



Center of Excellence: Supercomputing for AI & Big-Data

Processor Hardware and Instruction Set Architecture

by: Tassadaq Hussain

Director Centre for AI and BigData

Professor Department of Electrical Engineering

Namal University Mianwali

Collaborations:

Barcelona Supercomputing Center, Spain

European Network on High Performance and Embedded Architecture and Compilation

Pakistan Supercomputing Center

Topics

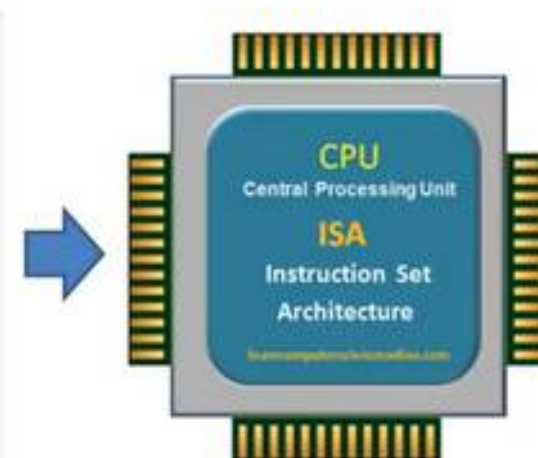
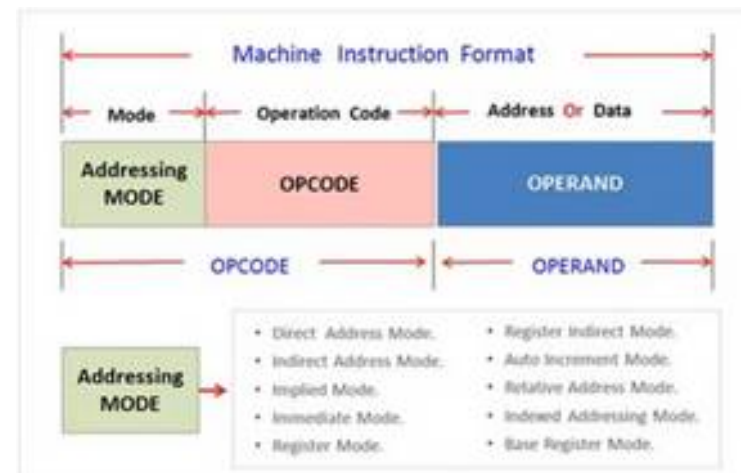
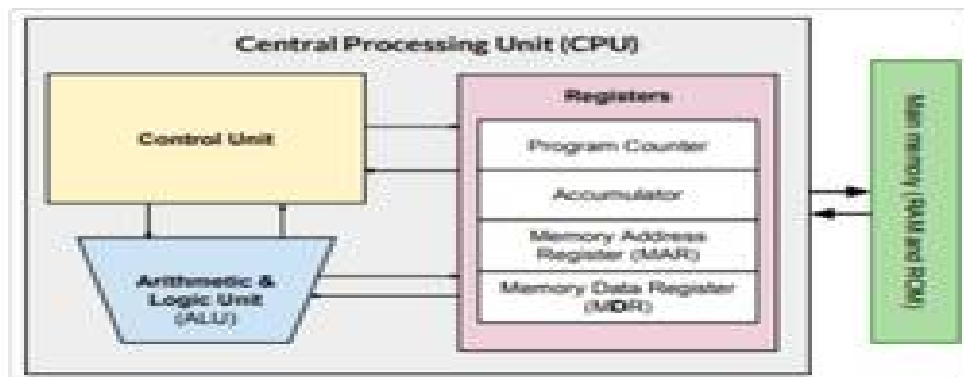
- 1. Basic Processor Architecture**
2. Different Types of Processor Architectures
3. RISC-V Processor Architecture
4. RISC-V Instruction Set Architecture
5. Programming RISC-V using assembly language

