JOHN ABBOTT
CEGEP/COLLEGE

Technical Support (420-1N6-AB)
Chipsets, CPU and RAM

Fall 2025

Outline

Last week we learned about how information is represented in computer systems (binary, logical circuits).

In particular, we saw how simple switches could be used to represent AND, OR, XOR, and NOT logic in electrical devices.

Combining these simple “gates” to produce more sophisticated devices allows computer engineers to created Integrated Circuits (ICs), or, as we more commonly call them, Chips.

We now know enough to learn finish learning about some of the most important ICs in a computer introduced in Week 2:

- Chipsets
- Central Processing Unit (CPU)
- Graphical Processing Unit (GPU)
- Memory and Storage Devices

Logic Gates -> Integrated Components (Chips)

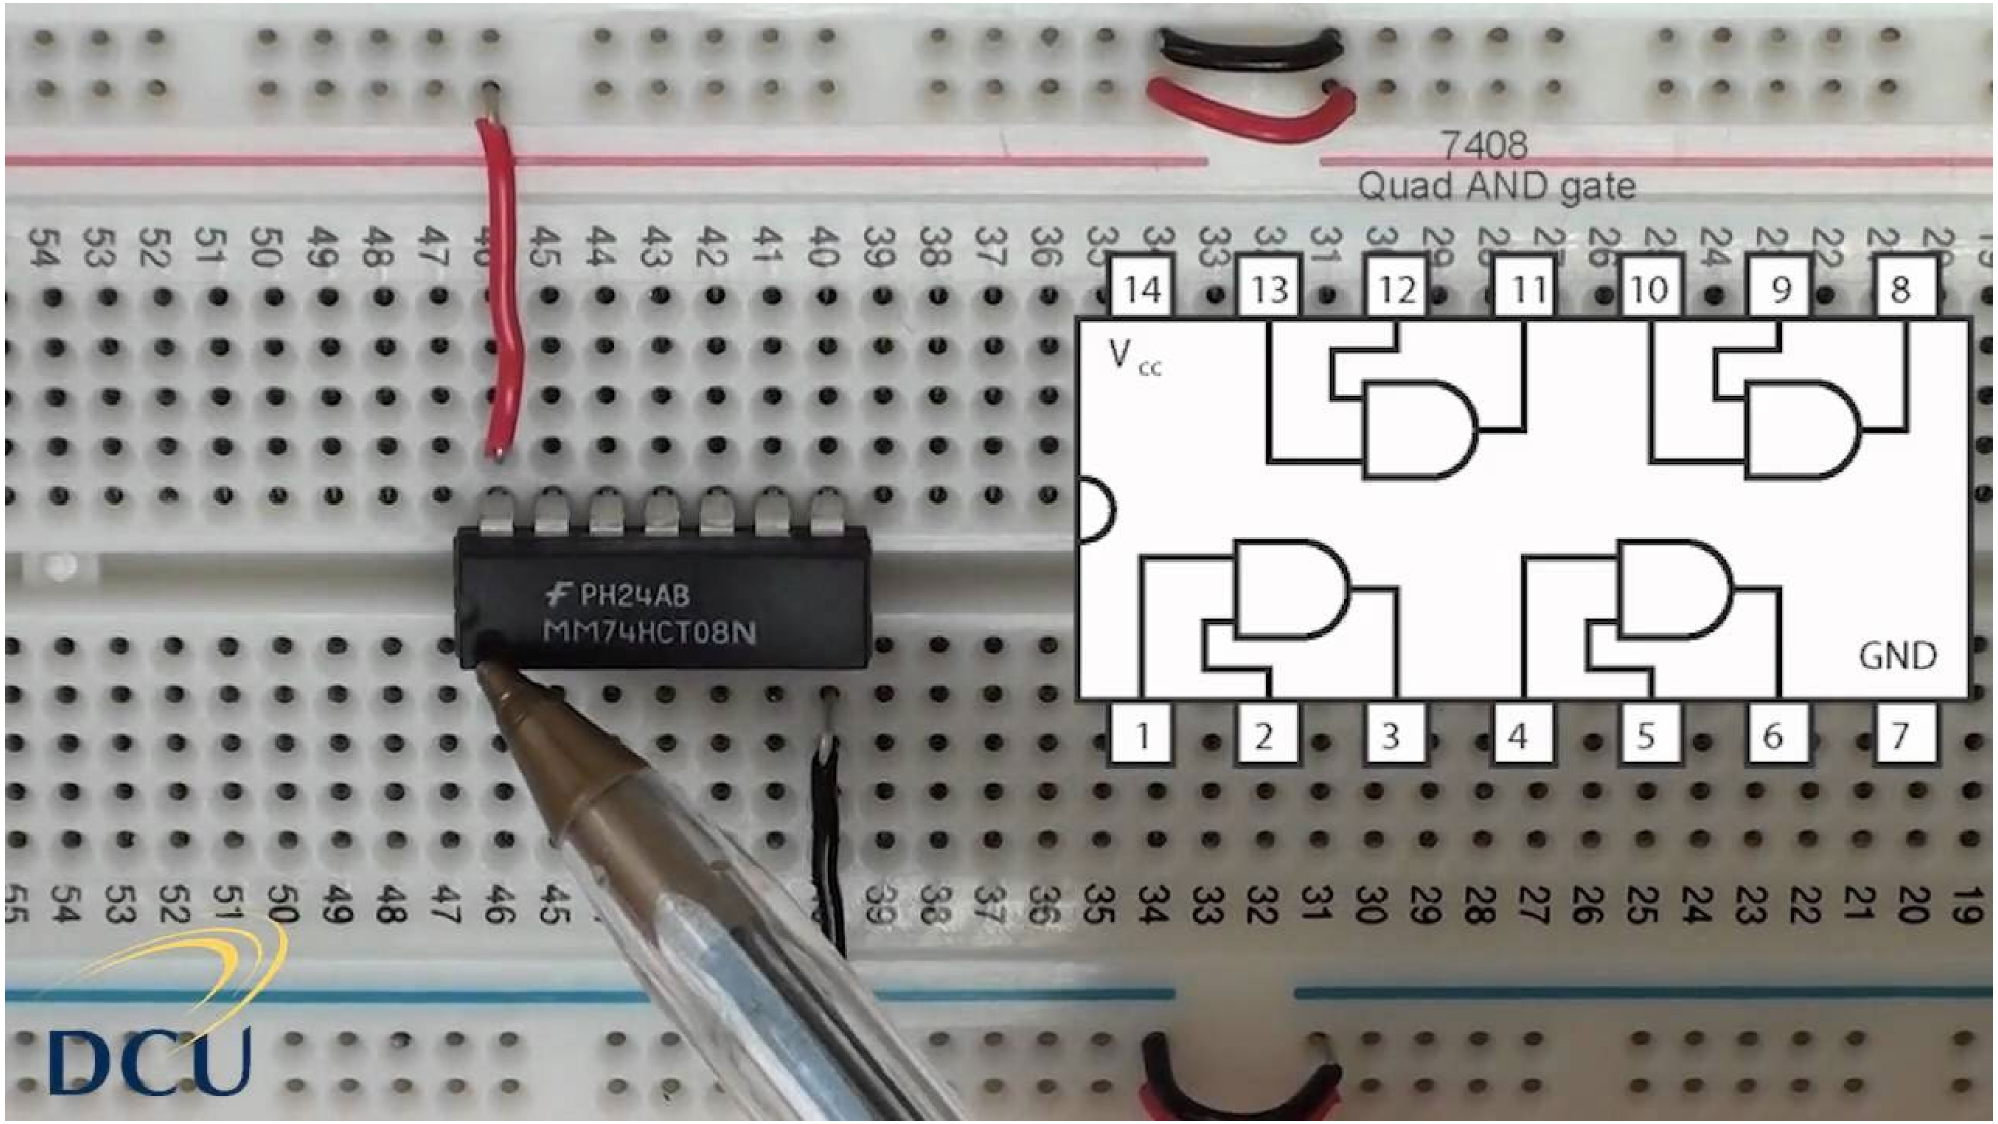

Integrated Components (Chips) -> Motherboard

Mainboard Layout

This section shows the major mainboard components.

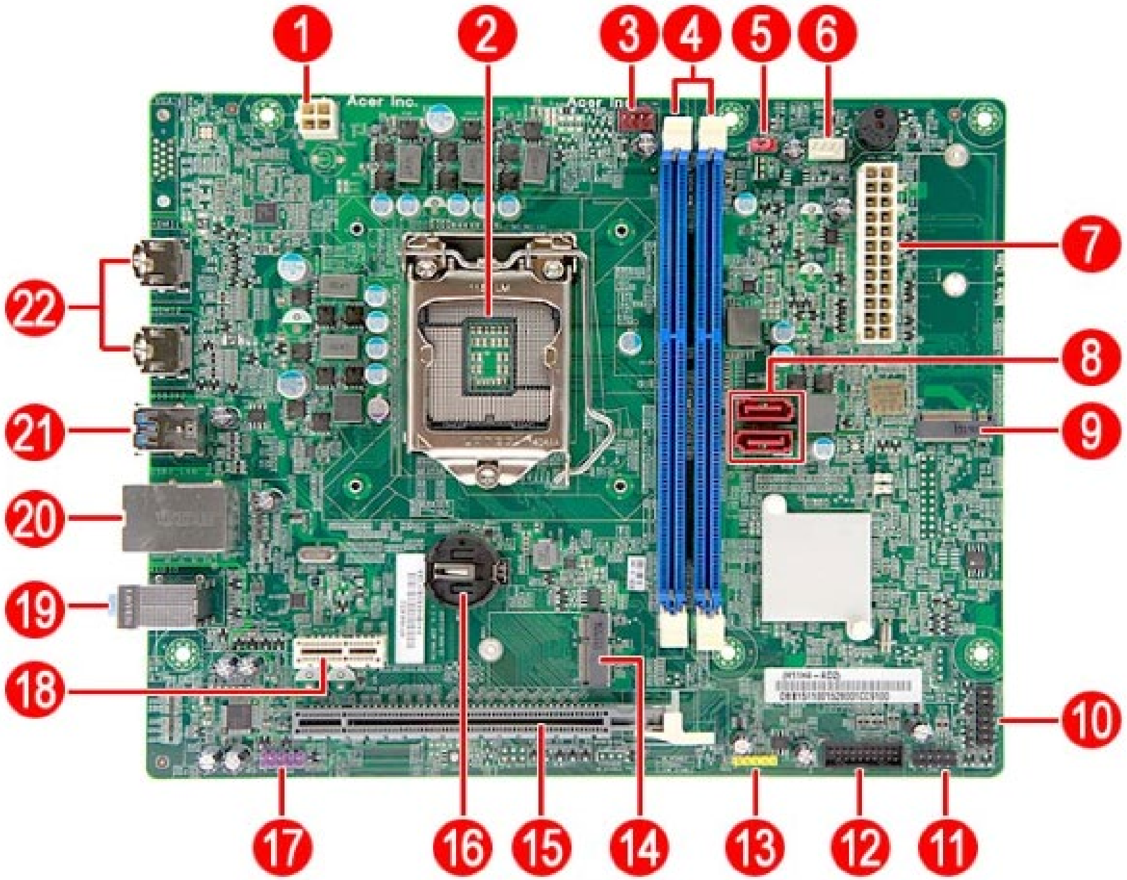

No.	Label	Description	No.	Label	Description
1	ATX_12V	4-Pin ATX power connector	12	USB3F1	Front panel USB 3.0 connector
2	CPU	CPU socket	13	F_USB2	Front panel USB 2.0 connector
3	CPU_FAN	CPU fan connector	14	M.2_1	M.2 WLAN connector
4	DIMM1~2	240-Pin DDR3L SDRAM slots	15	PCIE16X	PCIe x16 Slot
5	CLR_CMOS	Clear CMOS header with jumper	16	BAT1	Battery slot
6	SYS_FAN	System fan connector	17	F_AUDIO	Front panel audio connector
7	ATX_POWER	Standard 24-Pin ATX power connector	18	PCIE1X	PCIe x1 Slot
8	SATA1~2	SATA 6Gb/s connector	19	AUDIO	Rear panel audio ports
9	M.2_2	M.2 SSD connector	20	USB2_LAN	RJ45 + USB 2.0 ports
10	F_PANEL	Front panel connector	21	USB3X2	USB 3.0 ports
11	F_USB1	Front panel USB 3.0 connector	22	HDMI1~2	HDMI ports

Acer T3-710 Motherboard Schematic -- composed of many integrated components

More on Chipsets

The chipset

A **chipset** is a group of **microchips** on a computer's motherboard that manages **communication** between the CPU, memory, storage, and peripheral devices.

Key Functions

- Controls **data flow** between components
- Determines what **hardware** (RAM, CPU, GPU, etc.) is **supported**
- Manages **input/output** functions (USB, audio, networking)
- Affects system performance and upgrade options

The chipset –examples

- Intel **Z790** chipset – supports modern CPUs, PCIe 5.0, and DDR5 memory
- AMD **B650** chipset – designed for Ryzen processors

💡 **Fun Fact**

Earlier PCs had separate **Northbridge** and **Southbridge** chips.

Modern systems combine both into a **single chipset** for efficiency.

