

Chapter II: Processor Architectures

II Processor Architectures

The architecture of a processor refers to the internal organization of its components and the way they interact to execute instructions. It defines the relationship between the processor, memory, and peripherals, and directly influences the speed, energy efficiency, and programming complexity of an embedded system. Two major families dominate modern designs: the Von Neumann architecture and the Harvard architecture.

II.1 Neumann Architecture

The Von Neumann architecture is based on a single memory space that stores both the program instructions and the data being processed. The processor communicates with this memory through a single bus, meaning that instructions and data share the same communication path. This organization offers the advantages of simplicity and low implementation cost. It is particularly suitable for systems requiring high software flexibility, such as personal computers and workstations.

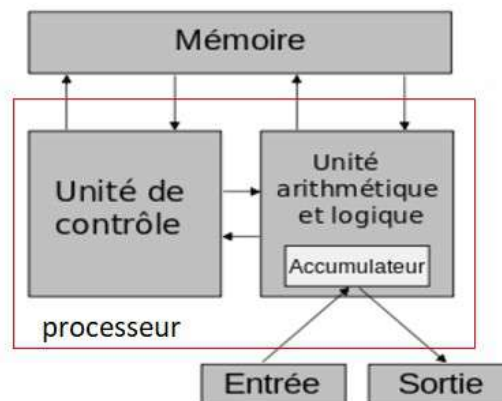


Figure II. 1 . Von Neumann Architecture

However, this architecture is constrained by a significant limitation known as the **Von Neumann bottleneck**, as the processor can access only one piece of information at a time—either an instruction or data—which slows down execution speed.

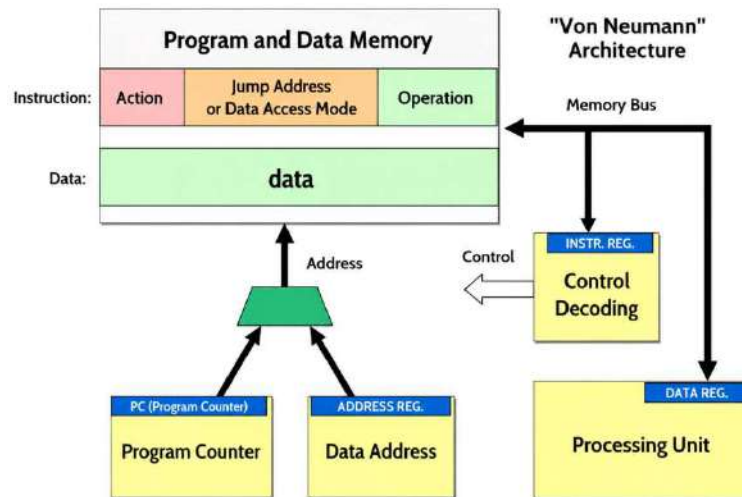


Figure II. 2 - Instruction Execution in a Von Neumann Architecture

In this architecture, the execution of a program follows the classical instruction cycle, often referred to as the **Von Neumann cycle**, which includes several successive stages:

- **Fetch (Instruction Read):** The processor reads the instruction to be executed from memory. Since there is only one bus, it can fetch only one instruction at a time.
- **Decode:** The instruction is analyzed by the control unit to identify the operation to be performed and the resources required.
- **Fetch Operand (Data Read):** If the instruction requires data, these are retrieved from memory via the same bus. This may introduce additional latency, as instructions and data share the same access path.
- **Execute:** The arithmetic and logic unit (ALU) or another component of the processor executes the requested operation.
- **Write Back:** The result is written back into memory or a register, again using the single bus.

This cycle clearly illustrates the main limitation of the Von Neumann architecture: since instructions and data cannot be transferred simultaneously, each stage must wait for bus availability. This results in reduced performance, especially in applications that require high data throughput.

