u3a Computing Group

Alan Hopwood, 6 November 2025

Agenda

Welcome
Current News, Issues and Questions
Next Meeting
Topic: Comparison of Processor Families
Discussion

Current News, Issues and Questions

Anything to discuss?

Next Meeting

November: Convergence of TV and Computers

and

Create list of new topics

Presentation Agenda

Comparison of Processor Families

- Which processors should we cover? - the microprocessor market
- Microprocessor building blocks
- Comparisons

Uses of Microprocessor

Comparison of Processor Families

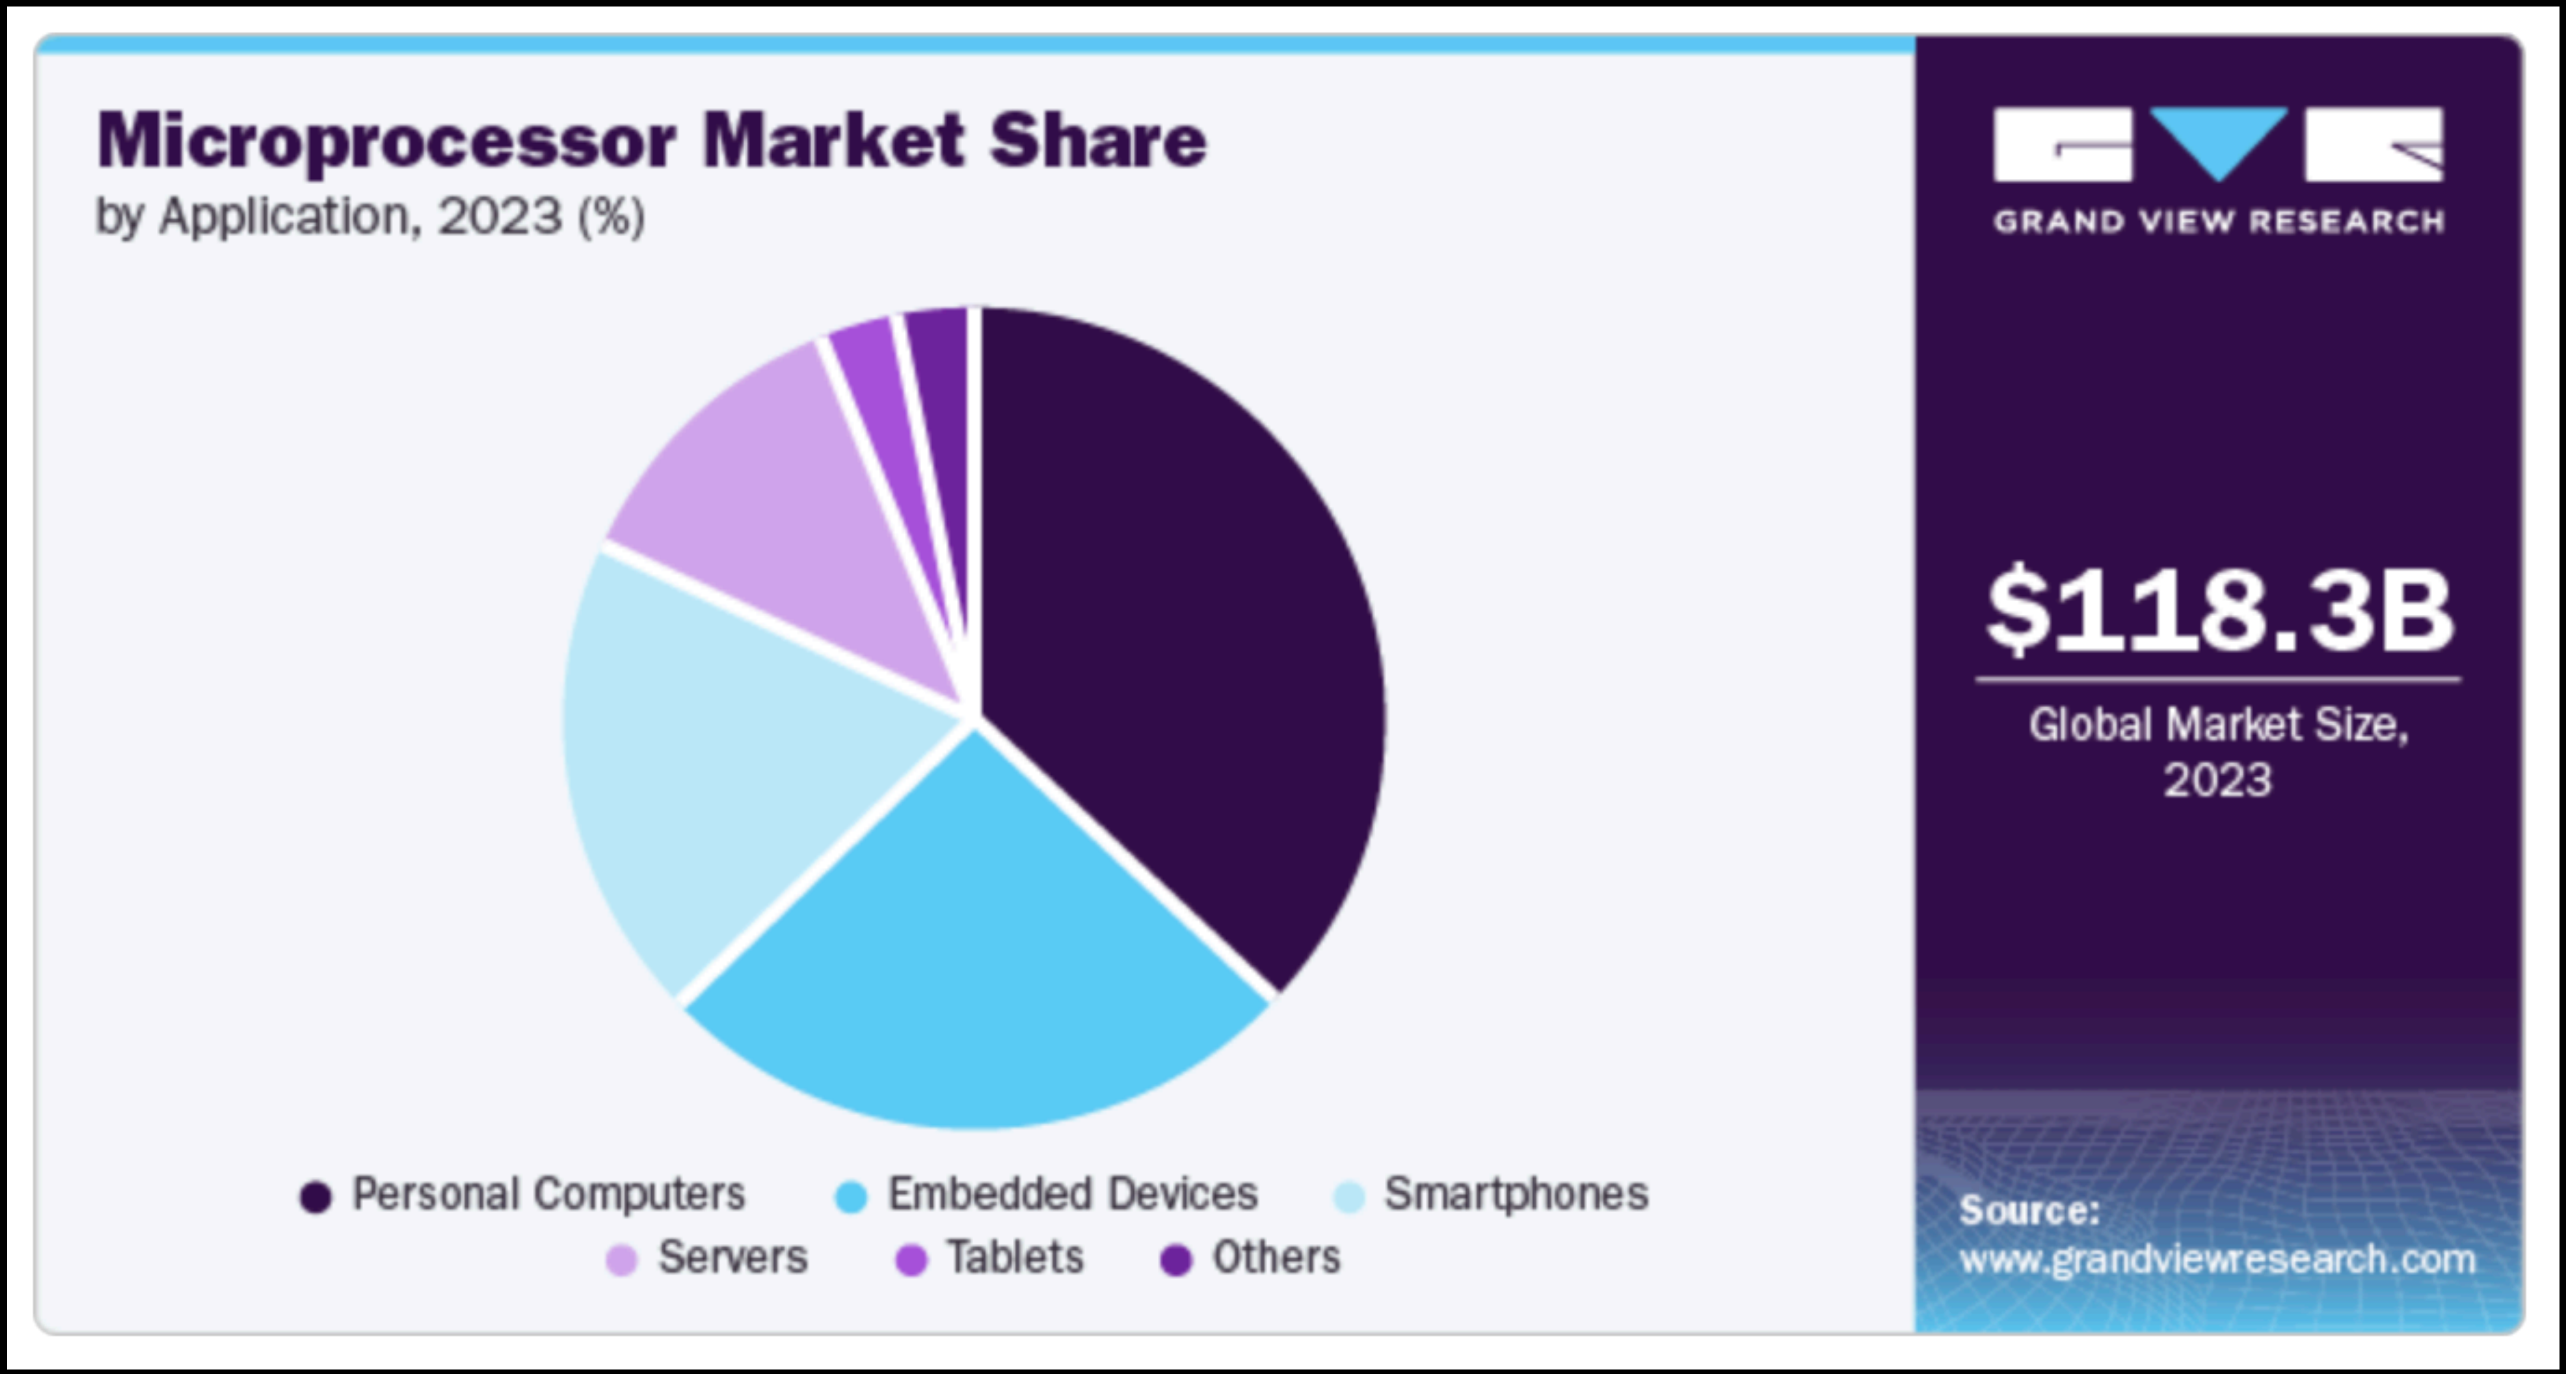

From Microprocessor Market (2024 - 2030), Grand View Research

The Microprocessor Market - Q4, 2023

- **Desktop and Laptop PCs** (x86 Market
  - **Intel** ~78% (of x86 CPUs) volume leader, but has ceded significant share over the last 5 years.
  - **AMD** ~22% (of x86 CPUs) Has peaked as high as ~35% in recent quarters. Its Ryzen CPUs are highly competitive, often leading in performance and efficiency.
  - When you include Apple's Macs with M-series chips, their share of the total client CPU market (by revenue) is estimated at around 14%
- **Mobile Devices** (Smartphones & Tablets)
  - **ARM** Architecture ~99% Virtually every smartphone processor in the world (Apple A-series, Qualcomm Snapdragon, MediaTek, Samsung Exynos, Google Tensor) is based on designs licensed from ARM Holdings.
  - **Qualcomm** Leader in Android market for high-end and mid-range chips.
  - **MediaTek** Leader in volume (units shipped) Leads in total volume by focusing on highly competitive mid-range and budget segments globally.
  - **Apple** 100% of iPhones/iPads Designs its own A-series and M-series chips using an ARM architecture license.
- **Server and Data Center CPUs**
  - **Intel** ~79% (of x86 server CPUs) The trend has been consistently downward as AMD gains traction.
  - **AMD** ~21% (of x86 server CPUs) Its EPYC processors have been a massive success.
  - **ARM**-based Growing (ex. Ampere, Amazon Graviton) While still a small single-digit percentage of the total server market, ARM-based CPUs are growing rapidly.