The chipset

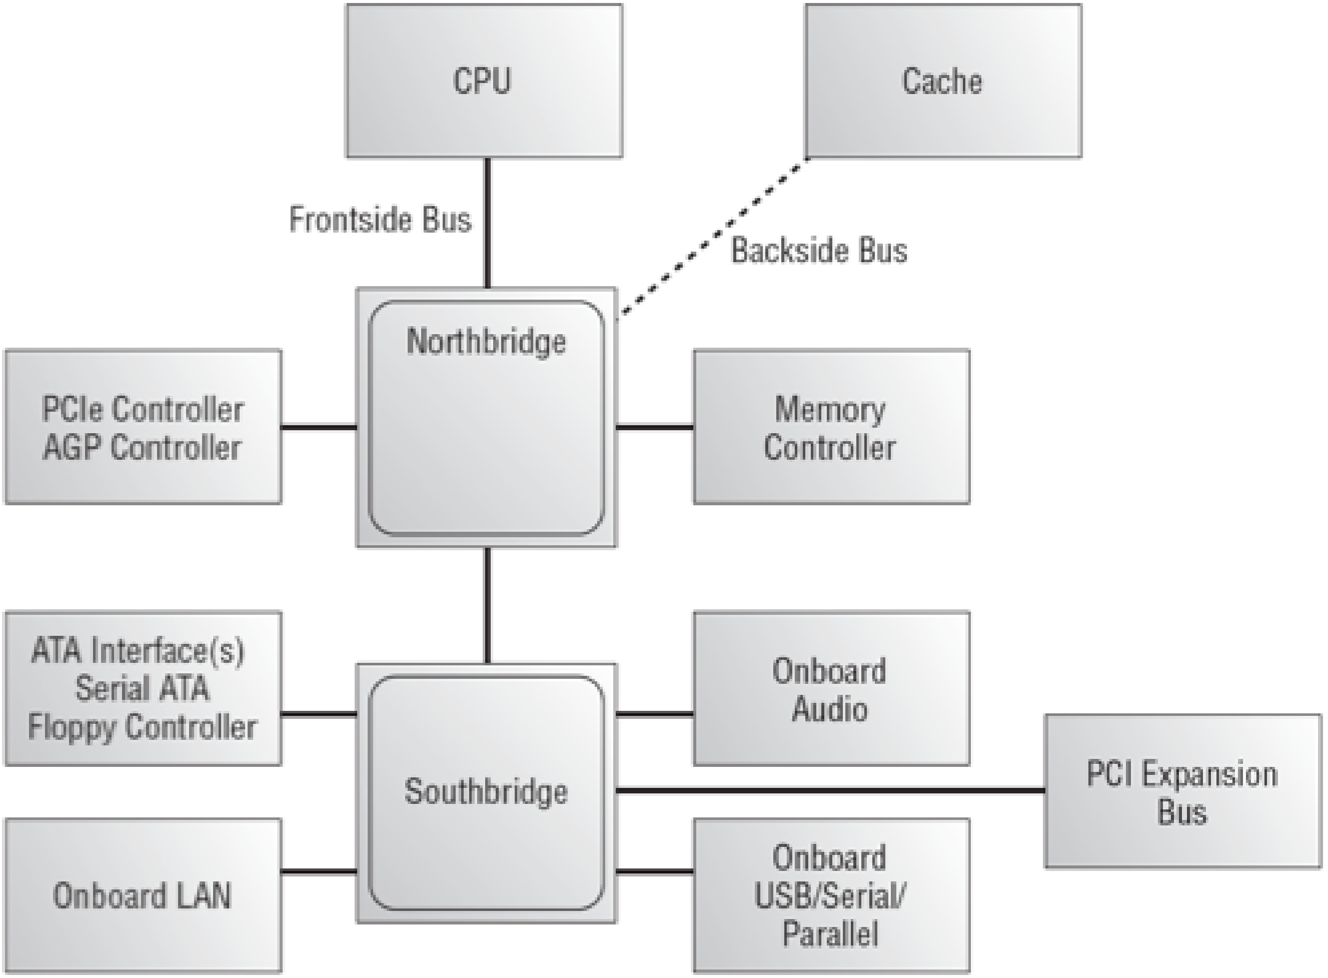
CPU
Cache
Frontside Bus
Backside Bus
Northbridge
PCIe Controller
AGP Controller
Memory
Controller
ATA Interface(s)
Serial ATA
Floppy Controller
Southbridge
Onboard
Audio
PCI Expansion
Bus
Onboard LAN
Onboard
USB/Serial/
Parallel

More on Chipsets

- Chipsets are a major component of the of the motherboard.
- A chipset controls the communication between the CPU, RAM, storage and other peripherals.
- A chipset dictates hardware compatibility across components.

How did Chipsets Evolve?

- Older motherboards were designed with a lot of **different chips scattered all over the motherboard.**

- There were a lot of different chips controlling different functions on the motherboard:
  - Chips for bus controllers, memory controllers, keyboard controllers, and so on.

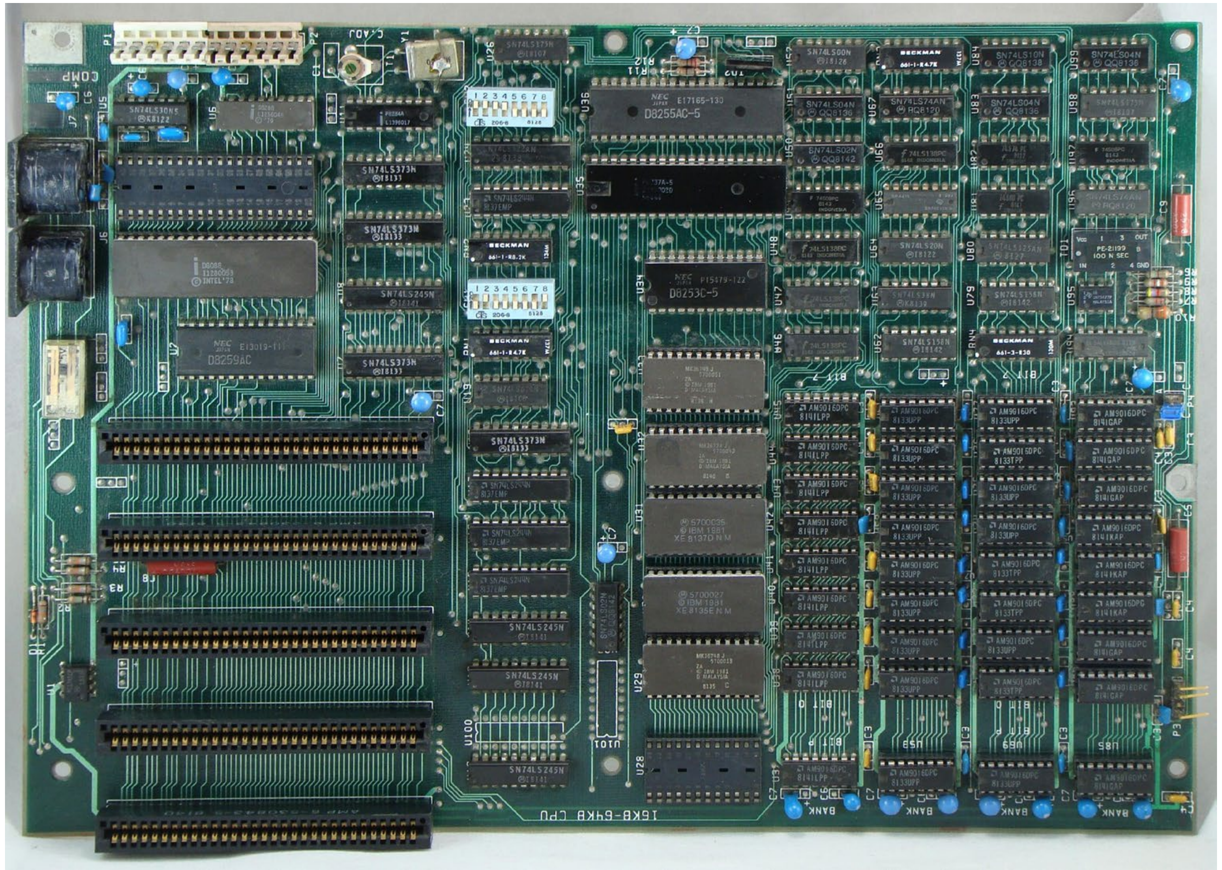

How did Chipsets Evolve? (cont.)

- As technology progressed, computer engineers decided to reduce the number of chips and have them more in a centralized location so instead of having these different chips scattered all over the motherboard controlling different functions.

- The number of chips was reduced to few chips that do the same job or what's now called a chipset.

So.. What is Chipset?

- A chipset is a smaller set of chips that has replaced a larger amount of chips.
- Chipsets job is to control data flow between the CPU, the peripherals, bus slots, and memory.
  - All the different computer hardware components communicate with the CPU through the chipset.
- Most chipsets are divided into two distinct components:
  - **Northbridge**
  - **Southbridge**

Northbridge

Function:

Connects the CPU to **high-speed components (like RAM**, **graphics card**,PCI-Express. and sometimes th

Southbridge).

Speed:

Works faster because it handles critical data transfers.

Location:

Usually found close to the CPU on the motherboard,

often **under a heatsink** (because it runs hot).

Examples of connections:

- CPU ↔ RAM
- CPU ↔ GPU (through AGP or PCIe)

Northbridge Chipset

- To communicate with any other components, the communication goes through the northbridge then the southbridge.
- Northbridge is **faster** than the southbridge.
  - Why? Northbridge is optimized for **speed** (actions per second) since it connects the **most used and most important components** of the motherboard.
  - Faster actions per section comes at a cost! The Northbridge is a source of heat and needs to be places by heatsinks and fans.
  - These days the Northbridge is often directly a part of the CPU, while the Southbridge remains separate.

Southbridge Chipset

Function:

Connects to slower peripheral devices.

It handles **input/output (I/O)** operations.

Speed:

Works slower than the Northbridge.

Location:

Found lower on the motherboard (hence "south").

Examples of connections:

- USB ports
- Hard drives (SATA/IDE)
- Audio, BIOS, network controllers

Southbridge Chipset

- Connected to the northbridge.
- Acts as a communication hub for less performance sensitive components.
  - Storage devices (SATA and IDE connectors).
  - PCI bus slots.
  - I/O Ports (USB ports, networking, sound card, etc.)
- Southbridge does not require high speed communication like the northbridge.
- To send data from the hard drive to the main memory:
  - Data goes to the southbridge.
  - Southbridge sends the data to the northbridge.
  - Northbridge will send the data to the main memory.

Modern Motherboards and Chipsets

Modern chipsets integrate the northbridge chips directly into the CPU chip, while the **remaining Southbridge functions** are handled by a single controller called the **Platform Controller Hub (PCH)**.

🚀 **New Design Advantages**:

- Sensitive hardware components **communicate directly with the CPU.**
- **Reduces system latency– data moves faster**.
- **Increases system responsiveness** and overall performance
- Less sensitive hardware (like storage and USB) **benefits from improved efficiency**

Bus Width: x86 versus x64

The chipset

When we speak about the motherboard bus width, we often refer to the **frontside bus**

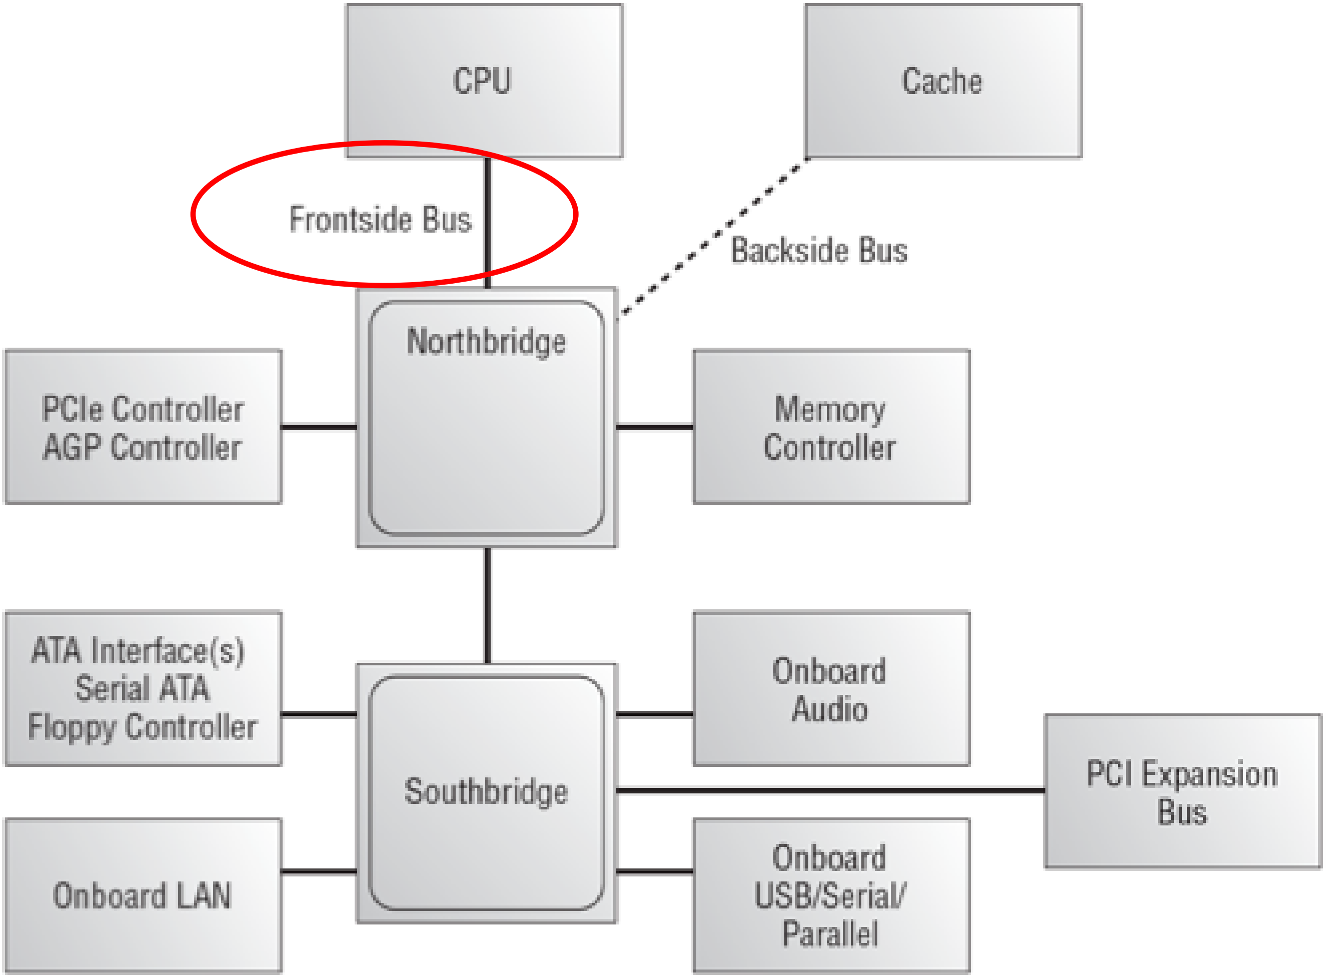

Why is the bus important?