II.2 Harvard Architecture

The **Harvard architecture** introduces a physical separation between the **instruction memory** and the **data memory**. The processor therefore uses **two distinct buses** — one dedicated to instructions and another to data. This organization enables the processor to fetch an instruction while simultaneously accessing data, which greatly increases execution speed and enhances parallelism.

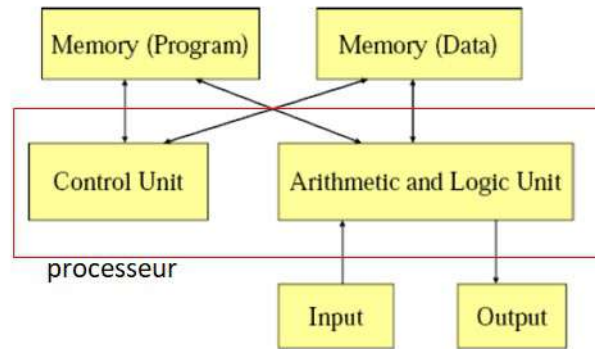


Figure II. 3 - Harvard Architecture

This structure is widely adopted in modern **microcontrollers** and **embedded systems**, where performance, energy efficiency, and responsiveness are essential. It also allows for better exploitation of **pipelining**, a technique that divides instruction execution into parallel stages. However, this design is more complex and slightly more expensive to implement than the Von Neumann architecture.

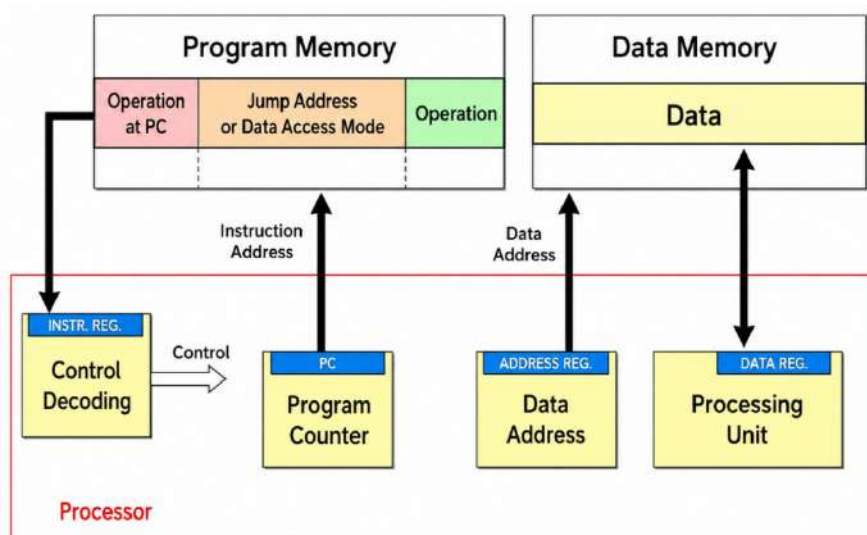


Figure II. 4 - Instruction Execution in a Harvard Architecture

The instruction cycle in a Harvard architecture is optimized by the presence of two separate memories and two buses, allowing faster and more efficient execution since multiple operations can occur in parallel:

- **Fetch (Instruction Read):** The processor retrieves the instruction from the program memory through the instruction bus.
- **Decode:** The instruction is interpreted by the control unit to determine the operation and required resources.
- **Fetch Operand (Data Read):** Simultaneously, using the separate data bus, the processor accesses the necessary data without waiting for the instruction read to complete.
- **Execute:** The arithmetic logic unit (ALU) or another processing unit performs the operation.

Write Back: The result is stored in data memory or a register using the dedicated data bus, while a new instruction can already be fetched.

This simultaneous operation significantly reduces idle time and enables pipeline organization, where several instructions are processed at different stages of execution. For instance, while one instruction is being executed, another can be decoded and a third fetched from memory. This approach eliminates the Von Neumann bottleneck and provides a substantial improvement in both processing speed and overall system throughput.