The Microprocessor Market

- **<u>Overall Revenue Leader</u>**: **Intel** still leads in overall x86 CPU revenue, but its historical dominance is being severely challenged.
- **<u>The Major Challenger</u>**: **AMD** has been steadily gaining market share across desktop, laptop, and the critically important server market.
- **<u>The Mobile Giant</u>**: **ARM** Holdings is the absolute dominant architecture in terms of units shipped, thanks to its use in all smartphones and tablets (via Apple, Qualcomm, MediaTek, etc.), but it's a different business model (ARM licenses its designs).
- **<u>The New Force</u>**: **Apple Silicon** (M-series chips) has completely disrupted the laptop and desktop CPU market, capturing a significant and highly profitable share.

Computer Architecture Community Favourites

Poll of the computer architecture community - 254 processor designers and engineers, researchers from academia and industry, faculty, and graduate students, identifying their favourite processors

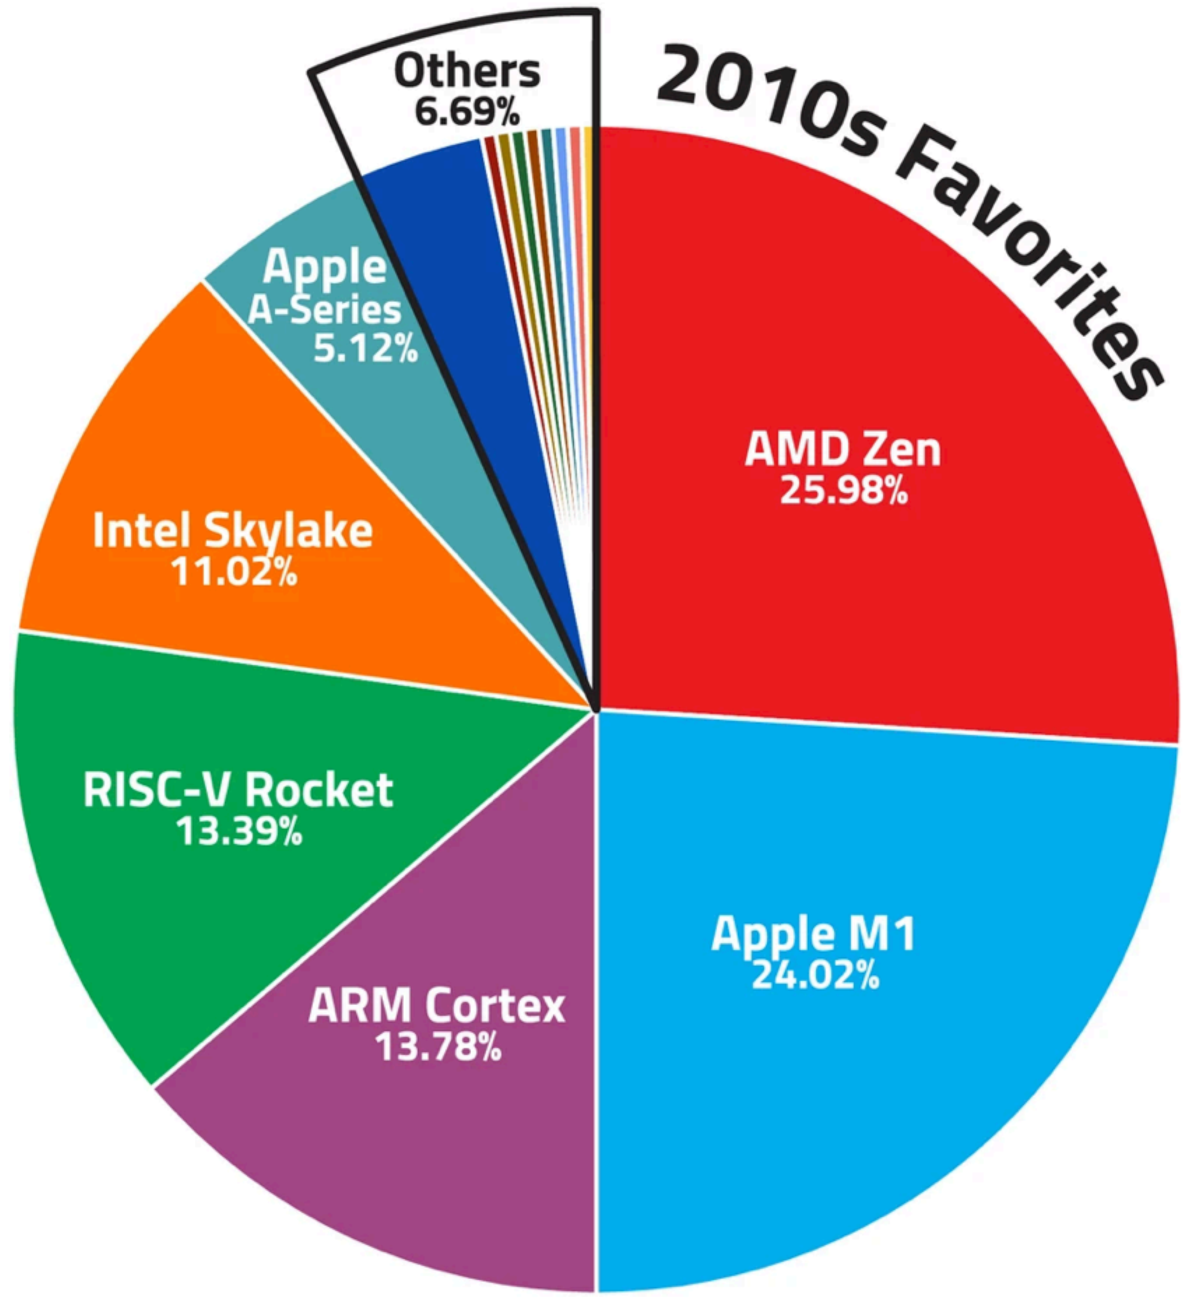

FIGURE 5. Favorite microprocessors from the 2010s.

IEEE Article

Leading Microprocessors

I have added Nvidia to the mix and will cover:

- AMD
- Intel
- Apple
- Nvidia

Will Compare:

- AMD vs Intel
- Intel vs Apple
- Apple vs Nvidia

But first, get some orientation by looking at the make up of a simple microprocessor

Processor Building Blocks

Comparison of Processor Families

- CMOS transistor
- NAND gate
- Memory Circuit
- Bus
- Clock
- Basic Processor
- Instruction Set

CMOS transistors

- Most Microprocessors are based on CMOS transistors
- Complementary Metal Oxide Semiconductor
- Wafer is silicon
- "Doping" is typically:
  - Boron (Group III element) for P (holes)
  - Phosphorus or Arsenic (Group V elements) for N (electrons)