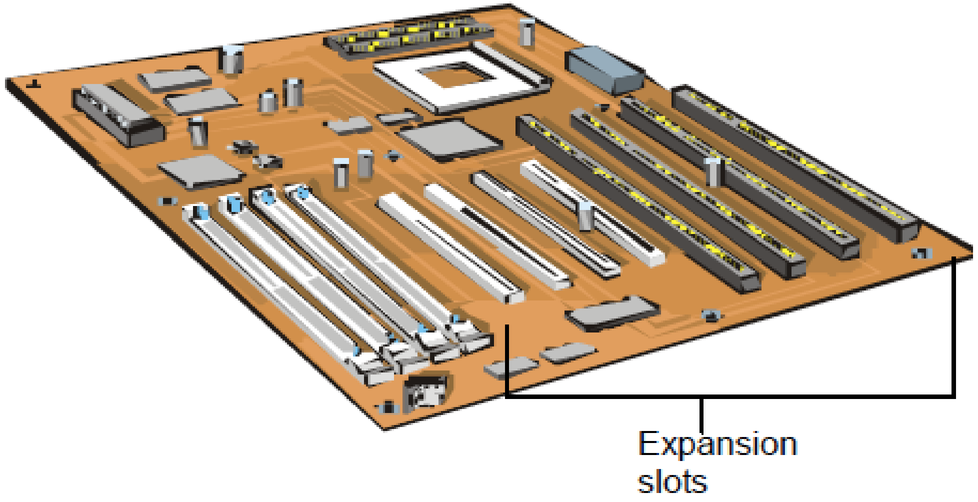

Highway

- The **number of lanes** determines how many cars can use the highway at once
- The **speed limit** determines how fast cars can **drive** on the highway

System Bus

- The **bus width** determines how much **information** can flow along the bus at a time
- The **bus speed** determines how fast information can travel

You can double check that you understand a scientific concept by checking what **units of measurement** that concept is measured by:

- Highway
  - **Number of lanes:** (a number)
  - **Speed limit:** (unit distance per time elapsed, e.g. km/h)
- System bus
  - **Bus width:** (a number of bits, eg a 32-bit processor)
  - **Bus speed:** (unit information per time elapsed, e.g. bits/second)

Highway analogy

- There's 6 lanes (6 cars can pass together)
- On each lane, 3 cars are passing per second
- How many cars are passing per second on this highway?

x86 versus x64

X86 – 32-bit systems

- 32 bits of data can travel at once in between system components
- 2^{32} ≈ 4 Billions bits at once

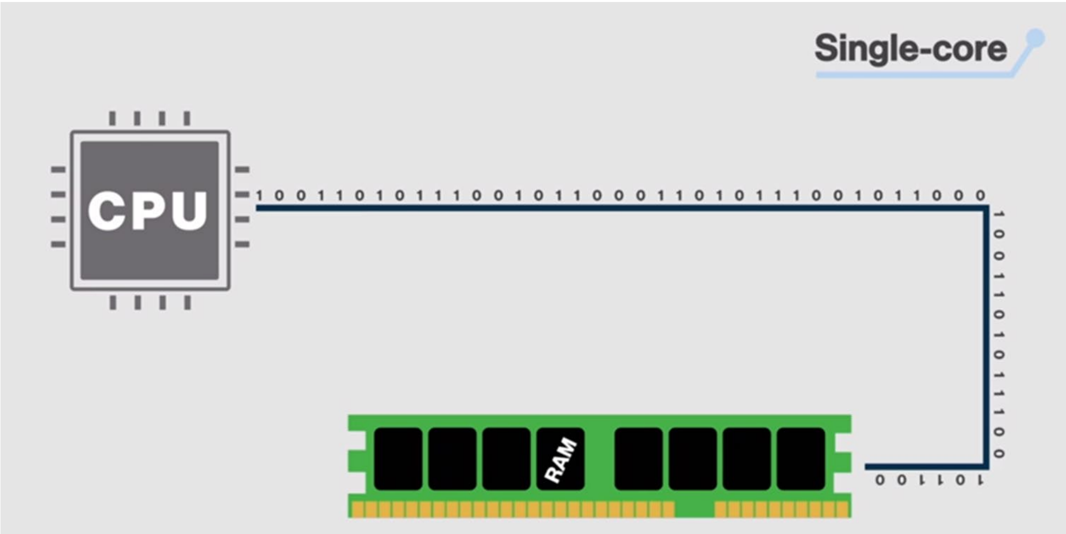

X64 – 64 bit system

- 64 bits can travel at once in our system components.
- X64: 2^{64} ≈ 4 Billions x 4 Billions bits at once.–that's roughly **18 quintillion bits!**

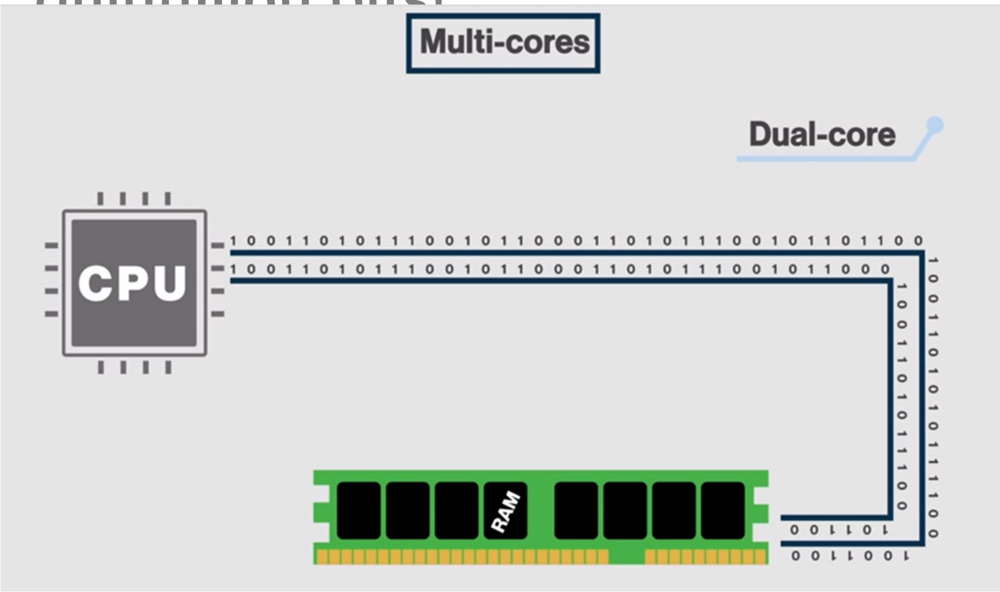

Modern Systems and x64 Architecture

Almost all **modern hardware and software** today use **x64 (64-bit)** architecture.

64-bits system offer:

- 🔒 **Better security** features
- 🎨 **Improved graphics and performance**
- ⚡ **Support for larger amounts of memory**

💡 But! Some older applications still rely on **32-bit (x86)** support — so many systems remain **backward compatible** to run older software.c).

Central Processing Unit

Central Processing Unit (CPU) – Recap

- Most important element of the computer in terms of computing power.
  - Brains of the computer.
  - Responsible for interpreting and executing the commands from the software.
  - The processor executes a program (sequence of stored instructions).
    - While the CPU is processing one step of the program, the remaining instructions and the data are stored in a special memory called cache.
- Almost all modern devices include a CPU:
  - Desktops, laptops, tablets, smartphones, TVs, TV boxes, Vending machines, traffic lights, Smart devices (Fridge, washer, light bulbs)

💻 What Do Cores Do in a CPU?

🧠 CPU cores are responsible for:

- ⚙️ **Executing program instructions**
- ➗ **Performing calculations**
- 📊 **Managing data flow** inside the CPU
- 🤝 **Coordinating with other components**

CPU Reminder (Lecture 4)

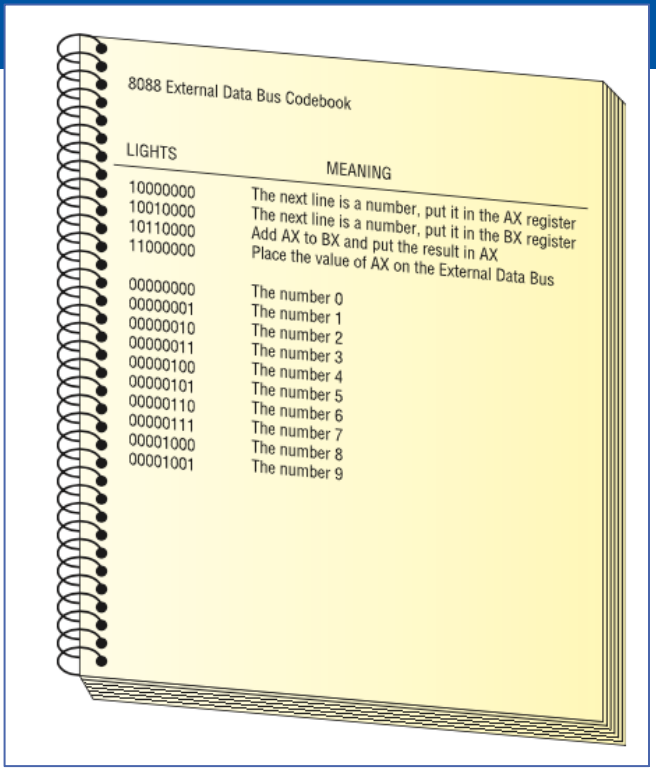

10000000	The next line is a number, put it in the AX register
00000010	The number 2
10010000	The next line is a number, put it in the BX register
00000011	The number 3
10110000	Add AX to BX and put the result in AX
11000000	Place the value of AX on the external data bus

A set of commands such as this is called a program.

A **program** is a series of command sent to a **CPU** in a specific order for the CPU to perform work.

CPU History – Recap

First CPU: **Intel 4004 processor:**

- Developed at Intel Corp by Ted Hoff.
- Introduced in 1971.

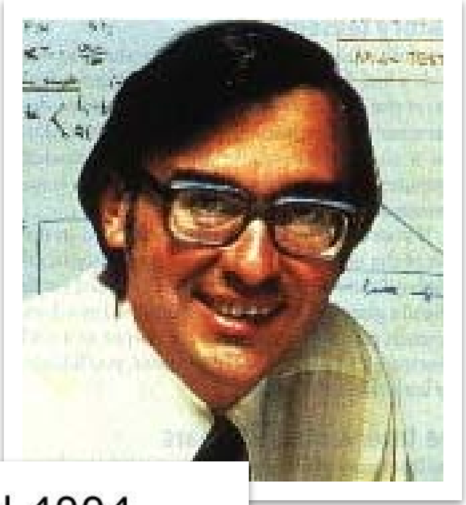

Design:

- **Clock Speed:** 108 kHz
- **Transistors:** 2,300
- **Performance:** 60,000 operations per second (OPS)
- **Memory Addressed:** 640 bytes
- **Cost at Launch:** $200

💡 **Fun Fact:**

The Intel 4004 was the first commercially available microprocessor, laying the foundation for all modern CPUs.

CPU Manufacturers – Reminder

For Desktops, Laptops, and Servers:

- **Intel** and **AMD** are the two most popular CPU manufacturers.

For Mobile Devices:

- **NVIDIA**, **Qualcomm**, **Samsung**, **and** **Apple** are major manufacturers of mobile CPUs.
- These CPUs are used in **smartphones, tablets, TV boxes, and other mobile devices**.

💡 **Note:**
Desktop/server CPUs focus on **high performance**, while mobile CPUs are designed for **power efficiency and compact size**.

CPU Manufacturers – Recap

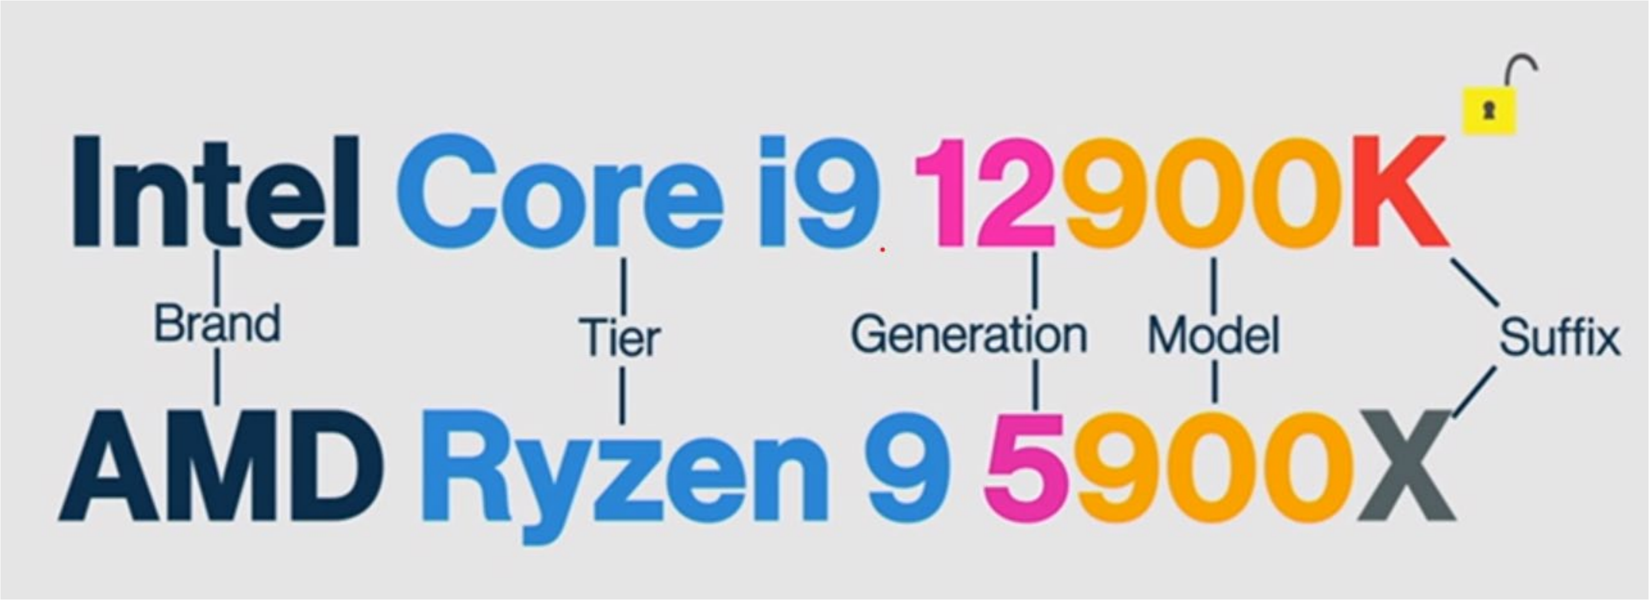

CPU Architecture

CPU Components

There are **two primary components** :

- Arithmetic Logic Unit – **ALU** :
  - Performs basic arithmetic operations like adding, subtracting, multiplication and division of numbers.
  - Performs logical operations like the comparison.
- Control Unit – **CU**:
  - Manages **all operations** within the processor
  - Instructs the ALU on which arithmetic/logical operation is to be performed.
  - Works in sync with the **system clock**
  - Routes **data internally**, ensuring it is **fetched and delivered** to the correct locations

Other CPU Components

- **Registers**: a temporary storage used to store data:
  - Coming from memory heading to the processor for execution.
  - Coming from the processor after processing
  - Being worked on by the processor.
- Internal Data Bus: connects the internal components of the processor to the motherboard.
- External Data Path: used to fetch data from memory to the processor.
- Address Lines: used to specify the exact location in memory where data can be found.