II.3 Comparison Between Von Neumann and Harvard Architectures

The Von Neumann architecture is characterized by its simplicity of design and flexibility of use. It relies on a single shared memory and one common bus for both instructions and data. This organization reduces hardware costs and simplifies programming, making it suitable for general-purpose computing. However, it suffers from a major limitation known as the Von Neumann bottleneck, which results from the processor's inability to access instructions and data simultaneously. As a consequence, execution speed can be significantly reduced, especially in applications requiring high-speed or intensive data processing.

In contrast, the Harvard architecture introduces a physical separation between instruction memory and data memory, each with its own dedicated bus. This design enables parallel

access to both instructions and data, eliminating the bottleneck and providing a substantial increase in processing speed. Additionally, it improves energy efficiency, making it particularly well-suited for embedded systems, real-time applications, and digital signal processing (DSP), where performance and resource optimization are critical.

Table II. 1- Comparison Between Von Neumann and Harvard Architectures

Characteristic	Von Neumann Architecture	Harvard Architecture
Memory Organization	Single memory for both instructions and data	Separate memories for instructions and data
Bus Configuration	Single shared bus	Distinct buses for instructions and data
Access to Instructions/Data	Sequential (one access at a time)	Parallel (simultaneous access possible)
Design Complexity	Simpler, lower cost	More complex, slightly more expensive
Execution Speed	Slower due to bottleneck	Faster, eliminates the bottleneck
Flexibility	High (easier programming and design)	More specialized
Typical Applications	General-purpose computers and systems	Embedded systems, DSPs, real-time applications
Energy Consumption	Higher under intensive workloads	More energy-efficient for targeted processing

II.4 Choosing Between Von Neumann and Harvard Architectures

The choice between **Von Neumann** and **Harvard** architectures primarily depends on the project constraints, performance requirements, and available resources. Each architecture offers specific advantages and limitations that directly affect the design and implementation of an embedded system.

II.4.1 Selection Criteria

The main criteria to consider when selecting an architecture include:

- **System Complexity:** The Von Neumann architecture is simpler to design and program, while the Harvard model requires more rigorous management of the separation between instructions and data.

- **Expected Performance:** For applications demanding high-speed or parallel processing, the Harvard architecture is more suitable. For general-purpose systems, Von Neumann remains sufficient.
- **Power Consumption:** Harvard is generally more energy-efficient, particularly in real-time and DSP systems, whereas Von Neumann tends to consume more power during intensive operations.
- **Cost and Accessibility:** Von Neumann architectures are more economical in terms of hardware design, while Harvard requires additional resources such as separate memory blocks and data buses.
- **Ease of Development:** Von Neumann simplifies programming and system integration, making it an ideal choice for prototypes, educational systems, and entry-level embedded projects.

II.4.2 Specific Applications for Each Architecture

The decision between Von Neumann and Harvard architectures is not merely theoretical; it directly influences the target application domains where each model proves most effective. The fundamental difference — a shared memory and bus in Von Neumann versus separate instruction and data paths in Harvard — naturally determines their areas of use. Some applications favor the flexibility and simplicity of Von Neumann, while others require the speed and energy efficiency of Harvard.

A) Von Neumann Architecture:

- **Personal and Laptop Computers:** Where application diversity (office software, multimedia, programming) demands flexibility.
- **General-Purpose Operating Systems (Windows, Linux, macOS):** Require architectures capable of managing multiple processes and resources simultaneously.
- **Educational Prototypes and Academic Projects:** Its simplicity in design and programming makes it an excellent foundation for learning.
- **Non-Critical Applications Where Speed Is Not a Priority:** Such as office automation tools, simple computing devices, or low-constraint embedded systems.

