (1/2):
- Hard Failures: • Permanent physical defect;
 - • Damaged memory cell(s) cannot reliably store data:
 - • become stuck at 0 or 1 or...
 - • ...switch erratically between 0 and 1
 - • Caused by:
 - • Harsh environmental abuse;
 - • Manufacturing defects;
 - • Wear and tear;

ERROR CORRECTION

- Semiconductor memory systems are subject to errors (2/2):
 - Soft Failures:
 - Random, nondestructive event
 - Alters the contents of one or more memory cells
 - Without damaging the memory.
 - Caused by:
 - Power supply problems;
 - Alpha particles (radioactive decay).
 - Hard and soft errors are clearly undesirable...So what can we do to address them?

ERROR CORRECTION




Figure: Error correcting code function (Source: (Stallings, 2015))

ERROR CORRECTION

- From the previous picture (1/3):
 - 1 M-bit data is to be written into memory;
 - 2 Computation $f(M)$ is performed on the data (K check bits);
 - 3 Both the data and the code ($M + K$ bits) are stored in memory;


ERROR CORRECTION

- From the previous picture (2/3):
 - 4 Eventually, the stored word will be read out;
 - 5 Computation $f(M)$ is again performed on the data (K check bits);
 - 6 A comparison is made between both check bits;

ERROR CORRECTION

- From the previous picture (3/3):
 - 7 The comparison yields one of three results:
 - • No errors are detected:
 - • Fetched data bits are sent out;
 - • An error is detected, and it is possible to correct the error:
 - • Data bits plus error correction bits are fed into a corrector
 - • Producing a corrected set of M bits to be sent out.
 - • An error is detected, but it is not possible to correct it:
 - This condition is reported.

-
- Error-correcting codes



CHAPTER 6

EXTERNAL MEMORY

TYPES OF EXTERNAL MEMORY

- Magnetic Disk
 - RAID
 - Removable
- Optical
 - CD-ROM
 - CD-Recordable (CD-R)
 - CD-R/W
 - DVD
- Magnetic Tape


MAGNETIC DISK

- Disk substrate coated with magnetizable material (iron oxide...rust)
- Substrate used to be aluminium
- Now glass
 - Improved surface uniformity
 - Increases reliability
 - Reduction in surface defects
 - Reduced read/write errors
 - Lower flight heights (See later)
 - Better stiffness
 - Better shock/damage resistance

READ AND WRITE MECHANISMS

- Recording & retrieval via conductive coil called a **head**
- May be **single** read/write head or **separate** ones
- During read/write, head is **stationary**, **platter rotates**
- Write
 - Current through coil produces **magnetic** field
 - Pulses sent to **head**
 - Magnetic pattern **recorded** on surface below
- Read (traditional)
 - Magnetic field moving relative to **coil** produces **current**
 - Coil is the same for **read** and **write**
- Read (contemporary)
 - **Separate** read head, close to write head
 - Partially shielded magneto resistive (MR) sensor
 - Electrical resistance depends on direction of magnetic field
 - High frequency operation
 - Higher storage density and speed


INDUCTIVE WRITE MR READ



DATA ORGANIZATION AND FORMATTING

- Concentric rings or tracks
 - Gaps between tracks
 - Reduce gap to increase capacity
 - Same number of bits per track (variable packing density)
 - Constant angular velocity
- Tracks divided into sectors
- Minimum block size is one sector
- May have more than one sector per block


DISK DATA LAYOUT




DISK VELOCITY

- Bit near centre of rotating disk passes fixed point slower than bit on outside of disk
- Increase spacing between bits in different tracks
- Rotate disk at constant angular velocity (CAV)
 - Gives pie shaped sectors and concentric tracks
 - Individual tracks and sectors addressable
 - Move head to given track and wait for given sector
 - Waste of space on outer tracks
 - Lower data density
- Can use zones to increase capacity
 - Each zone has fixed bits per track
 - More complex circuitry

DISK LAYOUT METHODS DIAGRAM



(a) Constant angular velocity




(b) Multiple zoned recording

FINDING SECTORS

- Must be able to identify start of track and sector
- Format disk
 - Additional information not available to user
 - Marks tracks and sectors

WINCHESTER DISK FORMAT SEAGATE ST506



CHARACTERISTICS

- Fixed (rare) or movable head
- Removable or fixed
- Single or double (usually) sided
- Single or multiple platter
- Head mechanism
 - Contact (Floppy)
 - Fixed gap
 - Flying (Winchester)

FIXED/MOVABLE HEAD DISK

- Fixed head
 - One read write head per track
 - Heads mounted on fixed ridged arm
- Movable head
 - One read write head per side
 - Mounted on a movable arm


REMOVABLE OR NOT

- Removable disk
 - Can be removed from drive and replaced with another disk
 - Provides unlimited storage capacity
 - Easy data transfer between systems
- Nonremovable disk
 - Permanently mounted in the drive


MULTIPLE PLATTER

- One head per side
- Heads are joined and aligned
- Aligned tracks on each platter form cylinders
- Data is striped by cylinder
 - reduces head movement
 - Increases speed (transfer rate)

MULTIPLE PLATTERS



TRACKS AND CYLINDERS



FLOPPY DISK

- 8", 5.25", 3.5"
- Small capacity
 - Up to 1.44Mbyte (2.88M never popular)
- Slow
- Universal
- Cheap
- Obsolete?