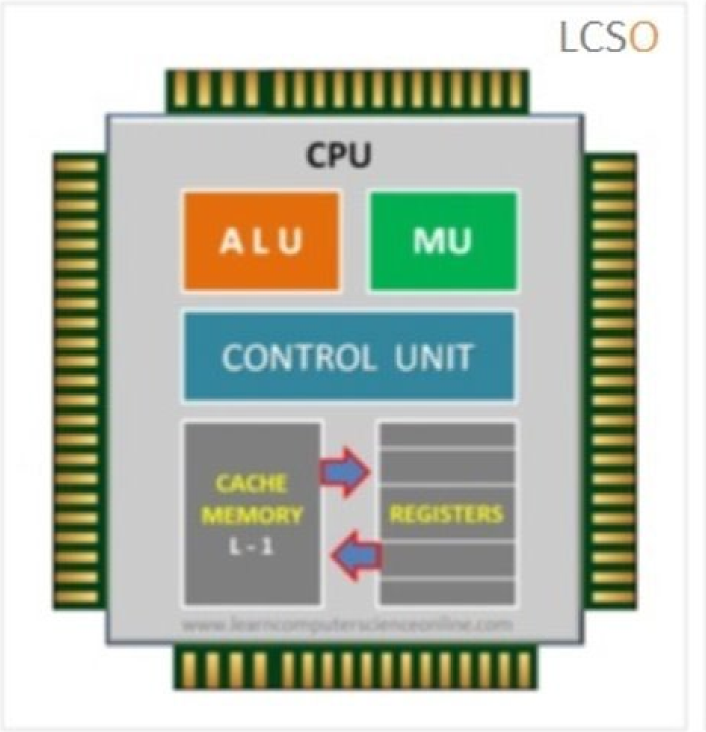

Pipelining - Reminder

- The CPU takes at least four stages (**fetch, decode, execute, write/repeat**) to do something. See https://en.wikipedia.org/wiki/Instruction_cycle for more detail

  Small discrete circuits inside the CPU handle each of these stages.
- By organizing these circuits in a conveyor-belt fashion (called pipelining), our "man-in-a-box" becomes "men-in-a-box", thus improving efficiency.
- Helps system to run faster.
  Not perfect and sometimes stalls.
  Some CPUs therefore have multiple pipelines.

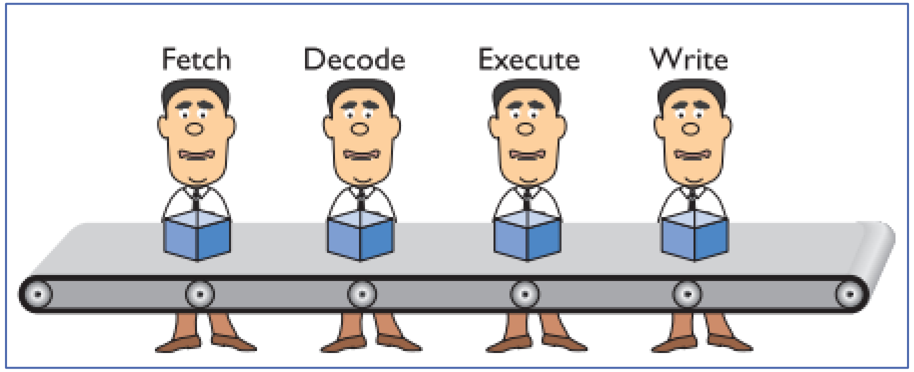

Pipelining

Pipelining several instructions processed **simultaneously** in different stages **within the same core**

A classic **5-stage pipeline** (used in many processors) includes:

- **IF – Instruction Fetch:** CPU fetches an instruction from memory.
- **ID – Instruction Decode:** CPU decodes the instruction and identifies what to do.
- **EX – Execute:** Performs arithmetic or logic operation (ALU).
- **MEM – Memory Access:** Reads or writes data to memory (load/store instructions only).
- **WB – Write Back:** Writes the result back to a CPU register.

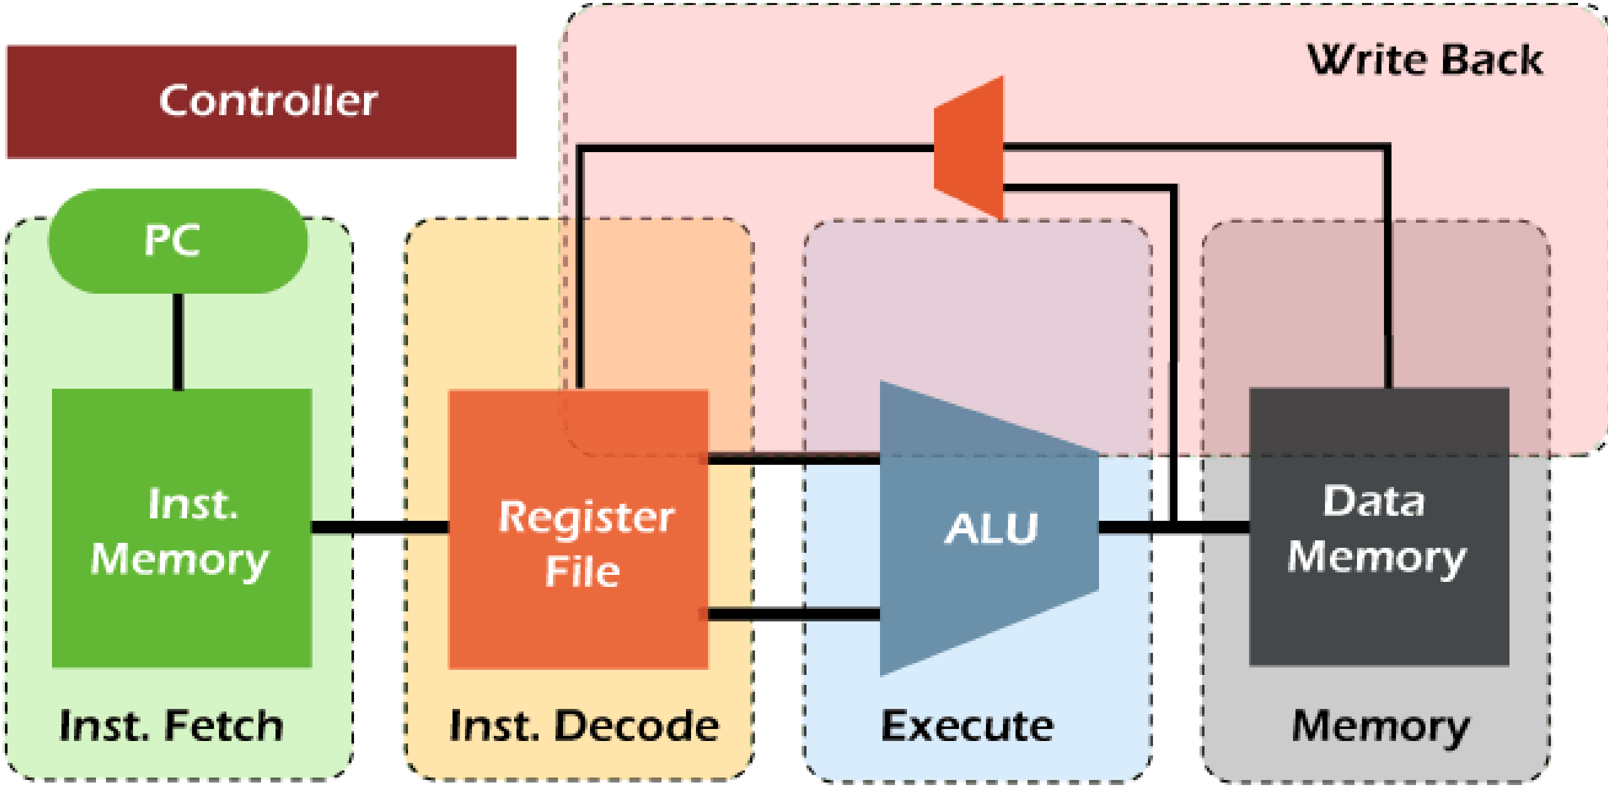

CPU Power

⚡ CPU power is determinded by **two main factors**: **speed** and **data width** (amt of data it **can process**.)

1 Speed:

- Measured in cycles per second,
  - megahertz (MHz) = millions of cycles per second
  - gigahertz (GHz).= billions of cycles per second
- **Hertz (Hz)** represents "per second," indicating how many cycles the CPU can perform each second.

2 Data Width (Bus Size)

- The amount of data the CPU can process at one time depends on the **Front Side Bus (FSB)**, also called the **CPU or processor bus**.
- **Bus width** is measured in **bits**, e.g., 32-bit or 64-bit FSB.
- A wider bus allows the CPU to **transfer more data per cycle**, increasing performance.

CPU Features – How to choose a CPU?

1) How to Connect a CPU?

- CPU connection on the motherboard is known as the CPU socket.
- The CPU is connected to the motherboard in one of two ways:

- **Pin Grid Array – PGA**
  - Pins are on the CPU (underside of the processor).
  - Socket on motherboard includes holes that line up with the CPU pins.

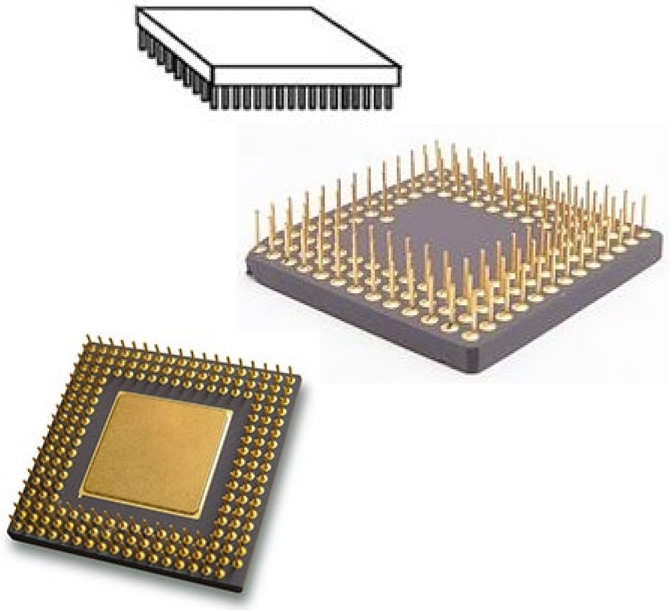

- **Land Grid Array – LGA**
  - Pins are on the motherboard (on socket) instead of on the processor.

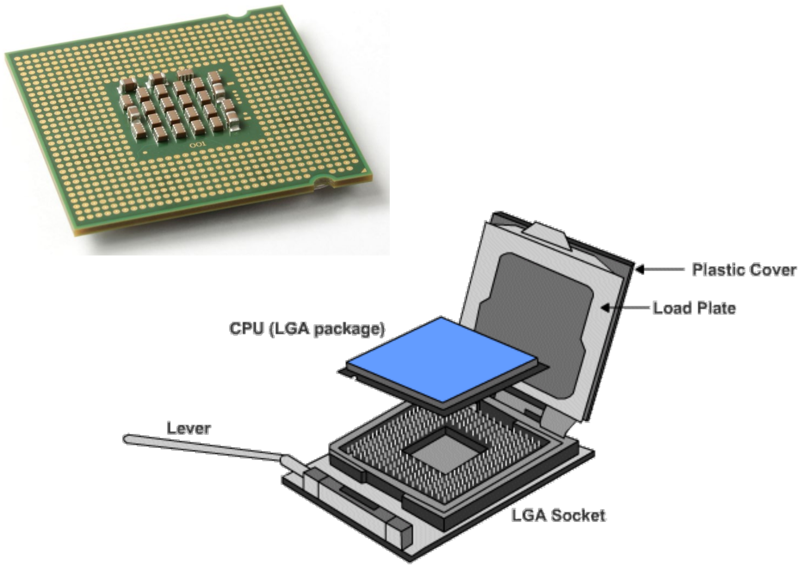

2) Clock Speed

- CPU clock speed: number of tasks per second it can perform.
- 3.5 Gigahertz (Ghz) = 3.5 billions of cycles per second.
- 1 Hz = 1 "thing"per second, it's a unit of **frequency**.
- 1 KHz = 1000 Hz
- 1 MHz = 1 million Hz
- 1 GHz = 1 billion Hz

2) Clock Speed (overclocking and throttling)

Overclocking:

- A technique to make a processor run **faster than its original specification**
- ⚠️ Not generally recommended as it can **increase heat and reduce CPU lifespan**

Throttling:

- When the processor runs **below its rated speed**
- Helps **conserve power** and **reduce heat**
- Commonly used in **laptops and mobile devices**

💡 **Summary:**

- **Overclocking → higher performance, more heat**
- **Throttling → lower performance, less heat/power consumption**

3) CPU Cores

Multi-Cores: More than one CPU core integrated into a single chip, each capable of running parallel processes

- Each core has its separate ALU unit, control unit and Registers.