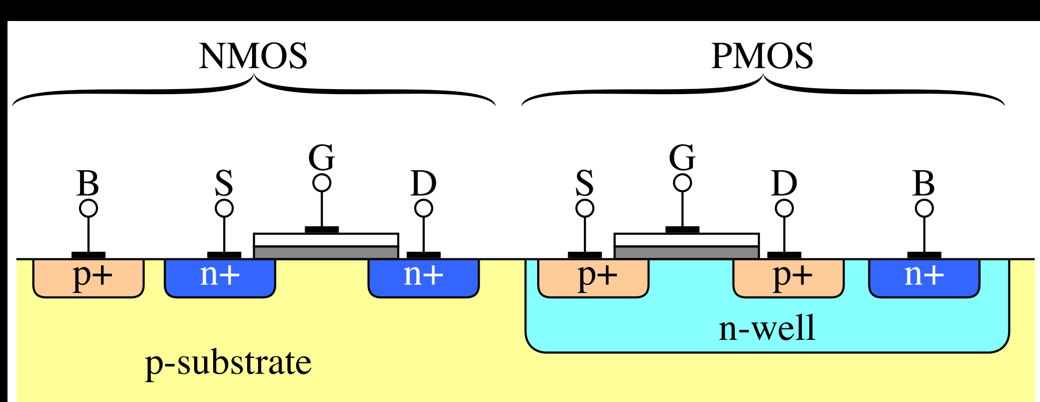

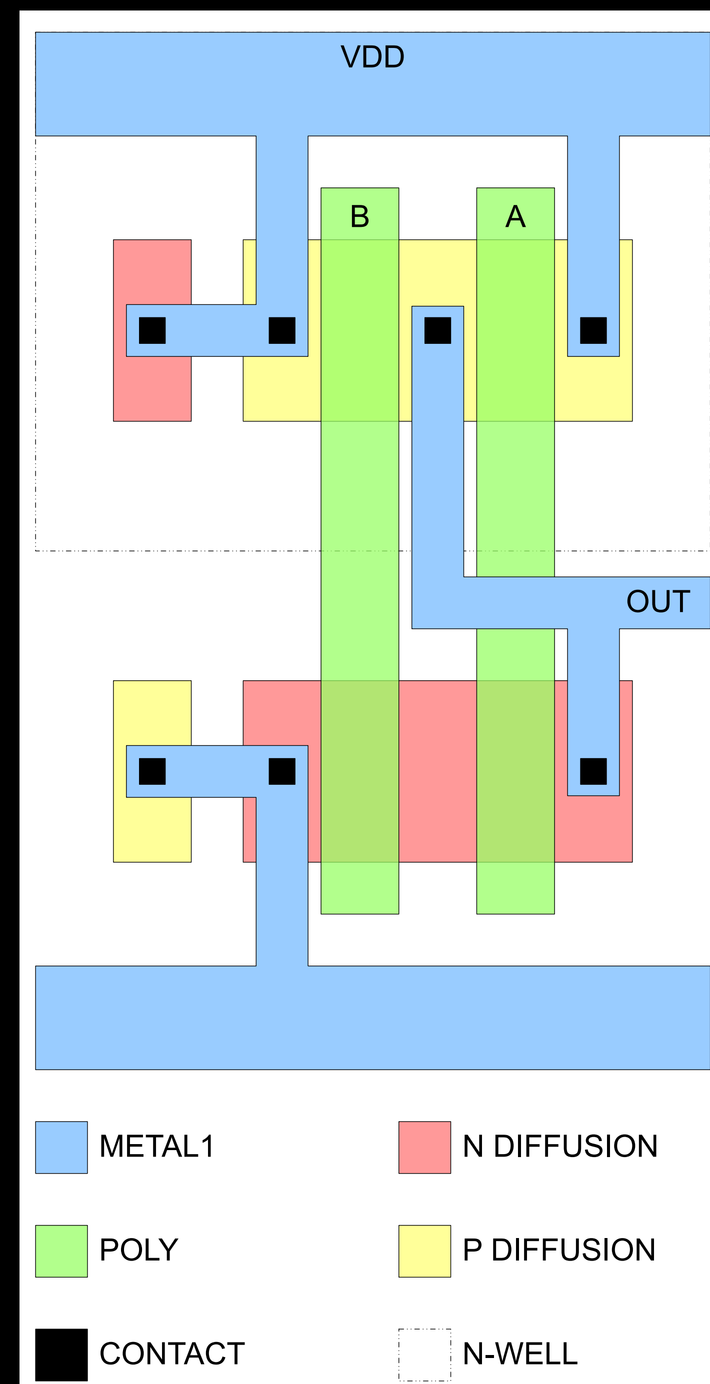

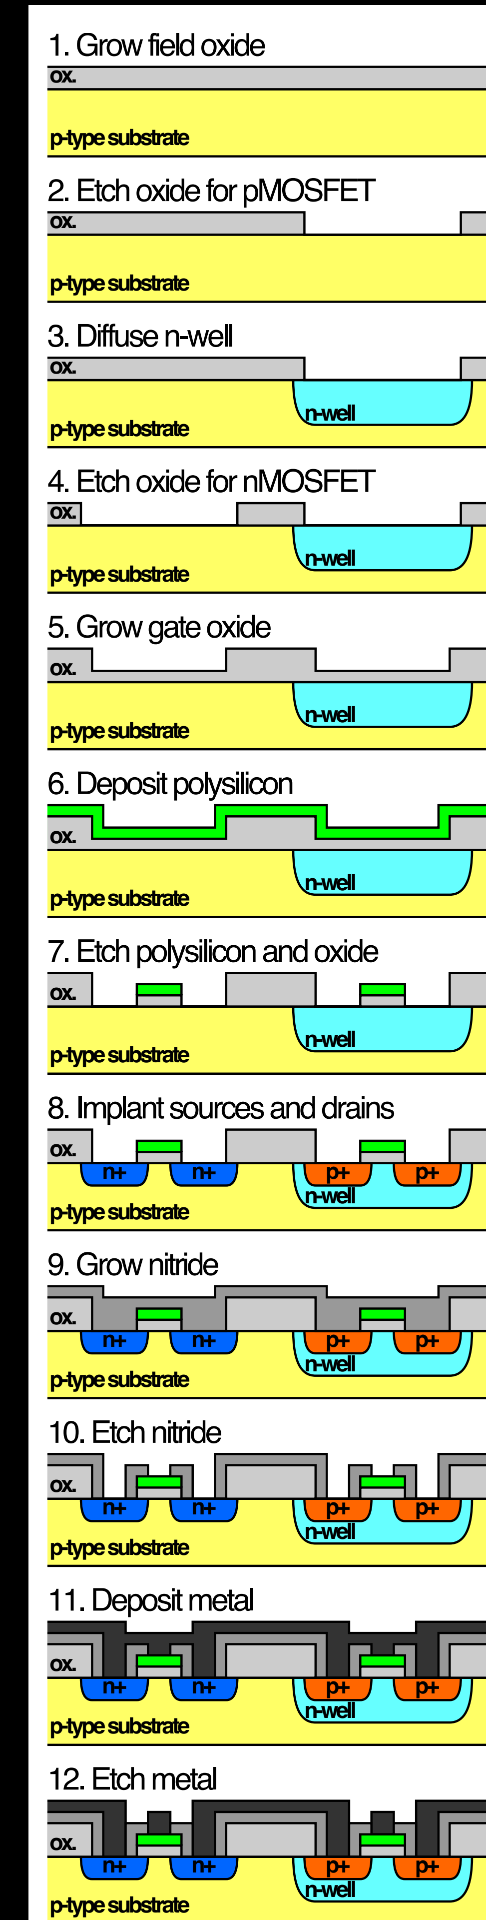

CMOS NAND Gate

Microprocessors are made up of logic gates like this NAND Gate

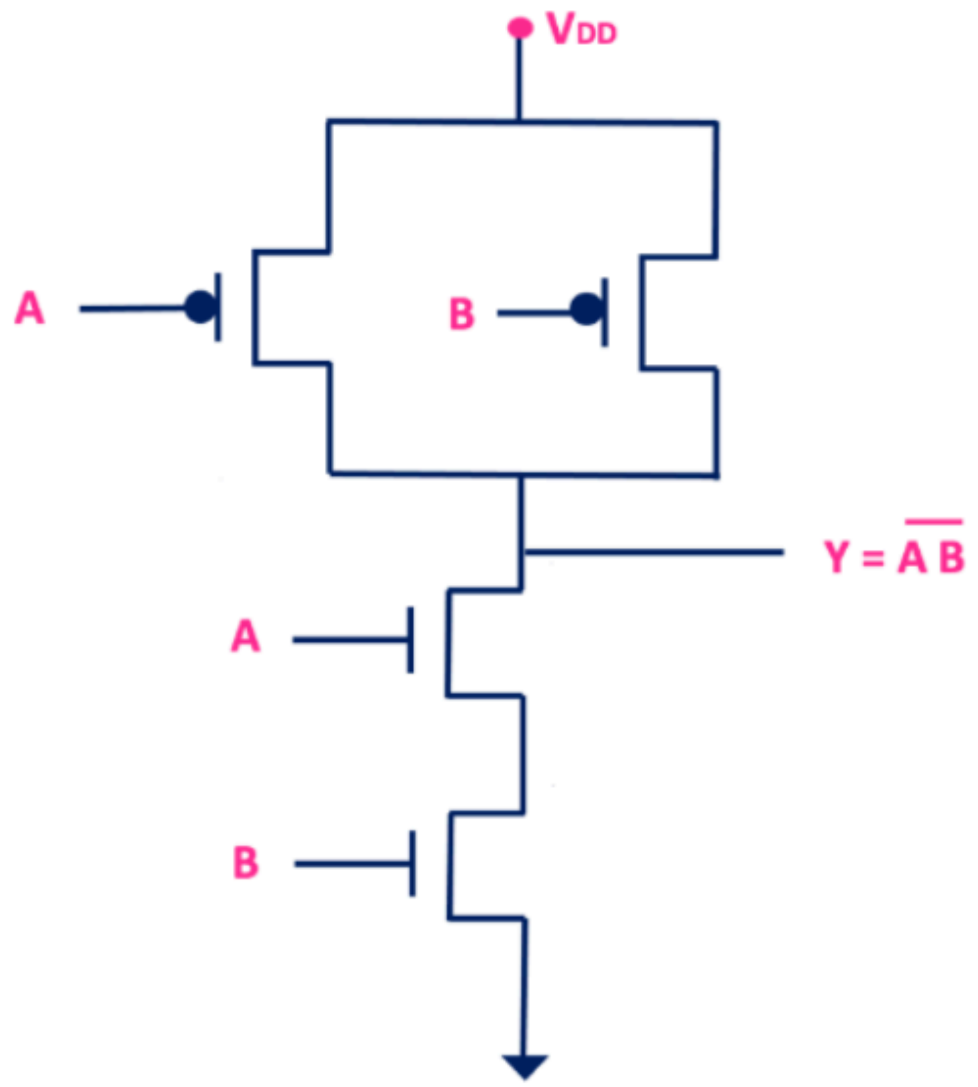

Truth Table

A	B	Y
0	0	1
0	1	1
1	0	1
1	1	0

CMOS RAM

Memory cells made in CMOS

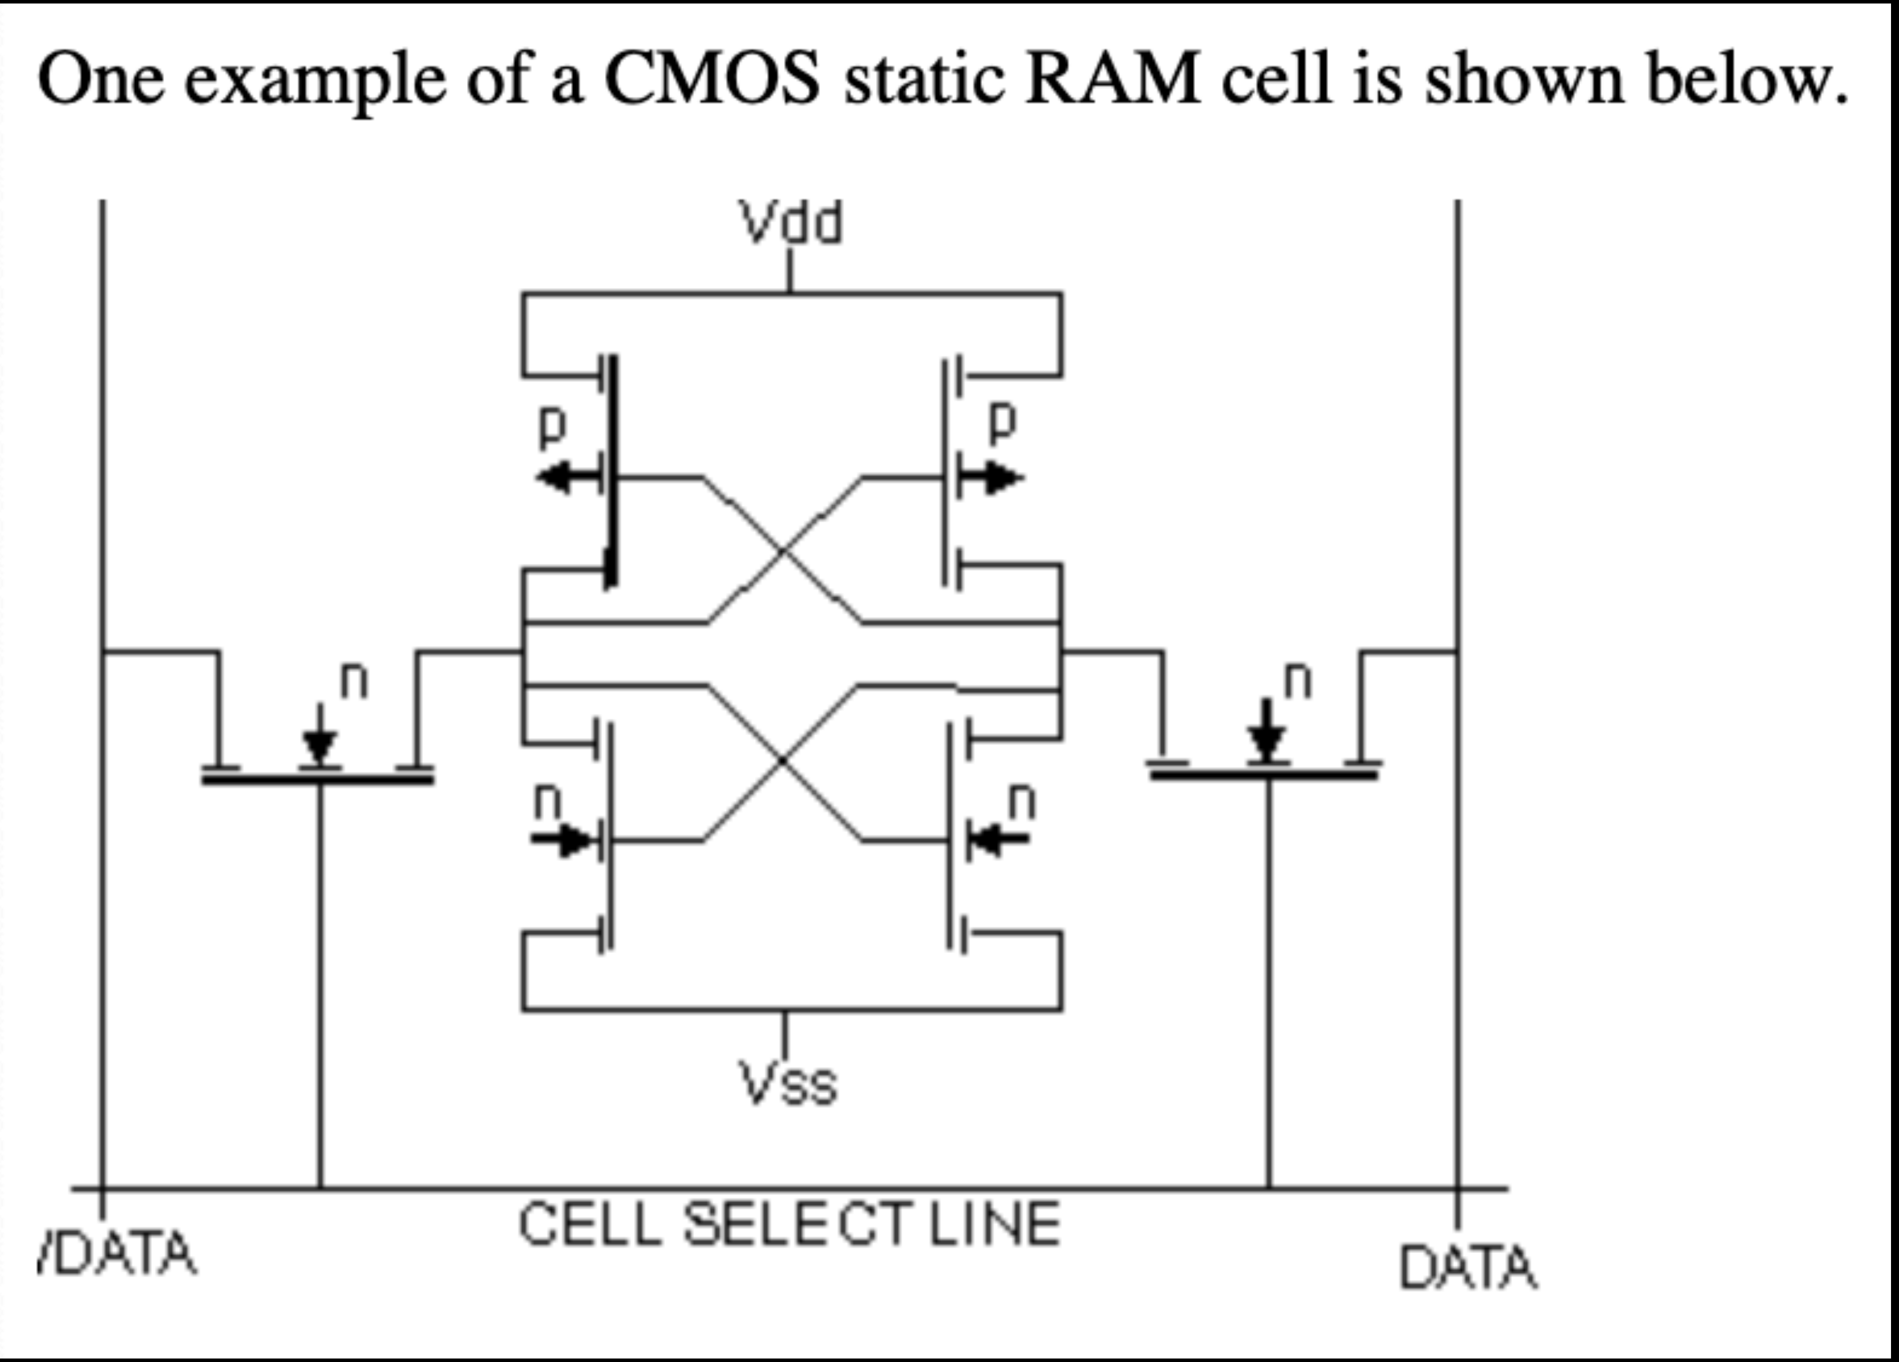

Processor Function

- Performs the same 3-step process over and over again