WINCHESTER HARD DISK (1)

- Developed by IBM in Winchester (USA)
- Sealed unit
- One or more platters (disks)
- Heads fly on boundary layer of air as disk spins
- Very small head to disk gap
- Getting more robust

WINCHESTER HARD DISK (2)

- Universal
- Cheap
- Fastest external storage
- Getting larger all the time
 - 250 Gigabyte now easily available

SPEED

- Seek time
 - Moving head to correct track
- (Rotational) latency
 - Waiting for data to rotate under head
- Access time = Seek + Latency
- Transfer rate

TIMING OF DISK I/O TRANSFER



RAID

- Redundant Array of Independent Disks
- Redundant Array of Inexpensive Disks
- 6 levels in common use
- Not a hierarchy
- Set of physical disks viewed as single logical drive by O/S
- Data distributed across physical drives
- Can use redundant capacity to store parity information

RAID 0

- No redundancy
- Data striped across all disks
- Round Robin striping
- Increase speed
 - Multiple data requests probably not on same disk
 - Disks seek in parallel
 - A set of data is likely to be striped across multiple disks


RAID 1

- Mirrored Disks
- Data is striped across disks
- 2 copies of each stripe on separate disks
- Read from either
- Write to both
- Recovery is simple
 - Swap faulty disk & re-mirror
 - No down time
- Expensive

OPTICAL STORAGE CD-ROM

- Originally for audio
- 650Mbytes giving over 70 minutes audio
- Polycarbonate coated with highly reflective coat, usually aluminium
- Data stored as pits
- Read by reflecting laser
- Constant packing density
- Constant linear velocity


CD OPERATION



CD-ROM DRIVE SPEEDS

- Audio is single speed
 - Constant linear velocity
 - 1.2 ms^{-1}
 - Track (spiral) is 5.27km long
 - Gives 4391 seconds = 73.2 minutes
- Other speeds are quoted as multiples
- e.g. 24x
- Quoted figure is maximum drive can achieve

CD-ROM FORMAT



- Mode 0=blank data field
- Mode 1=2048 byte data+error correction
- Mode 2=2336 byte data

RANDOM ACCESS ON CD-ROM

- Difficult
- Move head to rough position
- Set correct speed
- Read address
- Adjust to required location
- (Yawn!)

CD-ROM FOR & AGAINST

- Large capacity (?)
 - Easy to mass produce
 - Removable
 - Robust
-
- Expensive for small runs
 - Slow
 - Read only

OTHER OPTICAL STORAGE

- CD-Recordable (CD-R)
 - WORM
 - Now affordable
 - Compatible with CD-ROM drives
- CD-RW
 - Erasable
 - Getting cheaper
 - Mostly CD-ROM drive compatible
 - Phase change
 - Material has two different reflectivities in different phase states

DVD - WHAT'S IN A NAME?

- Digital Video Disk
 - Used to indicate a player for movies
 - Only plays video disks
- Digital Versatile Disk
 - Used to indicate a computer drive
 - Will read computer disks and play video disks
- Dogs Veritable Dinner
- Officially - nothing!!!


DVD - TECHNOLOGY

- Multi-layer
- Very high capacity (4.7G per layer)
- Full length movie on single disk
 - Using MPEG compression
- Finally standardized (honest!)
- Movies carry regional coding
- Players only play correct region films
- Can be “fixed”


DVD – WRITABLE

- Loads of trouble with standards
- First generation DVD drives may not read first generation DVD-W disks
- First generation DVD drives may not read CD-RW disks
- Wait for it to settle down before buying!

CD AND DVD



(a) CD-ROM - Capacity 682 MB



(b) DVD-ROM, double-sided, dual-layer - Capacity 17 GB

HIGH DEFINITION OPTICAL DISKS

- Designed for high definition videos
- Much higher capacity than DVD
 - Shorter wavelength laser
 - Blue-violet range
 - Smaller pits
- HD-DVD
 - 15GB single side single layer
- Blue-ray
 - Data layer closer to laser
 - Tighter focus, less distortion, smaller pits
 - 25GB on single layer
 - Available read only (BD-ROM), Recordable once (BR-R) and re-recordable (BR-RE)

MAGNETIC TAPE

- Serial access
- Slow
- Very cheap
- Backup and archive
- Linear Tape-Open (LTO) Tape Drives
 - Developed late 1990s
 - Open source alternative to proprietary tape systems