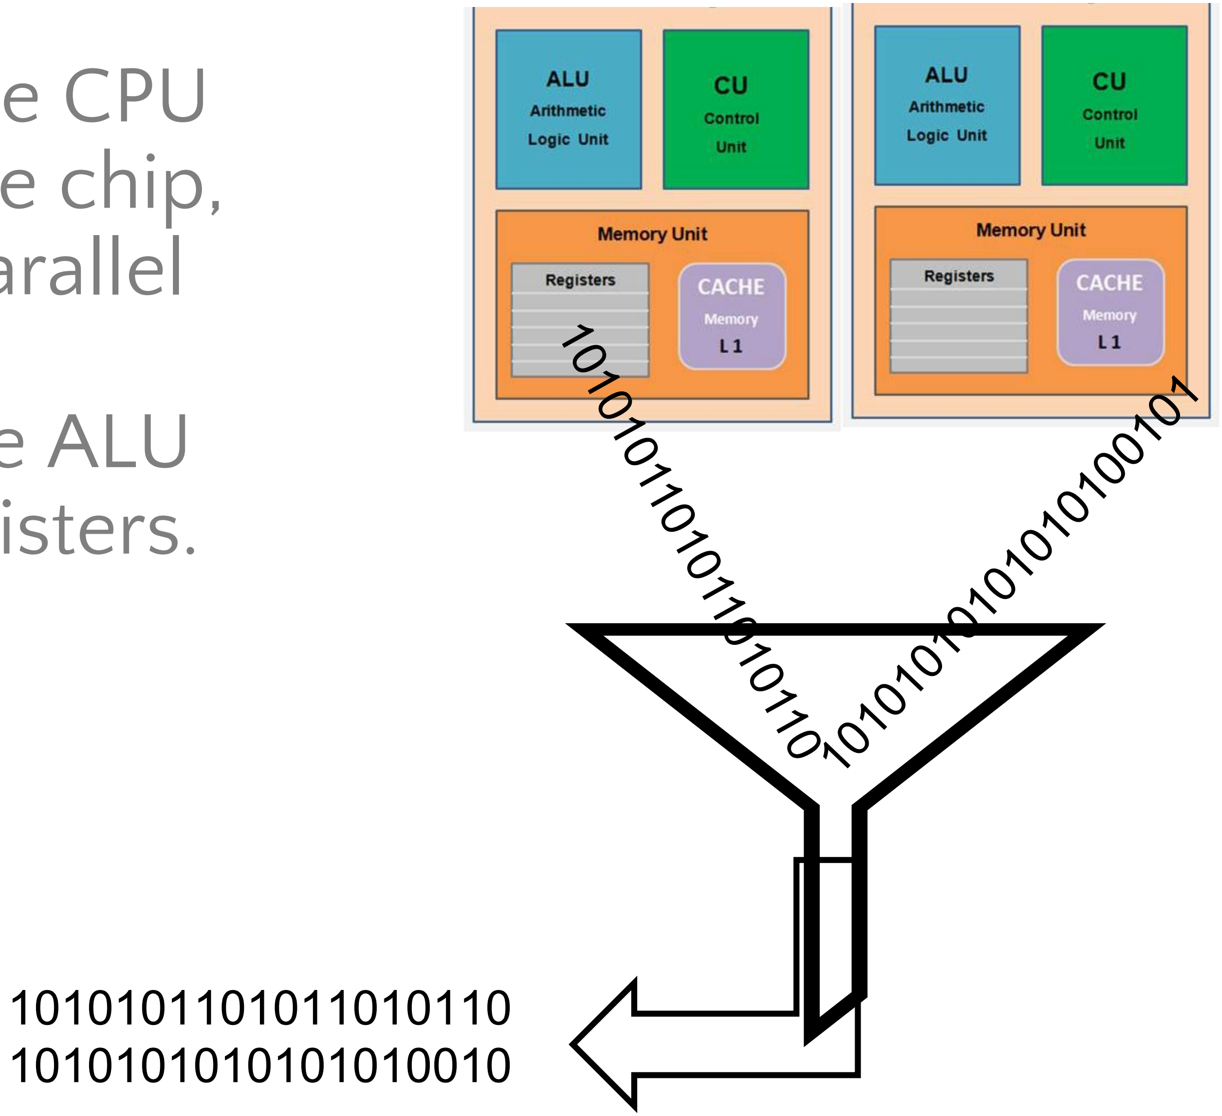

3) CPU Cores

Single Core

- Most CPUs are considered single Core.
- This means that only one data "lane" connects the CPU to the RAM.

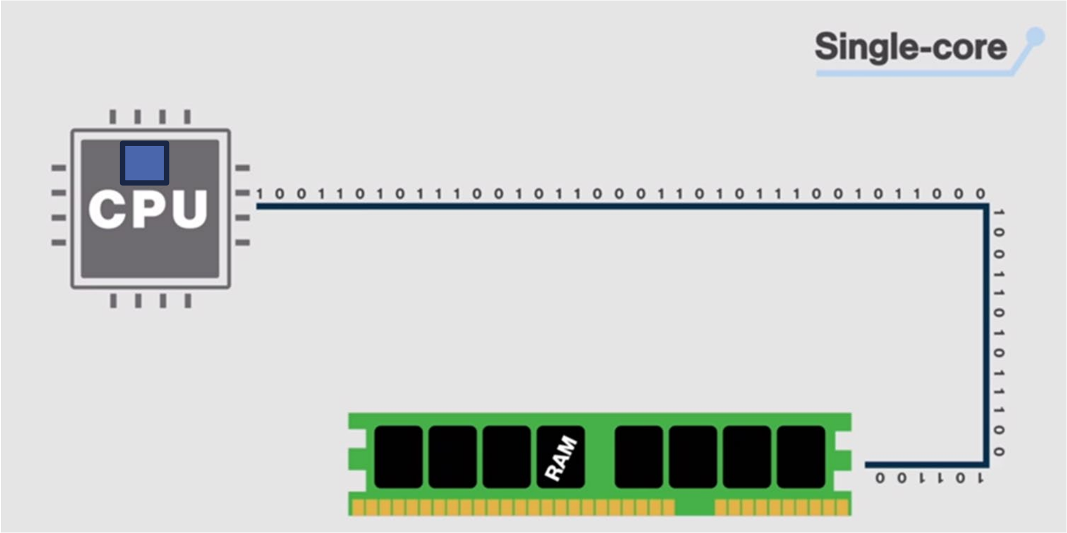

Dual- Core

- A Dual-Core CPU is a CPU with two lanes of data travelling to the RAM.
- 4 lanes ==> Quad-core

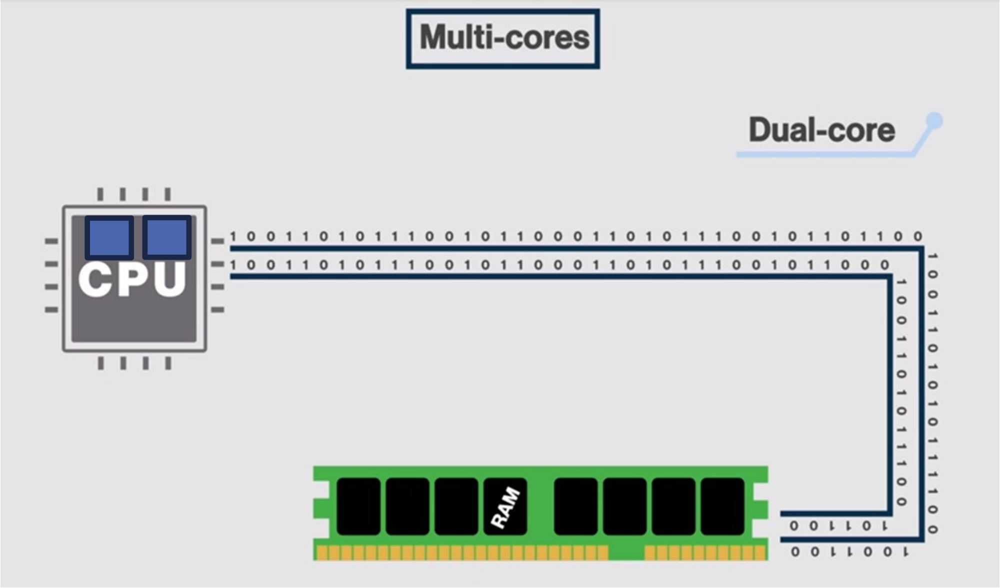

3) High-End CPU- (The threadripper)

- **Threadrippers:**
  - are CPUs with **32 and even 64 cores.**
  - Designed for **extreme multitasking**, content creation, and heavy computational workloads
- **Example:AMD EPYC 75F3**
  - **Clock Speed:** 2.95 GHz
  - **Cores / Threads:** 32 cores / 64 threads
  - **Cache:** 256 MB
  - **Power Consumption:** 280 W
  - **Cost:** CAD $15,699

4) CPU Cache

- The CPU store the most recently used data and instructions in its **cache**.
- This **speeds up the processing** as the cache is much closer to the ALU unit and increases the performance greatly.

Cache vs RAM

- **CPU Cache:** Uses **Static RAM (SRAM)**
  - Faster, more expensive, and retains data **without constant refreshing**
- **Main Memory (RAM):** Uses **Dynamic RAM (DRAM)**
  - Slower than cache, cheaper, and requires **constant refreshing** to maintain data

4) CPU Cache

- Cache = small, super-fast memory inside the CPU
- **SRAM in cache** provides quick access to data and instructions the CPU needs most
- **DRAM in main memory** provides larger, slower storage for running programs

SRAM vs DRAM

Static RAM

- Uses transistors memory cells.
- No need for power refreshing to keep data
- Low power consumption
- Very expensive ≈5000$/Gb
- Very fast read/write
- Volatile (data erased at power off)
- Only a few MB used in a PC

Dynamic RAM

- Uses transistors and capacitors
- Power refreshing required to keep data saved
- Higher power consumption
- Good balance price/amount ≈20-75$/Gb
- Fast, but not as fast as SRAM
- Volatile (data erased at power off)
- 4-32GB of RAM used in a PC

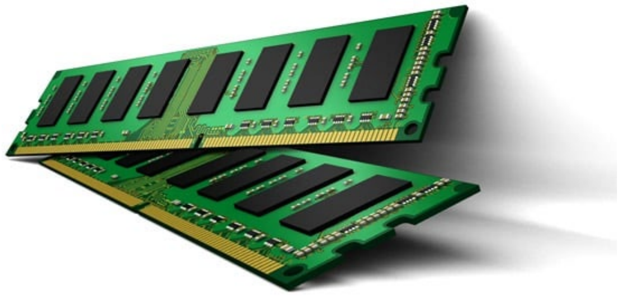

4) CPU Cache Levels (Static RAM – SRAM)

CPU caches are organized into **three levels** to balance speed and size: :

1 **L1 Cache – Internal Cache**

- Located inside each CPU core
- Very fast
- Typical size: 16 KB – 256 KB per core

2 **L2 Cache – External Cache**

- Fast, slightly larger than L1
- Dedicated to each core
- Typical size: 128 KB – 8 MB per core

3 **L3 Cache – Shared Cache**

- Largest and slowest cache
- Originally mounted on the motherboard, now integrated into the CPU
- Shared by all CPU cores
- Typical size: up to 64 MB

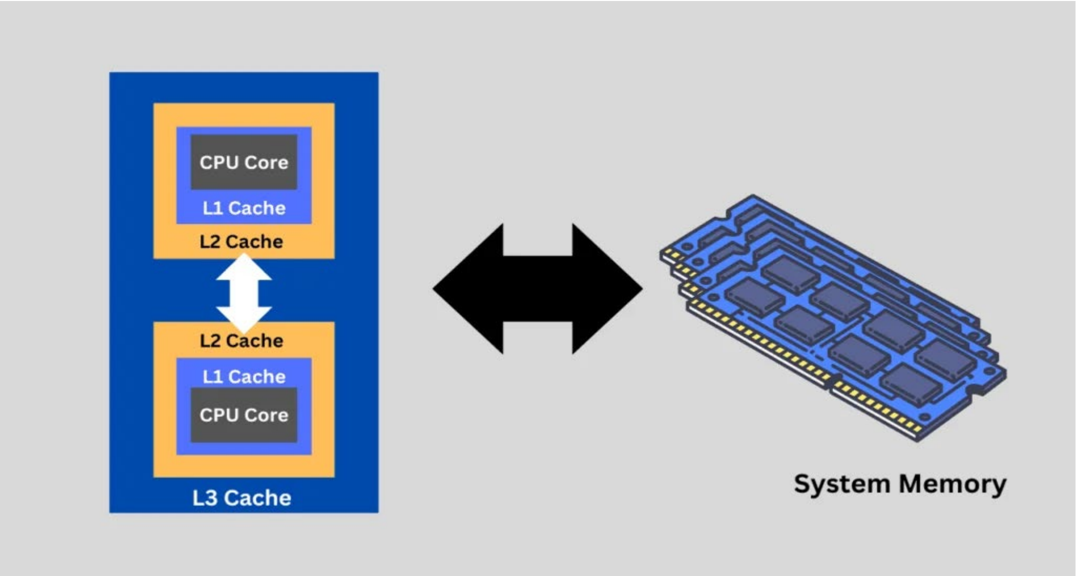

4) CPU Cache Levels (Static RAM – SRAM)

Analogy: Office Filing System

- Imagine the **CPU as a worker** trying to complete tasks.
- **L1 Cache:** Desk drawer
  - Holds the **papers you are currently working on**
  - Very **fast access**, but small capacity
- **L2 Cache:** Filing cabinet
  - Holds **papers you might need soon**
  - Slower than the desk, but bigger
- **L3 Cache:** Storage room
  - Holds **papers for everyone in the office**
  - Slowest among caches, but largest

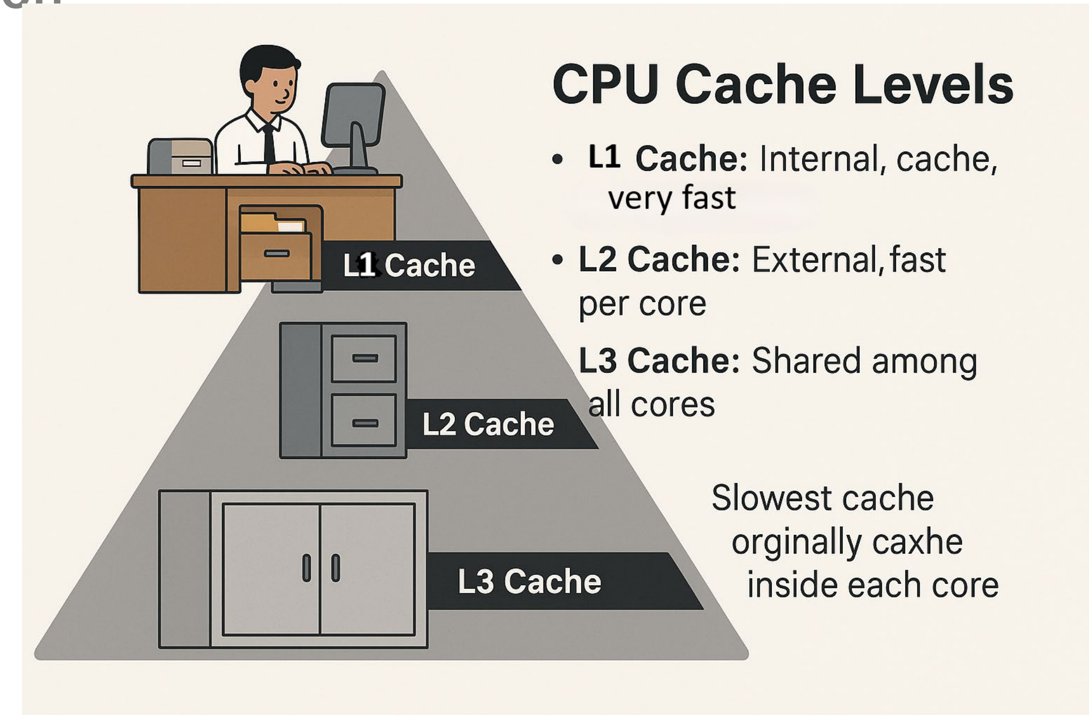

4) CPU Cache Levels (Static RAM – SRAM)

L1 → **fastest, smallest, core-specific**

L2 → **fast, medium-sized, core-specific**

L3 → **largest, slower, shared across cores**

Purpose: Each level provides a trade-off between **speed** and **capacity** to maximize CPU performance.