  - Fetch an instruction from memory
  - Decode the instruction
    - Is it an ADD, SUB, etc.?
  - Execute the instruction
    - Perform the specified operation
- This process is known as the **Instruction Cycle**

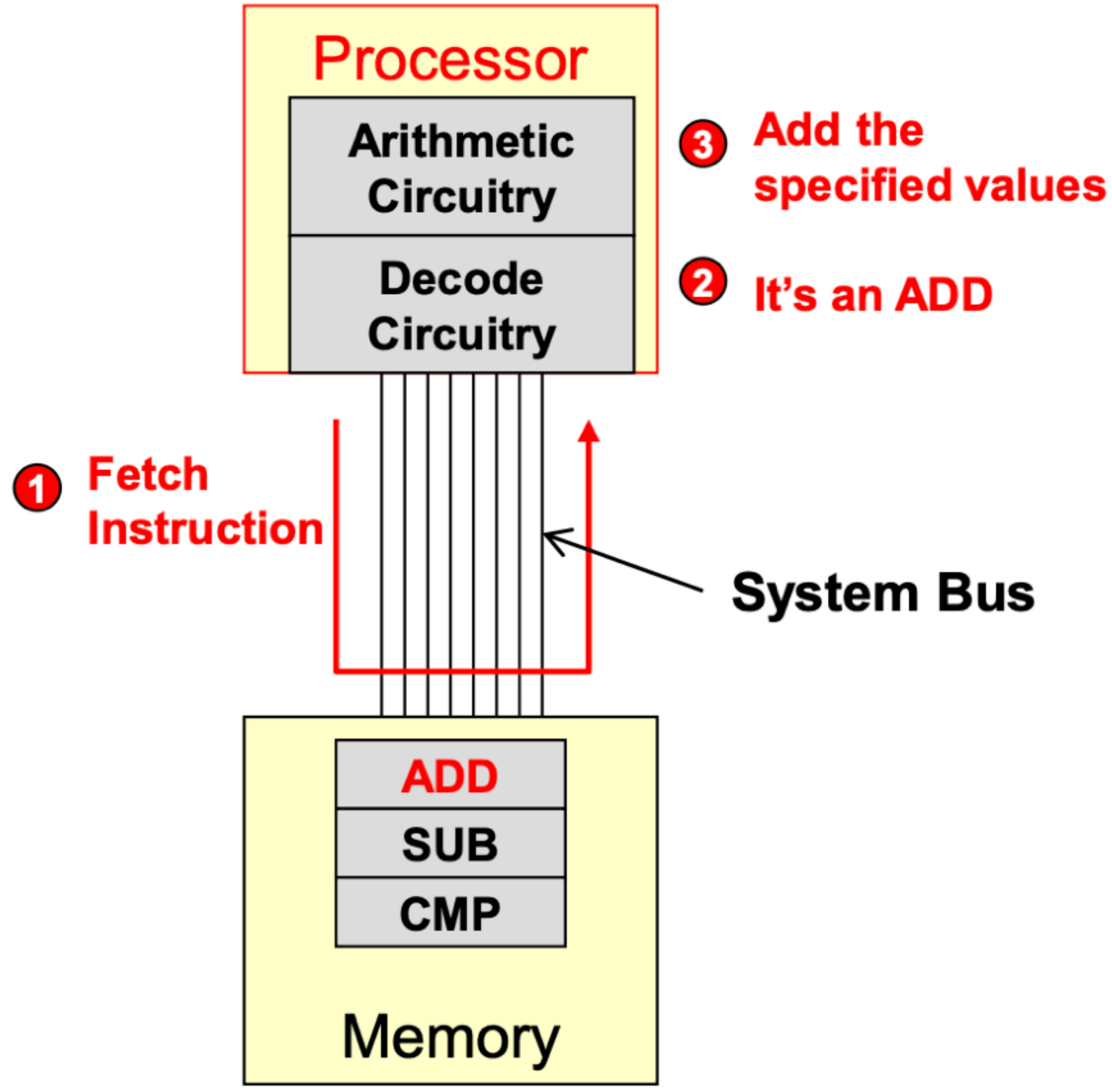

Processor Components

- 3 Primary Components inside a processor
  - ALU
  - Registers
  - Control Circuitry
- Connects to memory and I/O via **address, data**, and **control** buses (**bus** = group of wires)