B) Harvard Architecture:

- **Critical Embedded Systems (Aerospace, Automotive, Robotics):** Where reliability, high speed, and low power consumption are key.
- **Real-Time Applications:** Such as motor control, industrial process regulation, and automated control systems.
- **Digital Signal Processing (DSP):** Used in audio, video, telecommunications, and data compression.
- **Internet of Things (IoT) Devices:** Require high energy efficiency and fast local data processing within limited hardware resources.
- **Medical and Military Systems:** Where precision, responsiveness, and reliability are crucial for safety and operational performance.

II.5 Overview of Microcontroller Boards by Architecture

The study of processor architectures is not limited to theory; it finds direct application in microcontrollers and development boards widely used in embedded systems and educational environments. Depending on the adopted architecture — Von Neumann, Harvard, or a modified hybrid version — the performance, energy efficiency, and design flexibility of the system vary significantly.

II.5.1 Microcontrollers Based on the Von Neumann Architecture

Von Neumann–based microcontrollers use a single memory space for both instructions and data. This simplifies hardware design and programming but limits execution speed due to the shared data bus.

Examples:

- x86 processors (Intel, AMD): still rely on an optimized Von Neumann model.
- 8051 family: commonly used in simple, industrial, or educational applications.
- ARM Cortex-A series: used in Raspberry Pi boards, adopting a Von Neumann approach with modern enhancements (caches, pipelines, MMU).

Associated Boards:

- Raspberry Pi (Pi 3, Pi 4): ideal for learning, IoT projects, and applications requiring a full operating system (Linux).

- Industrial x86-based boards: such as embedded PCs and industrial mini-PCs.

II.5.2 Microcontrollers Based on the Harvard Architecture

In the Harvard architecture, instruction and data memories are physically separated, enabling parallel execution. This structure enhances performance, reduces latency, and optimizes energy efficiency — features that are essential in embedded systems.

Examples:

- AVR (Atmel, now Microchip): classic Harvard architecture, e.g., Arduino Uno.
- PIC (Microchip): widely used in industrial control applications.
- DSP (Digital Signal Processors): specialized for signal processing tasks.

Associated Boards:

- Arduino Uno / Mega (AVR): popular for education and rapid prototyping.
- PICKit boards (PIC-based): used in industrial and control systems.
- DSP development boards: employed in audio, imaging, and communication systems.

II.5.3 Microcontrollers Based on Modified Architectures (Optimized Von Neumann & Modified Harvard)

Most modern microcontrollers use a hybrid architecture to combine the strengths of both models.

- Optimized Von Neumann: adds cache memory, pipelines, and memory management units (MMU) to increase execution speed despite a shared bus.
- Modified Harvard: maintains separate instruction and data memories but allows certain resources to be shared through hierarchical buses or cache layers.

Examples:

- **ARM Cortex-M (STM32, nRF52, etc.):** modified Harvard architecture, ideal for embedded and IoT systems.
- **ESP32 (Espressif):** modified Harvard, integrating **Wi-Fi** and **Bluetooth** connectivity.

- **STM32 Discovery / Nucleo boards:** educational and industrial boards based on ARM Cortex-M.
- **Raspberry Pi:** optimized Von Neumann with ARM Cortex-A, suited for complex systems.

Table II. 2 - Summary Table

Architecture	Main Characteristics	Examples of Microcontrollers	Associated Boards
Von Neumann	Single memory, shared bus. Simple design but limited by bottleneck.	8051, ARM Cortex-A, x86	Raspberry Pi, industrial x86 boards
Harvard	Separate instruction/data memory, parallel execution, higher speed.	AVR, PIC, DSP	Arduino Uno/Mega, PICKit boards, DSP boards
Modified Harvard	Combines flexibility and speed; uses pipelines, caches, and energy optimization.	ARM Cortex-M (STM32), ESP32	STM32 Nucleo/Discovery, ESP32 DevKit
Optimized Von Neumann	Adds cache, pipeline, MMU, hierarchical bus. Improves flexibility and performance.	ARM Cortex-A, modern x86 processors	Raspberry Pi 3/4, embedded PCs