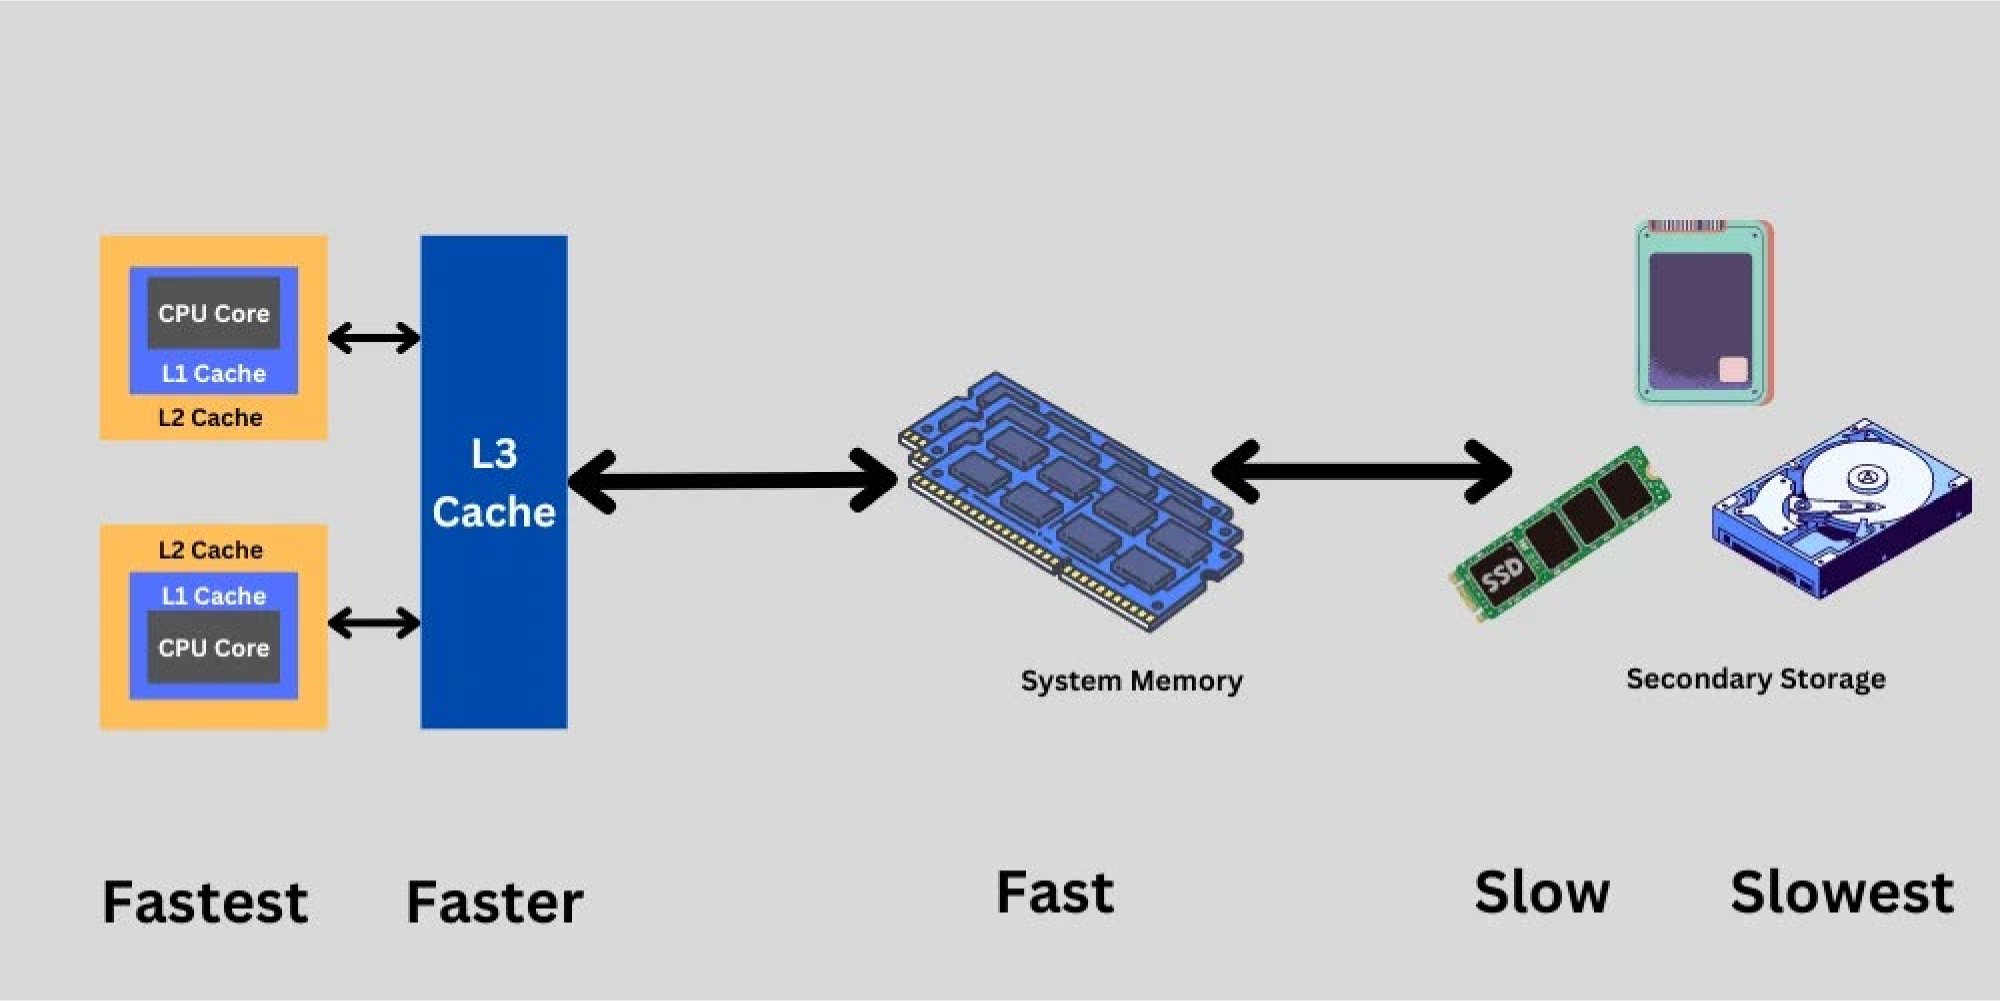
CPU Core
L1 Cache
L2 Cache
L2 Cache
L1 Cache
CPU Core
L3 Cache
System Memory
SSD
Secondary Storage
Fastest
Faster
Fast
Slow
Slowest