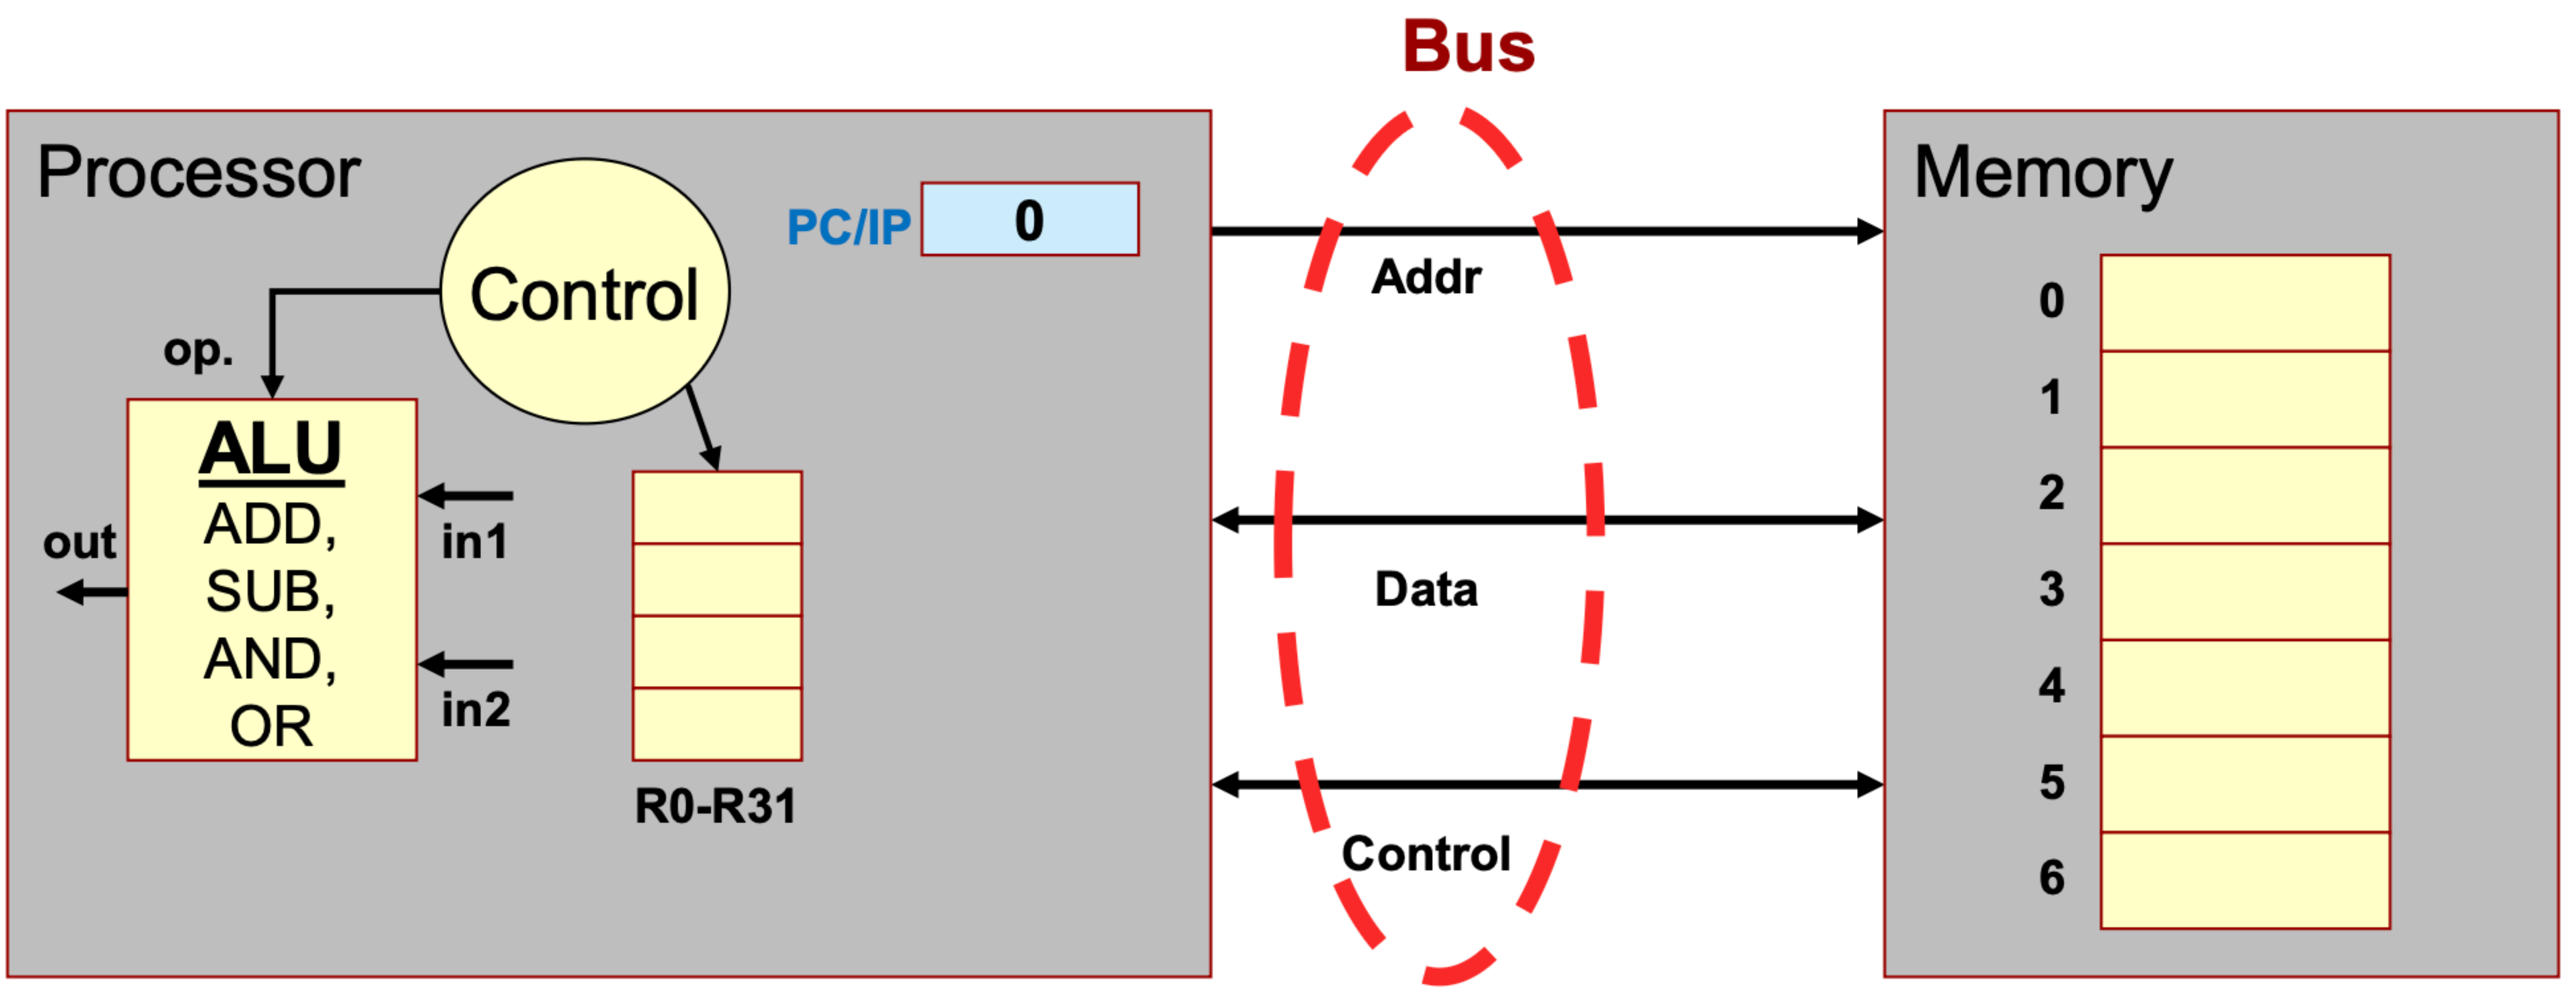

Decoding Instructions

- Control circuitry is used to decode the instruction and then generate the necessary signals to complete its execution
- Controls the ALU
- Selects Registers to be used as source and destination locations