4) CPU Cache – Data flow

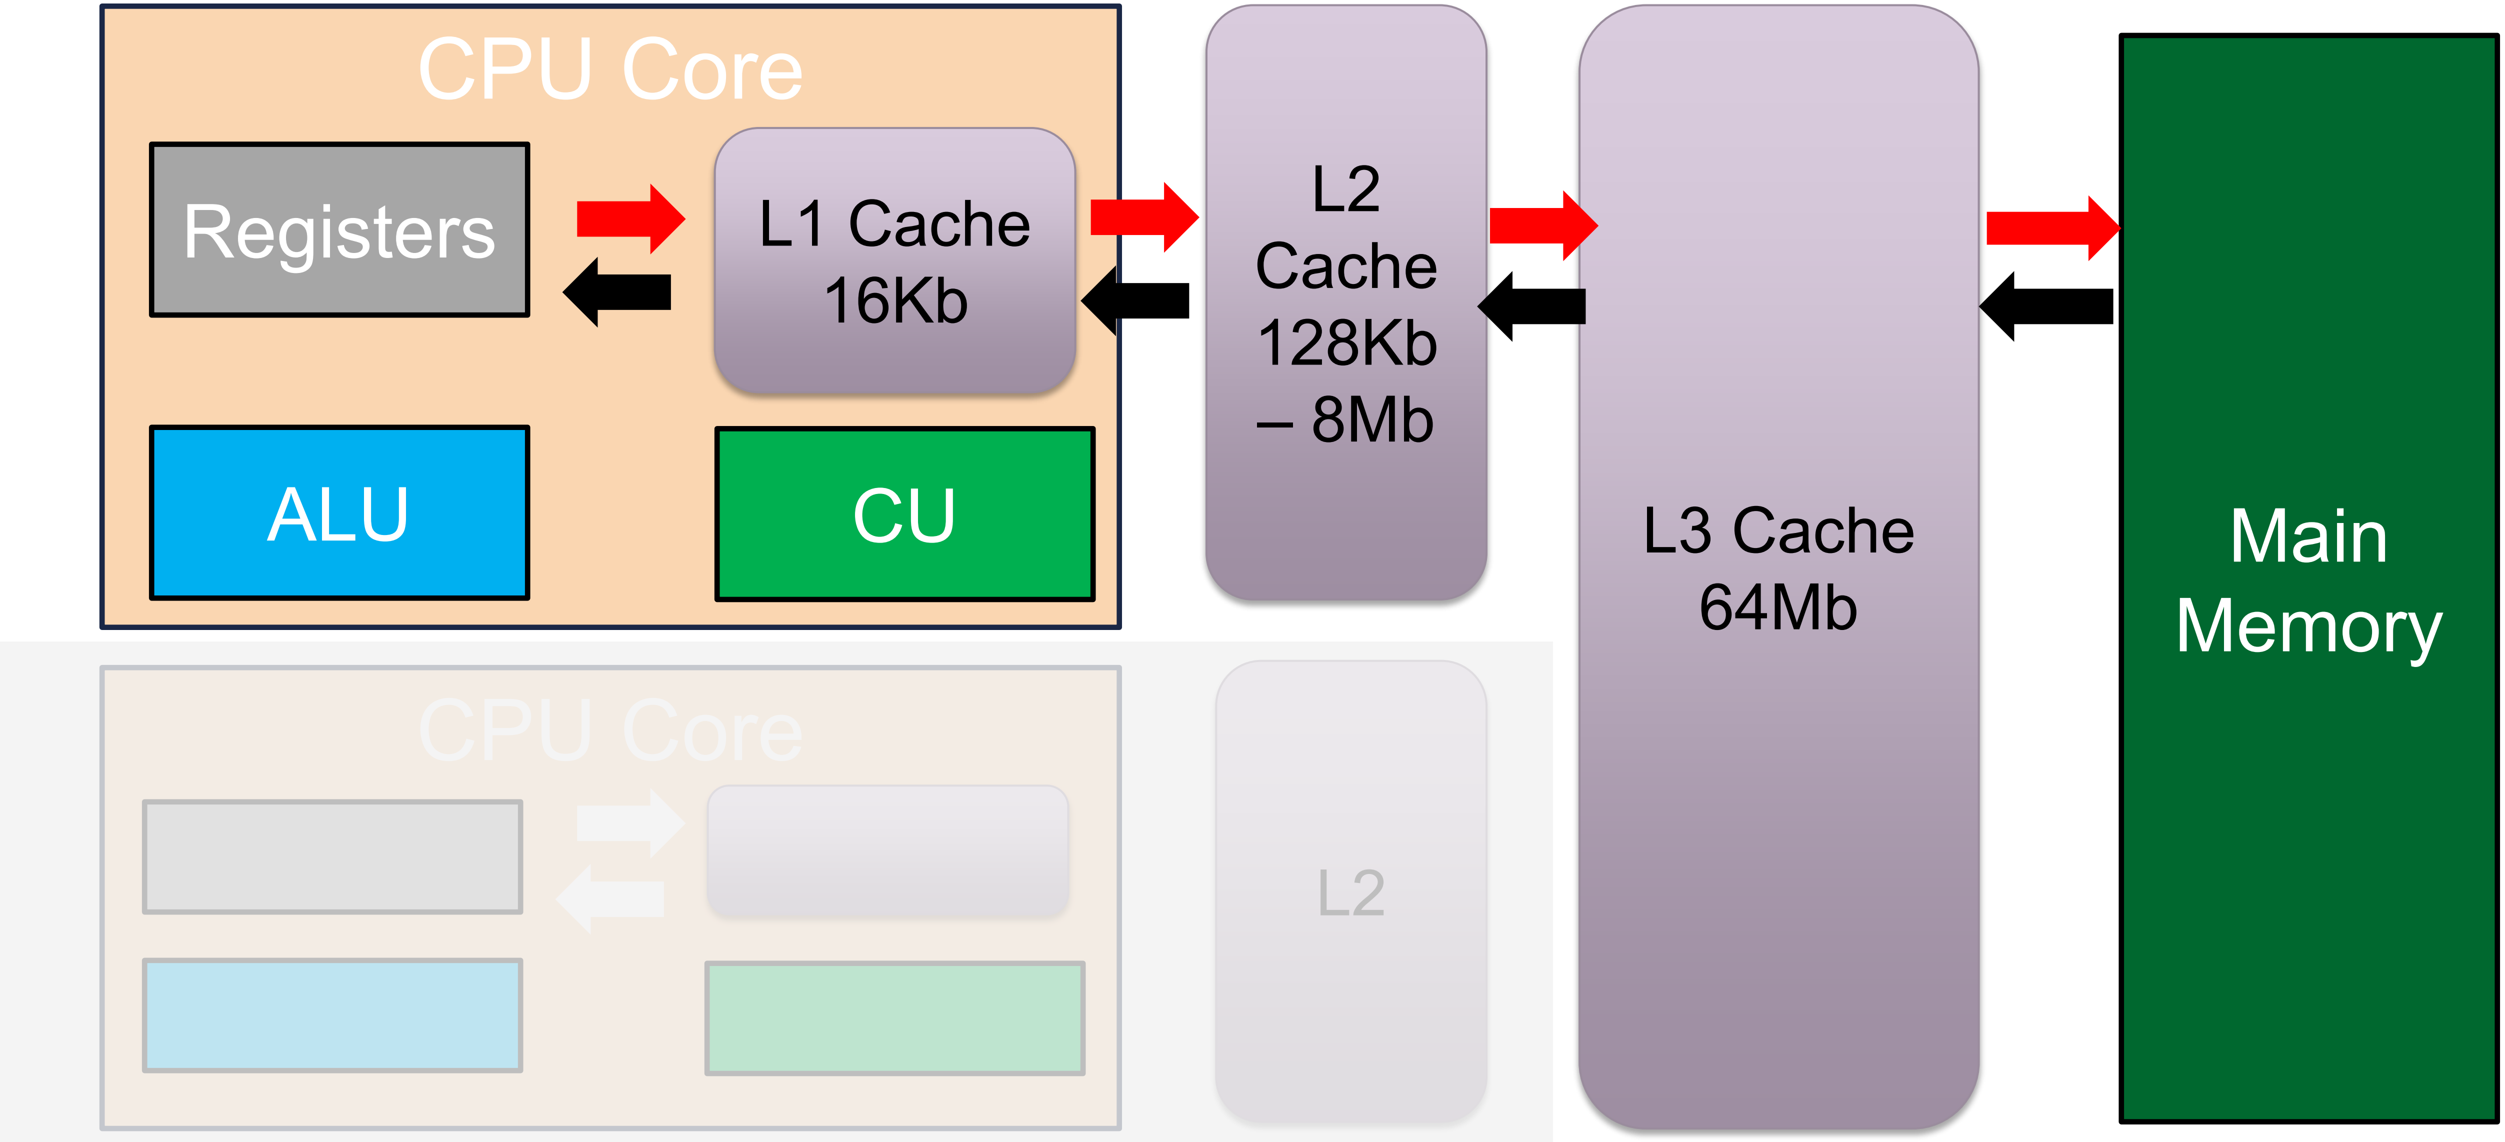

4) CPU Cache

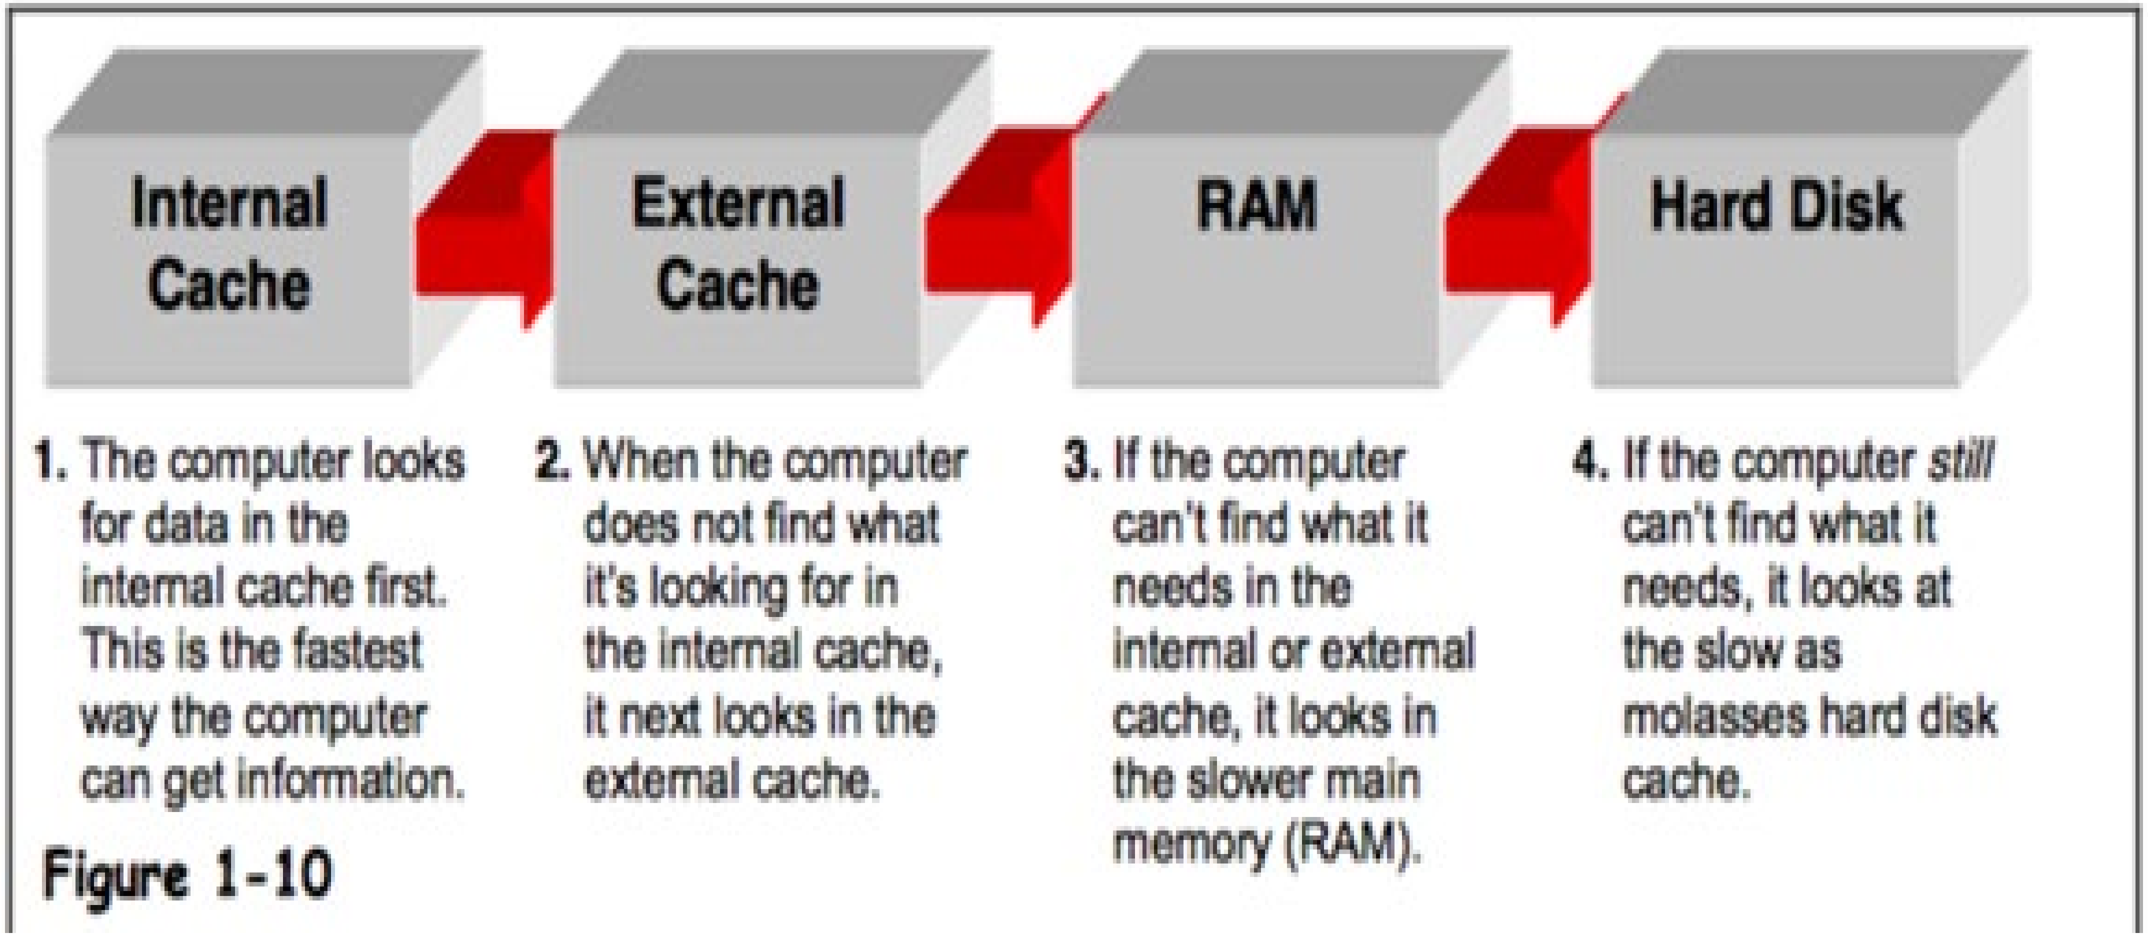
Internal Cache
External Cache
RAM
Hard Disk
1. The computer looks for data in the internal cache first. This is the fastest way the computer can get information.
2. When the computer does not find what it's looking for in the internal cache, it next looks in the external cache.
3. If the computer can't find what it needs in the internal or external cache, it looks in the slower main memory (RAM).
4. If the computer *still* can't find what it needs, it looks at the slow as molasses hard disk cache.
Figure 1-10

Activity: Bus Speed and Bus Width (CPU Power)

How to choose a CPU?

Let's say we have a 100-MHz, 32-bit bus CPU.

- How many bytes per second can the CPU process?

Now, let's compare the above CPU with a second 80-MHz 64-bit bus CPU.

- How many bytes per second can the second CPU process? Which of the two is faster at sending information?

First CPU: 400 Mbytes/s Second CPU: 640 Mbytes/s

Power consumption of the first CPU is bigger due to the higher clock rate.

Other Types of CPUs

ARM Chip

- Simplified CPUs called **Advanced Risk Machine.**
- More simplified instruction set
- Much higher speed of processing and reduced workload.
- Used in Arduinos and RasberryPis
- The Apple M1 Chip used in Macbook Air is an ARM chip

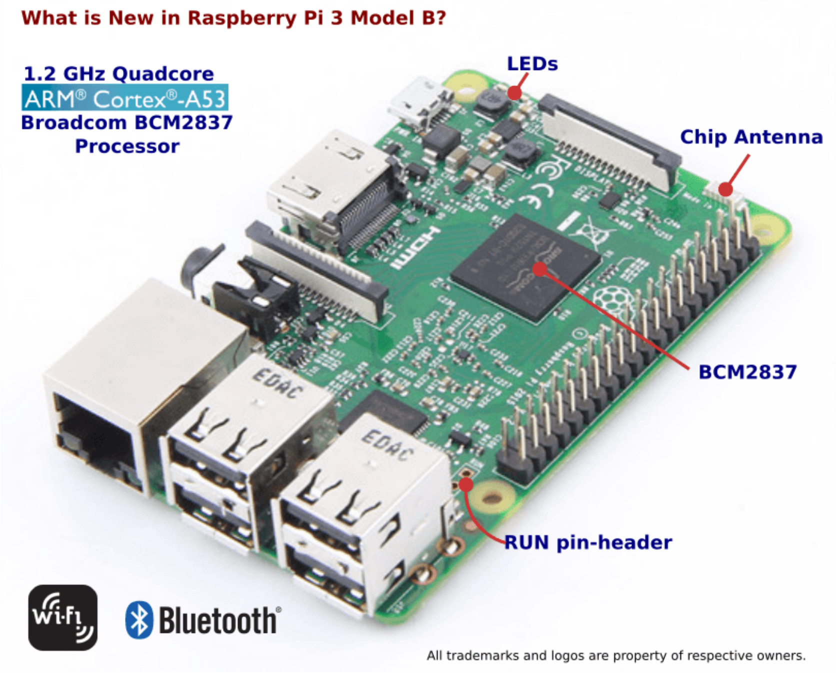

Accelerated Processing Unit (APU)

- APUs are CPUs with combined graphics processing capabilities.
- Not as powerful as a dedicated GPU.
- Great option for budget gaming.

Main Memory

What is the Main Memory?

- Referred to as RAM (Random Access Memory)
- RAM is a collection of integrated circuits (chips).
- Considered the “Primary” storage for the CPU.
- Temporary storage for data and programs that are being accessed by the CPU.
- Volatile memory: contents are erased when the computer is switched off. (or when power is lost)

Why use Main Memory?

- Did you ever thought why games take time to load?

This is the time it takes to copy all the files of the game from the storage device (long term memory) into the working memory of the main memory, which the CPU can access must faster and much easier.

Read/Write speed in RAM is 12 800 Mb/s
Read/Write speed in SSD is 456Mb/s

Why use Main Memory?

- More RAM means more capacity to hold and process large programs and files.

- Maximum installed is limited by the motherboard and the operating system.

Main Memory History

- Early computers had RAM installed on the motherboard as individual chips.
  - These were difficult to install and often became loose.

- Designers and engineers then created memory modules.
  - Special circuit boards on which the memory chips are soldered on.
  - Memory modules are attached to a motherboard using memory slots.
  - New design allows to add or replace memory modules.

Main Memory Features

1) Data Rate

SDRAM

- Original Main memory was the Synchronous Dynamic Random Access Memory (SDRAM)
- It's **synchronized** with the motherboard speed.
- For a 32-bit bus, if the motherboard goes up to 100Mhz, the SRAM speed was also 100MHz x 32bits of data send/received.