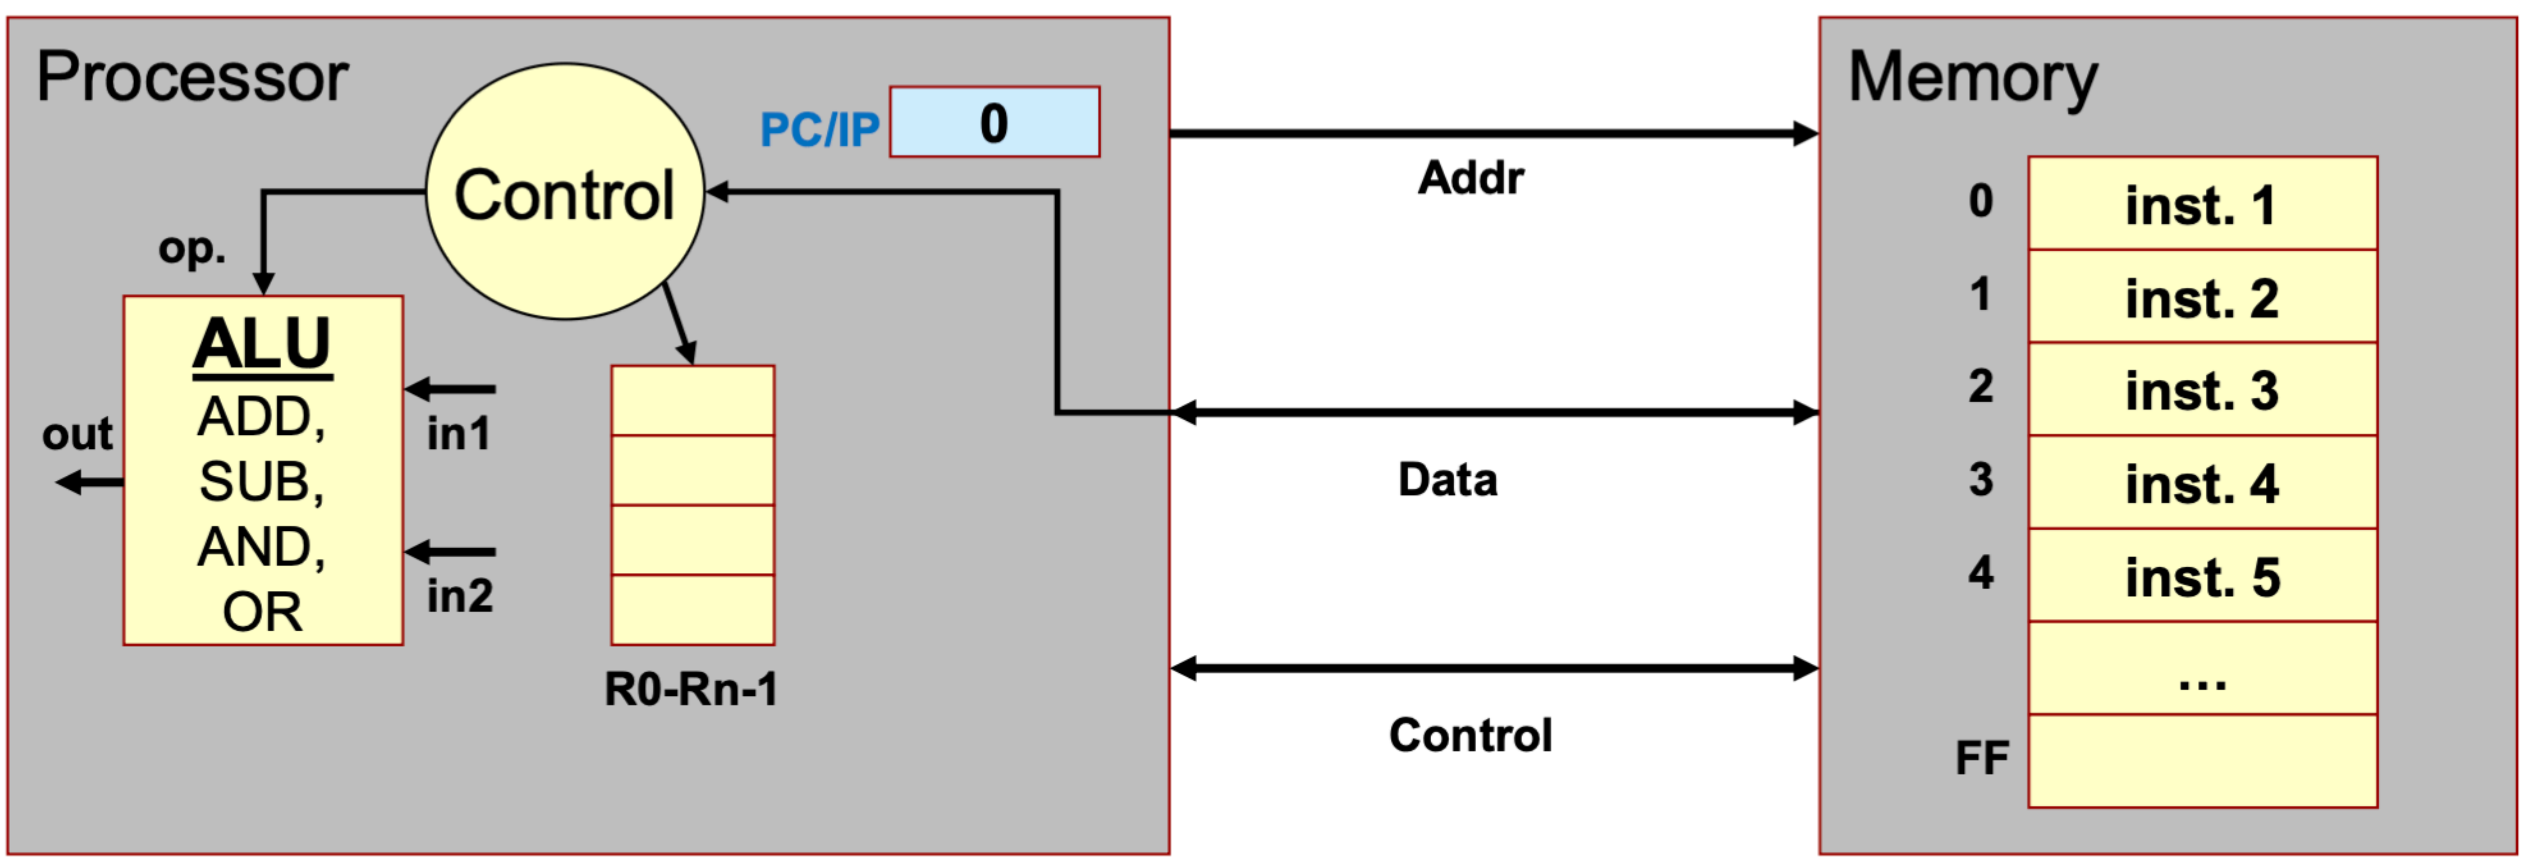

Simplified Microprocessor

Simple block diagram of 8085 microprocessor

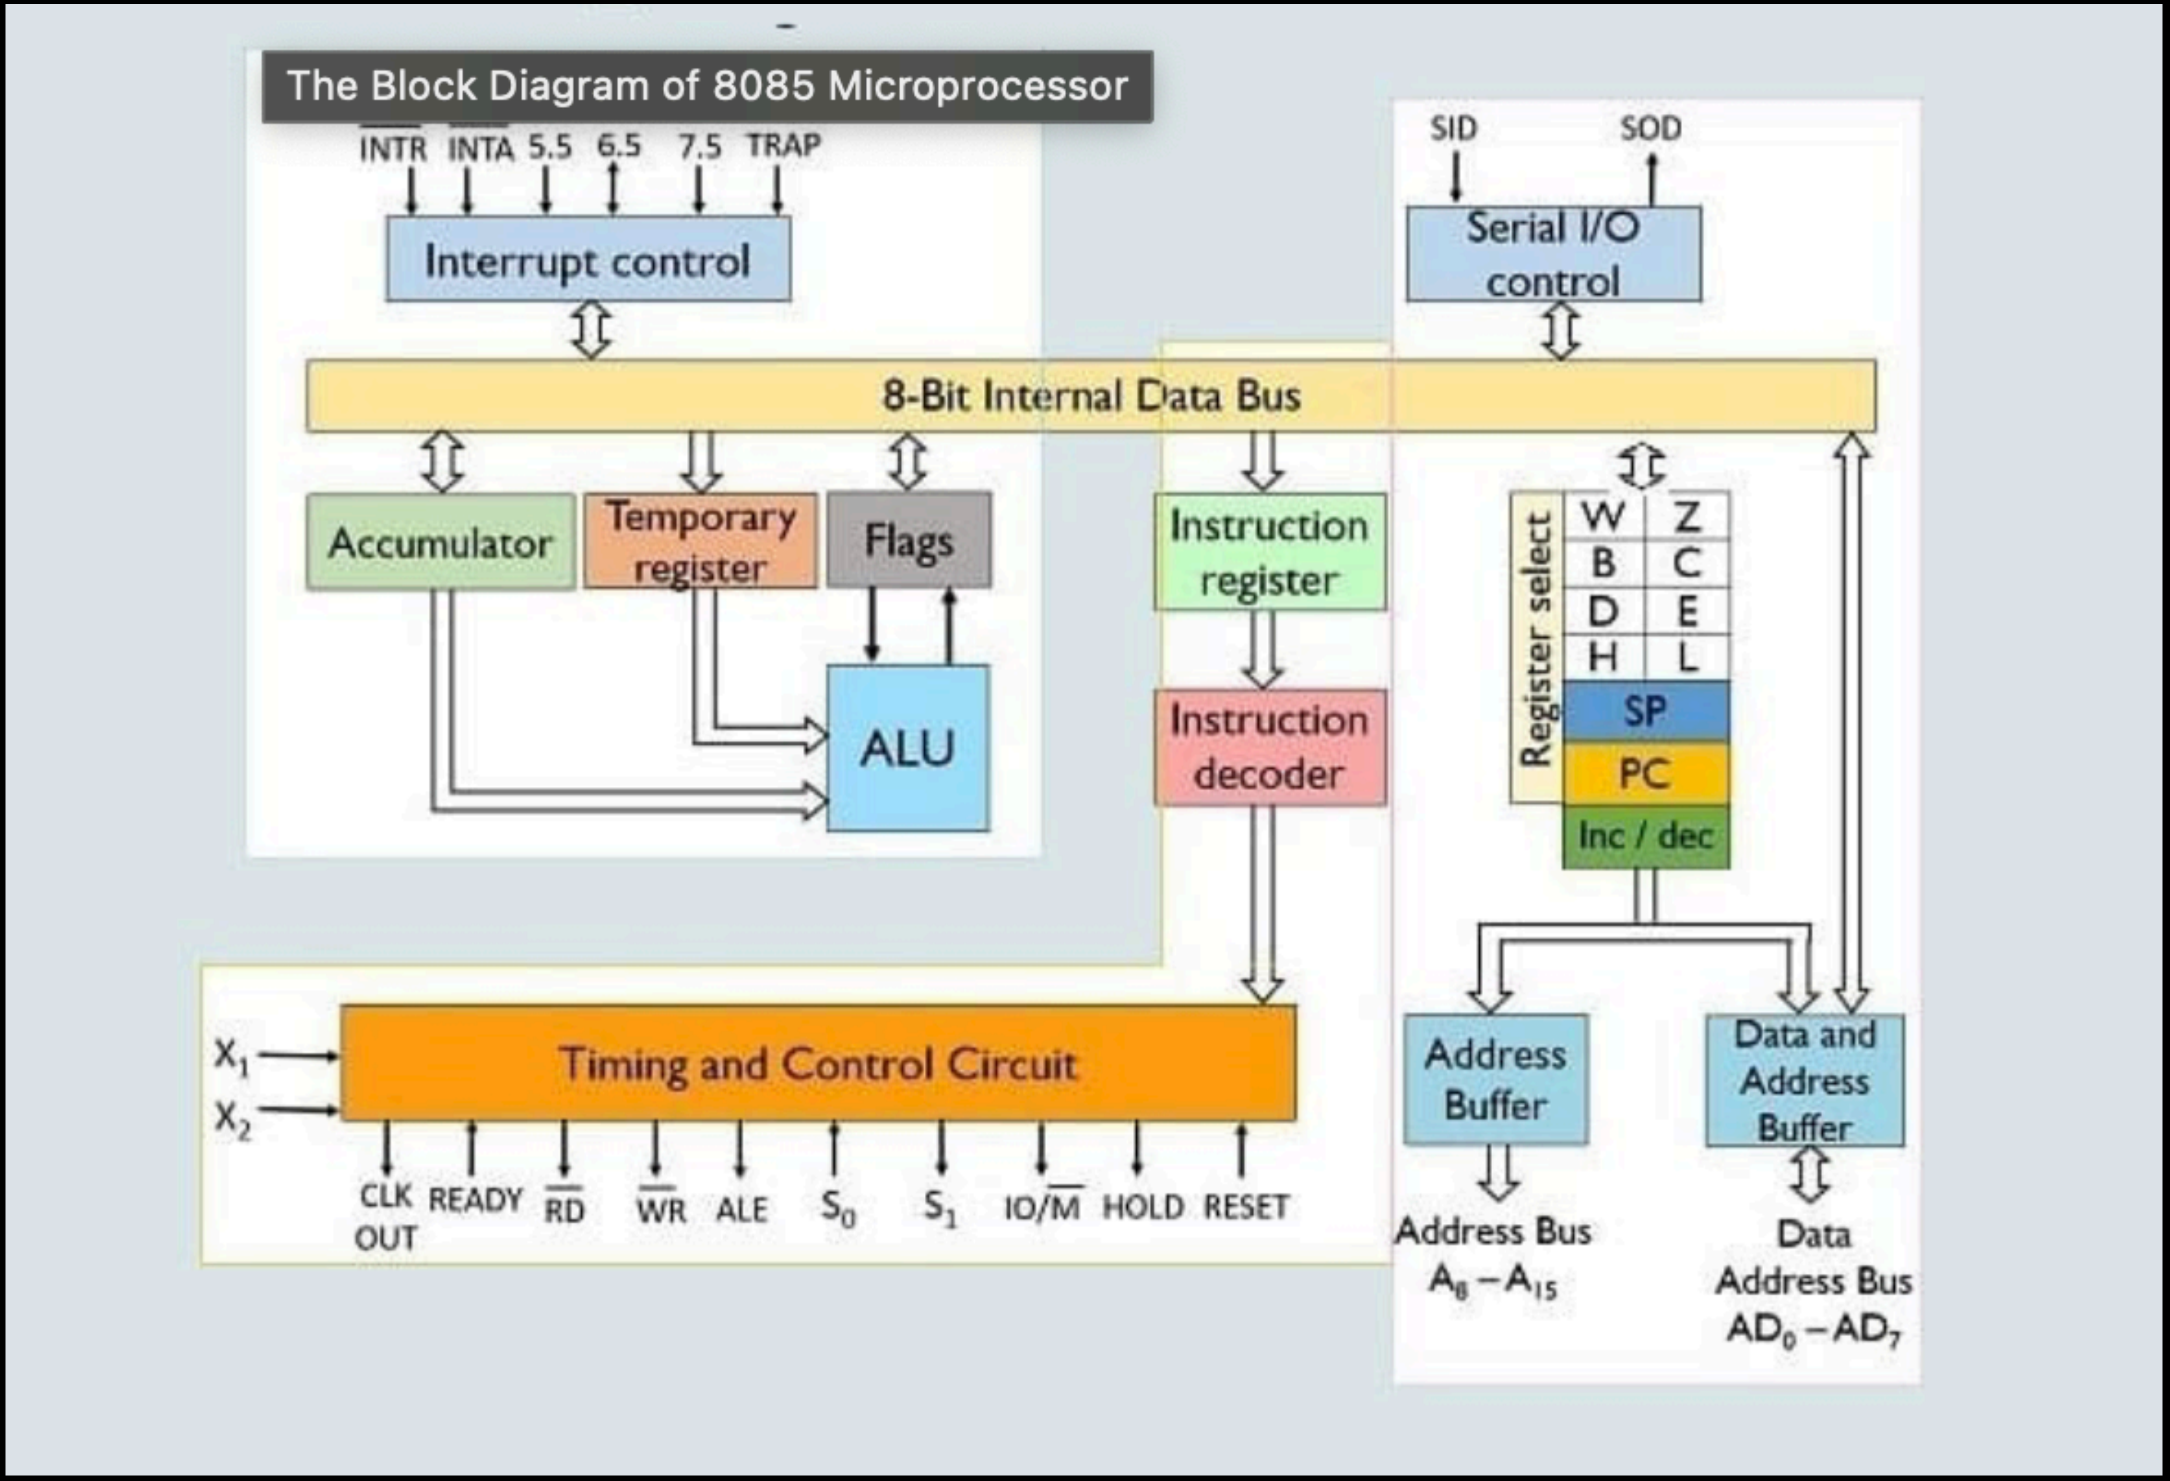

MOS 6502

Used in BBC MICRO

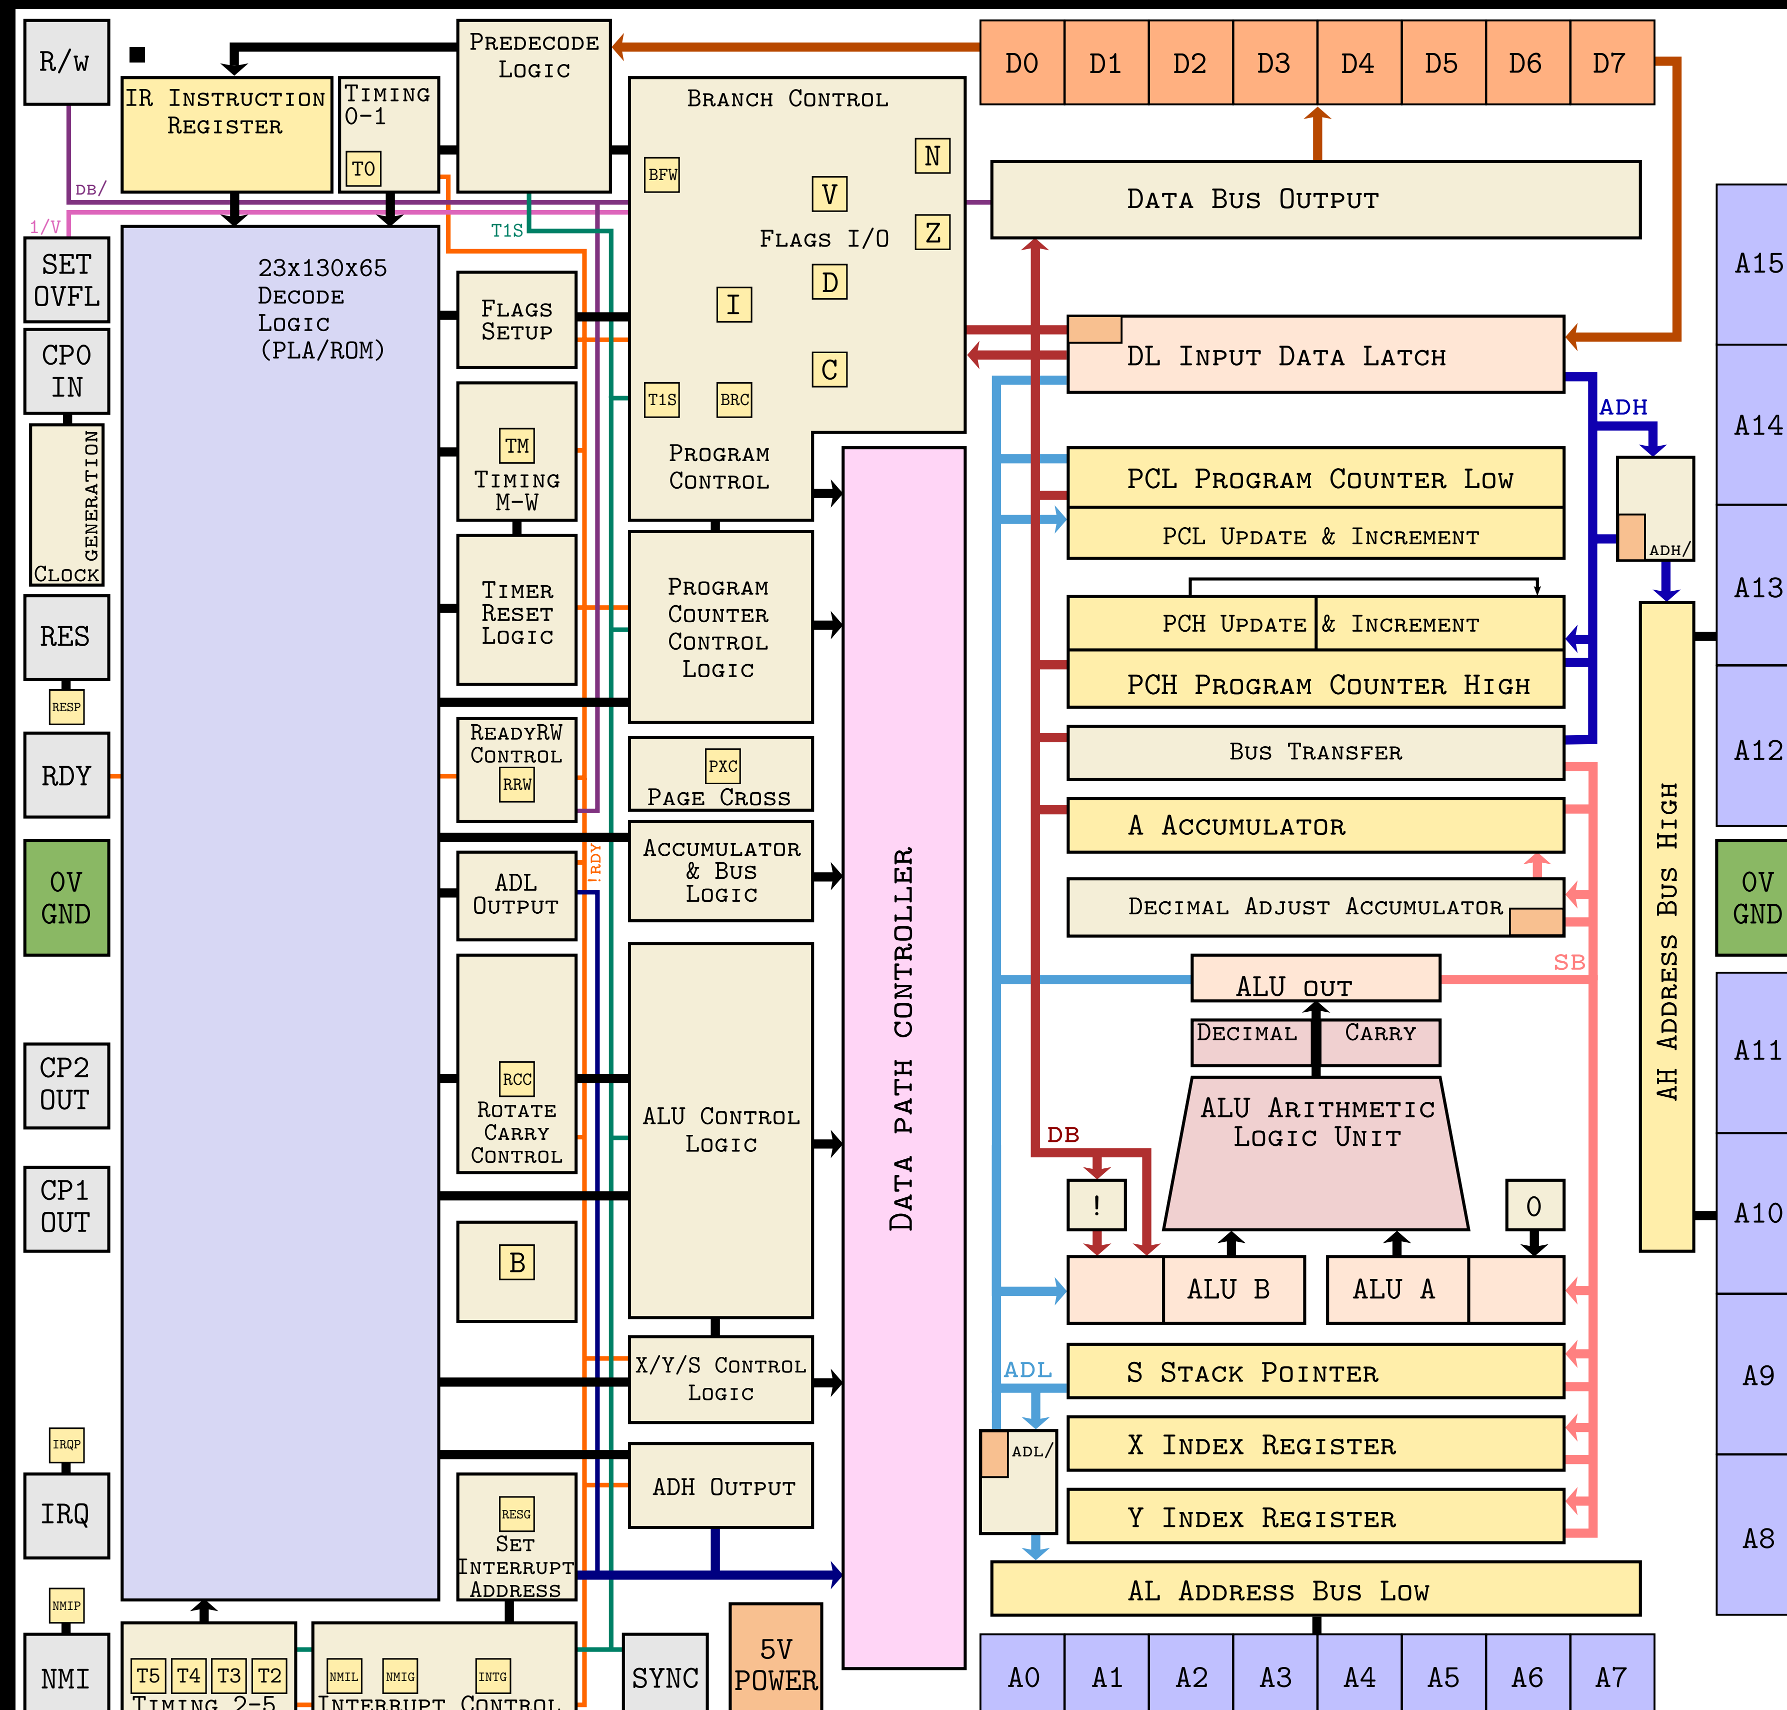

Microprocessor circuit diagram

MOS 6502

Used in BBC MICRO

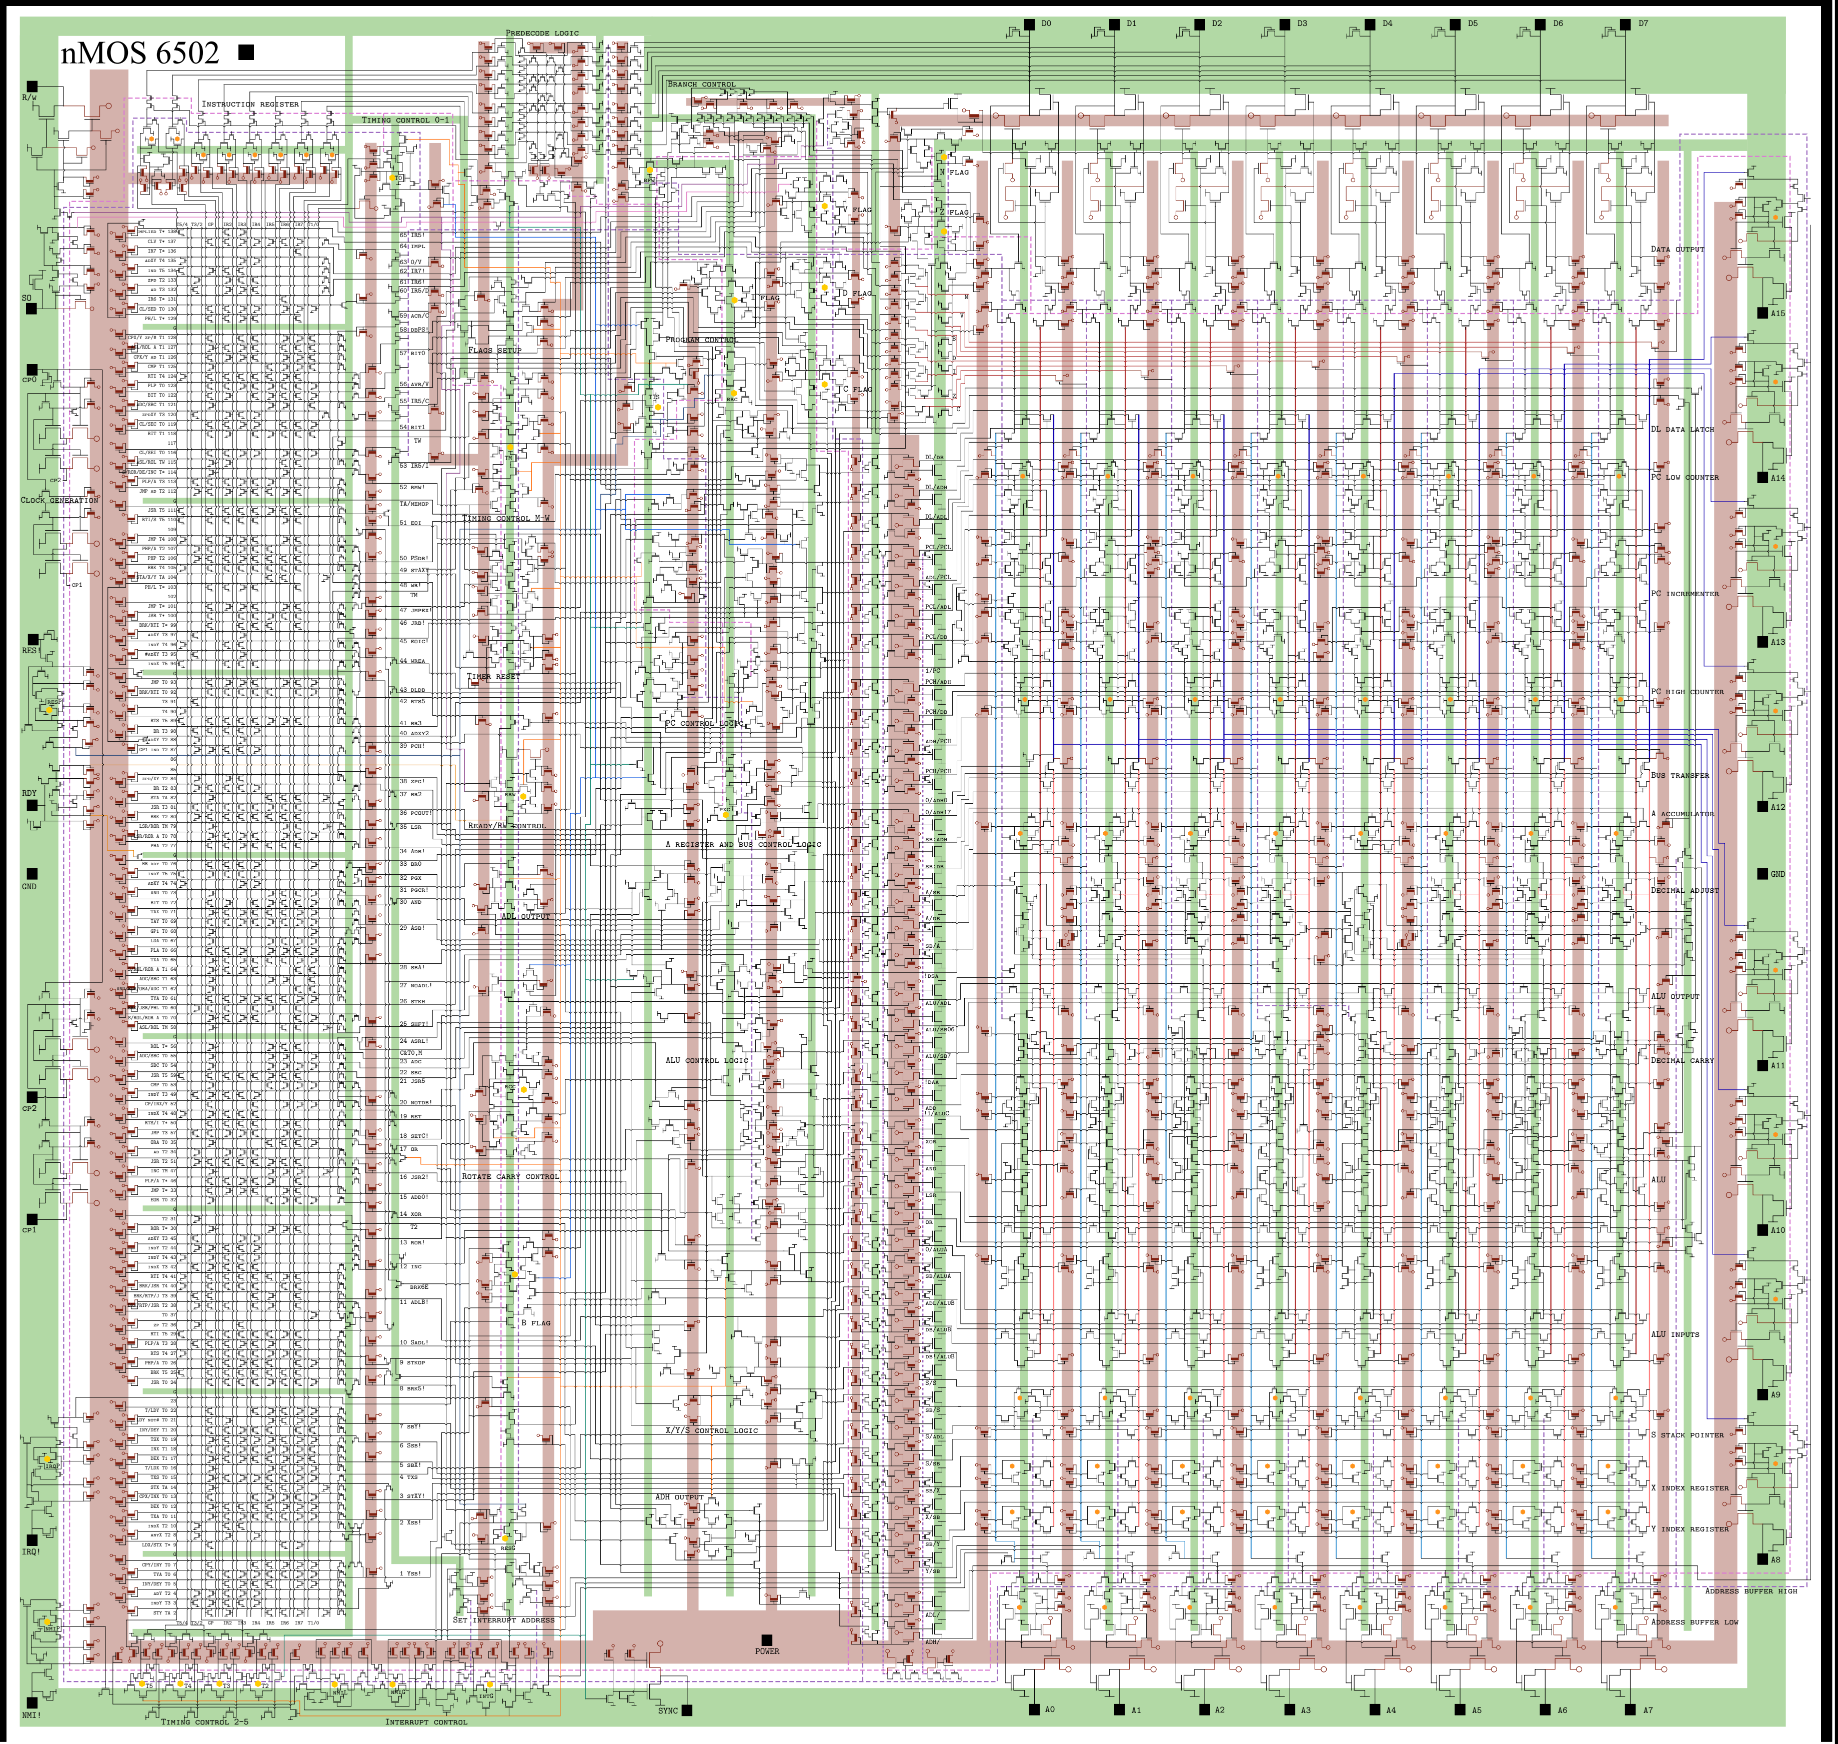

Microprocessor Instruction Set

- 3,510 transistors
- CPU clock 1 - 3 MHz
- Instruction set 56

https://www.masswerk.at/6502/6502_instruction_set.html

HI	LO-NIBBLE															
	-0	-1	-2	-3	-4	-5	-6	-7	-8	-9	-A	-B	-C	-D	-E	-F
0-	BRK impl	ORA X,ind				ORA zpg	ASL zpg		PHP impl	ORA #	ASL A			ORA abs	ASL abs	
1-	BPL rel	ORA ind,Y				ORA zpg,X	ASL zpg,X		CLC impl	ORA abs,Y				ORA abs,X	ASL abs,X	
2-	JSR abs	AND X,ind			BIT zpg	AND zpg	ROL zpg		PLP impl	AND #	ROL A		BIT abs	AND abs	ROL abs	
3-	BMI rel	AND ind,Y				AND zpg,X	ROL zpg,X		SEC impl	AND abs,Y				AND abs,X	ROL abs,X	
4-	RTI impl	EOR X,ind				EOR zpg	LSR zpg		PHA impl	EOR #	LSR A		JMP abs	EOR abs	LSR abs	
5-	BVC rel	EOR ind,Y				EOR zpg,X	LSR zpg,X		[illegible]	EOR abs,Y				EOR abs,X	LSR abs,X	
6-	RTS impl	ADC X,ind				ADC zpg	ROR zpg		[illegible]	ADC #	ROR A		JMP ind	ADC abs	ROR abs	
7-	BVS rel	ADC ind,Y				ADC zpg,X	ROR zpg,X		[illegible]	ADC abs,Y				ADC abs,X	ROR abs,X	
8-		STA X,ind			STY zpg	STA zpg	STX zpg		[illegible]		TXA impl		STY abs	STA abs	STX abs	
9-	BCC rel	STA ind,Y			STY zpg,X	STA zpg,X	STX zpg,Y		TYA impl	STA abs,Y	TXS impl			STA abs,X		
A-	LDY #	LDA X,ind	LDX #		LDY zpg	LDA zpg	LDX zpg		TAY impl	LDA #	TAX impl		LDY abs	LDA abs	LDX abs	
B-	BCS rel	LDA ind,Y			LDY zpg,X	LDA zpg,X	LDX zpg,Y		CLV impl	LDA abs,Y	TSX impl		LDY abs,X	LDA abs,X	LDX abs,Y	
C-	CPY #	CMP X,ind			CPY zpg	CMP zpg	DEC zpg		INY impl	CMP #	DEX impl		CPY abs	CMP abs	DEC abs	
D-	BNE rel	CMP ind,Y				CMP zpg,X	DEC zpg,X		CLD impl	CMP abs,Y				CMP abs,X	DEC abs,X	
E-	CPX #	SBC X,ind			CPX zpg	SBC zpg	INC zpg		INX impl	SBC #	NOP impl		CPX abs	SBC abs	INC abs	
F-	BEQ rel	SBC ind,Y				SBC zpg,X	INC zpg,X		SED impl	SBC abs,Y				SBC abs,X	INC abs,X	

$48
Sym.: PHA
Instr.: Push accumulator onto stack
Addr.: implied
Click for details...

Size and Complexity

Compare the Microprocessors being reviewed with the MOS 6502

	Transistors
MOS 6502	3,510
AMD Ryzen 9	16,650,000,000
Intel Core R9	10 - 20,000,000,000
Apple M5	28,000,000,000
Nvidia Blackwell B200 AI GPU	208,000,000,000

AMD vs Intel Architecture

- *The processor architectures of Intel and AMD both use the x86-64 instruction set, meaning they can run the same software.*
- **Intel** uses a hybrid core architecture in many recent desktop CPUs, combining Performance-cores (P-cores) for demanding tasks and Efficient-cores (E-cores) for background tasks and lower power consumption. P-cores use Hyper-Threading to handle two threads per core.
- **AMD** (Ryzen): Historically has used a more symmetrical core design where all main cores are high-performance and use Simultaneous Multithreading (SMT) (similar to Hyper-Threading) to handle two threads per core. AMD generally offers a higher core and thread count than Intel at similar price points, excelling in multi-threaded workloads.

Core = Processing Unit

AMD vs Intel Architecture

Feature	**AMD** (Ryzen - Zen Architecture)	**Intel** (Recent Generations)
Core Layout	**Symmetrical (Uniform high-performance cores)**	**Hybrid (P-cores + E-cores)**
Manufacturing	**Primarily relies on TSMC's advanced process nodes (e.g., 7nm, 5nm, 4nm).**	**Historically used proprietary nodes; more recently adopting a mix.**
Chip Design	**Uses a highly successful chiplet design (called Zen architecture) which connects separate core/cache groups, allowing for higher core counts and better yield.**	**Moving toward a chiplet-based design similar to AMD, separating components (1/O, Compute, etc.).**
Single-Thread Performance	**Very competitive, with recent generations (Zen 3 and later) significantly closing the gap.**	**Often maintains an edge due to higher Instructions Per Clock (IPC) or high clock speeds on P-cores.**
Integrated Graphics (iGPU)	**Stronger integrated graphics in their APUs (‘G' series), often suitable for entry-level or budget gaming.**	**Generally included in most CPUs (non-‘F' SKUs), often considered suitable for basic display and media.**
Overclocking	**Generally more flexible, with most Ryzen CPUs allowing overclocking on various motherboards.**	**Often restricted to specific chipsets (e.g., Z-series) and 'K' or 'KF' designated CPUs.**

- **Core** = processing unit
- **Process node** = semiconductor manufacturing process
- **Chiplet** = small modular IC with subset of processor functionality
- **Overclocking** = clocking above standard rate

Intel vs Apple Architecture

- The architectures of Intel Core processors and Apple's M-series (M1, M2, M3, M4) represent the two major opposing philosophies in modern computing—the established x86 Complex Instruction Set Computing (CISC) model and the modern, highly integrated ARM Reduced Instruction Set Computing (RISC) System on a Chip (SoC) model.
- Apple's Unified Memory (UMA) is a fundamental difference
  - Concept: The CPU, GPU, Neural Engine, and other specialised processing blocks on the M-series chip all share access to the same pool of high-bandwidth physical memory (RAM), which is physically located right next to the processor on the chip package.
  - Advantage: Eliminates the need to copy data between separate memory pools (e.g., from system RAM to discrete GPU VRAM). This dramatically reduces latency and increases efficiency, making tasks like video editing and machine learning very fast.

Intel vs Apple Architecture

Feature	Intel Core Processors (x86-64)	Apple M-series (ARM64)
Instruction Set	x86-64 (CISC): Variable length, highly complex instructions. Designed for maximum raw power and legacy compatibility.	ARM64 (RISC): Fixed-length, simpler instructions. Designed for high efficiency and rapid execution.
Design Type	CPU (Central Processing Unit): Primarily focuses on the CPU with separate chips often used for GPU, RAM, etc.	SoC (System on a Chip): Integrates the CPU, GPU, RAM, Neural Engine, and 1/0 controllers onto a single chip.
Core Architecture	Hybrid Cores: Combines Performance-cores (P-cores) and Efficient-cores (E-cores) on a single die.	Hybrid Cores (Symmetrical): Combines large, highly powerful Performance Cores and highly efficient Efficiency Cores.
Manufacturing	Primarily uses its own fabrication nodes (though increasingly uses external foundries).	Exclusively uses TSMC's cutting-edge, high-density process nodes (e.g., 5nm, 3nm).