DDRAM

- The first leap in RAM technology was the Double Data RAM DDRRAM.
- For every tick of the motherboard clock, DDRAM sends/receives 2 bits of information instead of just 1.
- This "doubles" the speed: 2 x 100MHz x 32 bits

1) Data Rate

- How can we make the difference between SDRAM, DDR, DRR2, etc.
- The number of pins and the position of the notch

	Number pins	Number notch	Clock Speed	DDR Speed	PC Speed (in bits)
SD-RAM	168	**2**	100 MHz	SD-100	PC-800
DDRAM	184	1	100 MHz	DDR-200	PC-1600
DDR2	240	1	100 MHz	DDR2-400	PC2-3200
DDR3	240	1	100 MHz	DDR3-800	PC3-6400
DDR4	288	1	100 MHz	DDR4-1600	PC4-12800
DDR5	288	1	100 MHz	DDR5-3200	PC5-25600

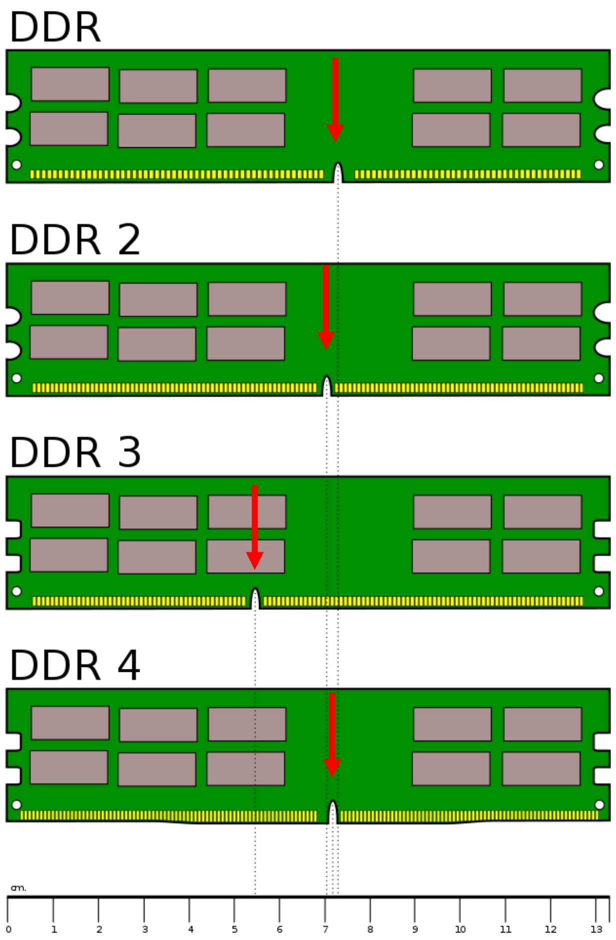

2) RAM Capacity

- Every Stick of a RAM has a capacity (in GB) which is the amount of information that the RAM can store.
- The capacity depends on the size of the chips (1Gb, 4Gb, 16Gb, etc)
- The capacity depends if the RAM is SIMM or DIMM

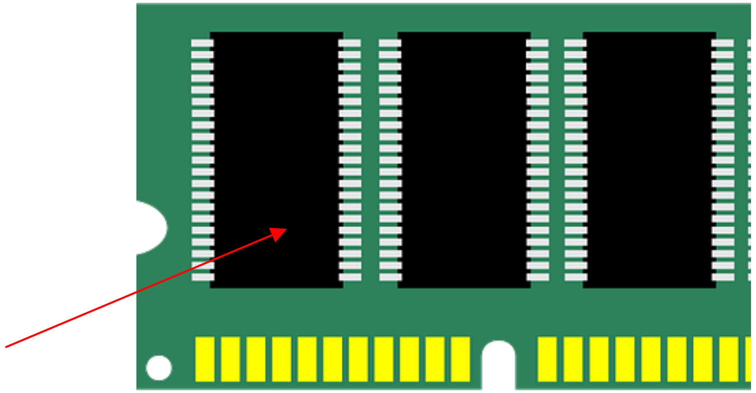

Serial Presence Detect (SPD)

How can we get all the specs of currently installed RAM sticks?

- Each RAM stick includes an **SPD chip (Serial Presence Detect)** that stores detailed information about the memory — such as speed, size, manufacturer, and timings.

.

- A tool like **CPU-Z** can **read data from the SPD chip**, allowing you to view all the RAM specifications **without physically opening your computer**.

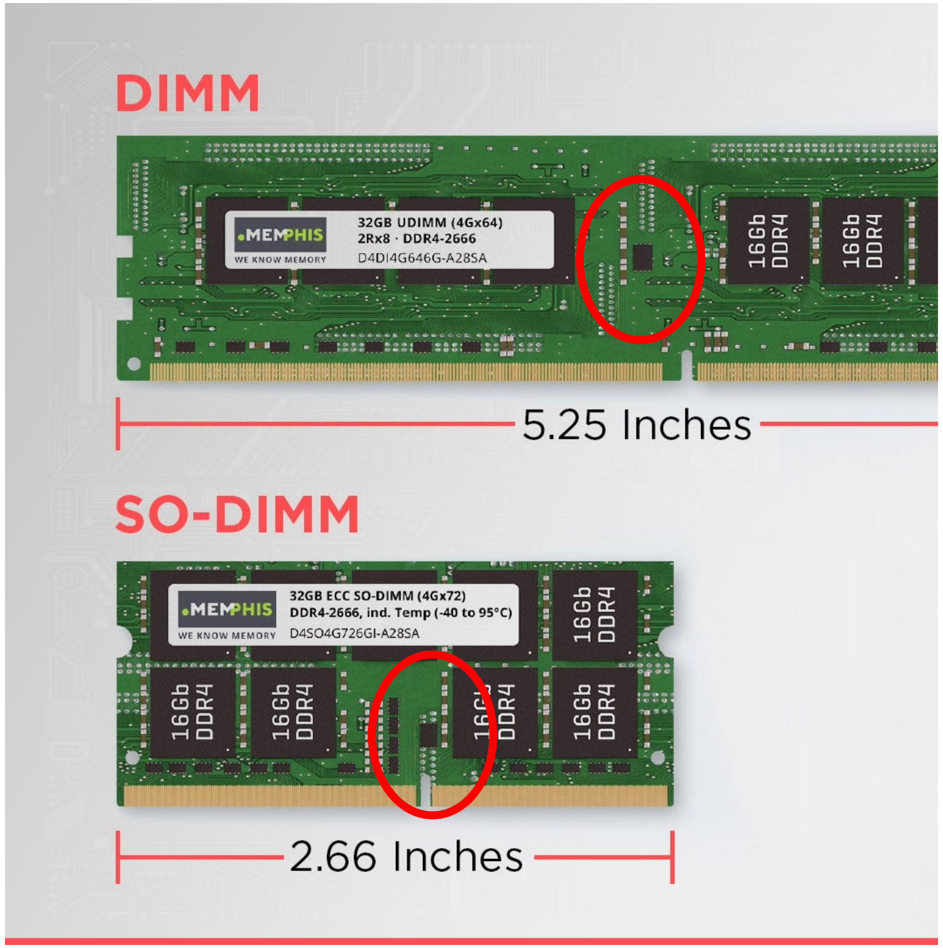

2) RAM Capacity: SIMM vs DIMM

SIMM (Single Inline Memory Module)

- Older type of memory module.
- Has **a single row of memory chips** and **one set of electrical contacts**.
- Both sides of the stick are **connected to the same circuit**.
- Common in early computers (e.g., Pentium I).

DIMM (Double Sided RAM)

- Modern type of RAM used in most computers today.
- Has separate electrical contacts on each side of the module.
- Can store memory chips on **both sides**, effectively doubling the memory capacity (e.g., 2 GB → 4 GB, 8 GB → 16 GB).
- Supports wider data paths (64-bit vs 32-bit in SIMM), leading to faster performance.

3) RAM Channel

- Most motherboards come with limited **number of RAM slots** —usually **2 or 4 slots** in total.
- RAM is typically sold in **pairs of RAM sticks** to take advantage of **dual channel**

Performance.

- Most modern motherboards support **dual-channel memory architecture** — this allow the CPU to access data from **two RAM sticks at the same time** instead of just one.

  ➡️ This **doubles the data bandwidth**, making your system faster and more efficient.

3) RAM Channel

Motherboards have multiple RAM slots, usually four, labeled something like:
A1, A2, B1, B2.

- Slots A1 and A2 belong to one memory channel.
- Slots B1 and B2 belong to the second channel.
- To activate dual channel, you must install the RAM sticks in matching colored or labeled slots, e.g. A2 and B2.

- For the best performance, the **RAM sticks in the same channel should be identical** — same size, speed, and brand if possible.

If you place the RAM sticks in the incorrect way on the motherboard, the dual channel will not be active, and you won’t use the PC to its best capacity.

3) RAM Channel

🚫 What happens if installed incorrectly

If you put both RAM sticks in the **same channel** (e.g., A1 and A2), the motherboard will treat them as **single-channel** memory.

That means:

- The system can only access one stick at a time.
- Memory bandwidth is cut in half.
- Overall performance drops (especially noticeable in gaming, video editing, or multitasking).

4) Memory Redundancy / Error Checking

Memory errors happen when **data is not stored or retrieved correctly** from the RAM chips.

Computers use different types memory to detect and correct data errors in memory.

- **Non-parity memory:** Cannot detect memory errors.
- **Parity memory:** Can detect errors, but cannot correct them.
- **ECC (Error-Correcting Code) memory:** Can both **detect and correct** common memory errors automatically.

4) Parity vs ECC

Non-parity memory	Cheapest level of redundancy, works fine for most applications. If a memory error occurs, the OS has no way of knowing. Data can be therefore be corrupted without warning.
Parity memory	Also called true parity. Parity memory contains extra chips that act as parity chips. A memory error can be detected, and the OS is informed.
ECC memory	Also called error checking and correcting, most expensive. Detects multiple-bit memory errors and can correct 1-bit errors.

5) RAM Latency

Latency refers to the **small delay** that occurs when the memory controller begins accessing a line of data from RAM.

RAM with **lower latency** can respond **faster**, while RAM with **higher latency** takes slightly longer to deliver data.

> 💡 **NOTE:** Latency is a general concept in physics and engineering. The Wikipedia article on Latency provides a good two sentence explanation:
>
> “Latency from a general point of view is a time delay between the cause and the effect of some physical change in the system being observed… Latency is physically a consequence of the limited velocity which any physical interaction can propagate.”
>
> Remember from Class 2: CRT monitors have lower latency than LCD/LED monitors due to electrons travelling faster in a vacuum thanthrough digital circuits!

Other RAMs

SO-DIMM

- Smaller RAM sticks
- Used in laptops

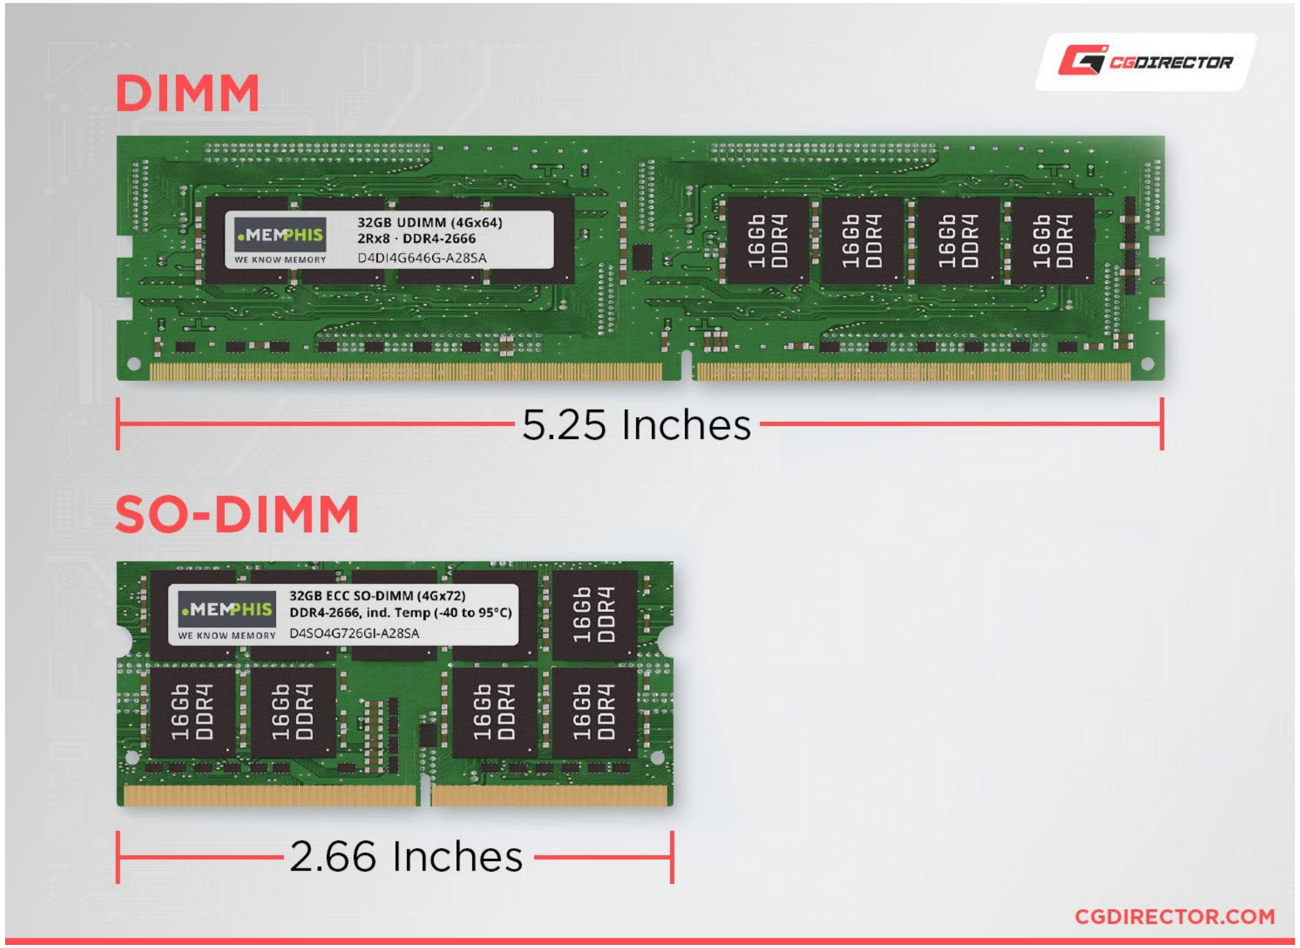

Virtual Memory

Every computer can use a portion of its **hard drive** as **temporary RAM** when the system runs out of physical memory.
This is called **virtual memory**.

- Virtual memory uses a section of the hard drive known as a **swap file** or **page file**.

- When your computer's real RAM becomes full, the operating system **moves (swaps)** some data or inactive programs from RAM to the page file.
  This **frees up space** in RAM for the programs you're currently using.

- Exam question: Why is virtual memory not ideal?

Virtual Memory: Example

Assume your computer has **4 Gb RAM.**

- And you have several programs installed on your hard drive (**A, B, C, D**).
- As programs load they take up RAM. Let's say you currently run programs **A, B and C**.
- At a certain point there is no more free RAM to run any more programs, but you want to run program **D**.

- This is where virtual memory comes in.
  - A **page file** is created on your hard drive which acts as a temporary storage box.
  - Whichever program is least used **(A, B or C**) is swapped out of RAM to the page file.
  - Let's assume this is **B**, so **B** is removed from RAM and into the page file.
  - Program **D** can now run and occupy the freed space in RAM.
- Later if you try to access **B** again,
  - it cannot run from the page file.
  - Another less-used program in RAM needs to swap over to the page file.
  - Program **B is then** swapped back into RAM and can run

Note:

- This swapping is invisible to the user, but takes time (and Windows will slow down).
- If page file is accessed too frequently, user may notice the hard drive LED going crazy.
- This is called "**disk thrashing**" and indicates the computer would benefit from more RAM.