Ecosystem	Open/Modular: Designed to work with various motherboard chipsets, operating systems (Windows, Linux, macOS via emulation), and third-party components (like discrete GPUs and standard RAM).	Closed/Integrated: Designed specifically for the macos/iOs ecosystem, allowing for deep, proprietary hardware/software optimization.

Apple vs Nvidia Architecture

- The comparison between a NVIDIA GPU and an Apple M-series processor is primarily a comparison between a Dedicated Graphics Processing Unit (dGPU) and a highly advanced Integrated Graphics Processor (iGPU) within a System on a Chip (SoC). They are designed for different priorities and operate in fundamentally different architectural environments.

Apple vs Nvidia Architecture

Feature	Apple M series	NVIDIA GPU
Design Type	Integrated SoC: The GPU is built into the same physical chip package as the CPU and Neural Engine.	Dedicated/Discrete GPU: A standalone component, separate from the CPU.
Memory Architecture	Unified Memory Architecture (UMA): The GPU shares the same pool of high-speed system RAM (LPDDR) with the CPU.	Dedicated VRAM: Uses its own large pool of high-speed memory (GDDR), entirely separate from the system's main RAM.
Power & Thermal	High Efficiency: Designed for performance-per-watt (often 40W–80W under heavy load), prioritizing battery life and thermal limits.	High Power/Heat: Designed for maximum power draw (often 200W–450W+) to achieve peak performance.
Performance Metric	Low Latency & Efficiency: Focused on minimizing the time it takes for data to move between the CPU and GPU.	Raw Throughput: Focused on maximum TFLOPS and memory bandwidth (VRAM) for heavy, sustained workloads.
Compute Ecosystem	Metal: Uses Apple's proprietary Metal API, optimized specifically for Apple hardware and its Unified Memory.	CUDA: Relies on NVIDIA's proprietary CUDA platform, which is the industry standard for AI, machine learning, and high-performance computing (HPC).

Summary

AMD Ryzen (High-End R9):

- **Strength**: Often delivers superior multi-threaded performance and an excellent price-to-performance ratio (especially with 3D V-Cache models).
- **Weakness**: Single-threaded performance can sometimes lag behind Intel's high-end offerings.

Apple M-Series:

- **Strength**: Exceptional power efficiency (performance per watt) and highly optimised, integrated performance for macOS ecosystem. Leads in tasks like video editing.
- **Weakness**: Limited upgradeability/customisation (integrated architecture) and potential software compatibility issues for non-macOS applications.

Intel Core i9:

- **Strength**: Offers class-leading single-threaded performance due to high clock speeds. Excellent for competitive gaming and workloads dependent on a few fast threads. Broad compatibility and robust features.
- **Weakness**: Typically has higher power consumption and thermal output under heavy load compared to competing desktop CPUs.

NVIDIA Blackwell B200 (AI/HPC GPU):

- **Strength**: Unprecedented AI/HPC performance with massive parallel processing power, huge memory (192 GB HBM3e), high memory bandwidth (8 TB/s), and advanced Tensor Core architecture (FP4/FP6/FP8 support).
- **Weakness**: Highly specialised and expensive; it is a dedicated accelerator (GPU) designed for data centres and is not a general-purpose desktop CPU. The full software stack may take time to mature.

Thank You