Basic Processor Architecture

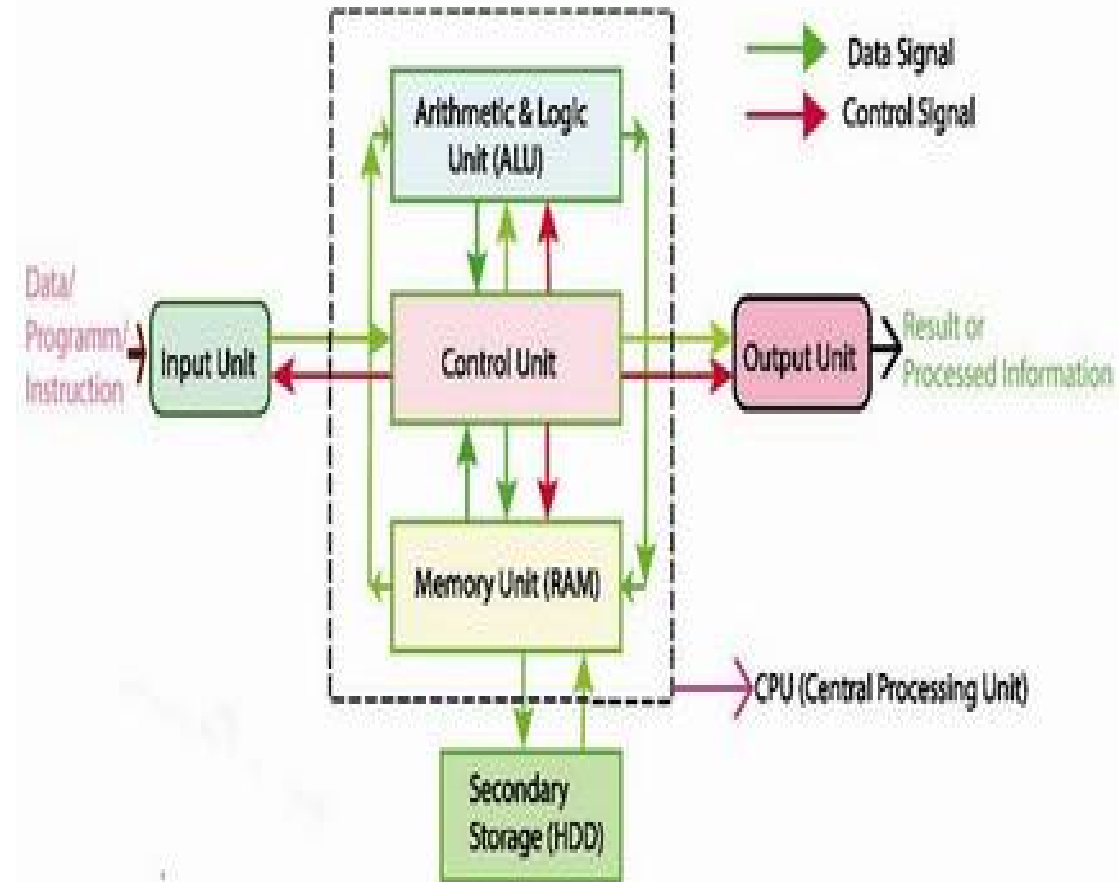
Processor Architecture refers to the design and organization of a processor's central processing unit (CPU).

Components of Processor:

- **Arithmetic and Logic Unit:** Performs mathematical calculations.
- **Control Unit:** Control the overall processing of the processor.
- **Decoders Unit:** Convert coded instructions into signals that can control other components.
- **Registers:** Hold data, instructions, and addresses temporarily during processing.
- **Buses:** Electrical pathways that transmit data and signals between components. Types include the data bus, address bus, and control bus.



- **Clock:** Generates timing signals to synchronize the operations of the CPU components. The clock speed determines how many instructions per second the CPU can execute.
- **Instruction Set Architecture (ISA):** Defines the set of instructions the CPU can execute
- **Cache:** Stores frequently accessed data and instructions to speed up processing.
- **Memory Management Unit (MMU):** Handles the translation of virtual addresses to physical addresses. Manages memory protection and caching.
- **Input/Output (I/O) Interfaces:** Allow the CPU to communicate with peripheral devices. Include ports and controllers for devices such as keyboards, mice, and storage.
- **Power Control Unit:**



Arithmetic Logic Unit ALU:

In computing, an arithmetic logic unit (ALU) is a combinational digital circuit that performs arithmetic and bitwise operations on integer binary numbers.

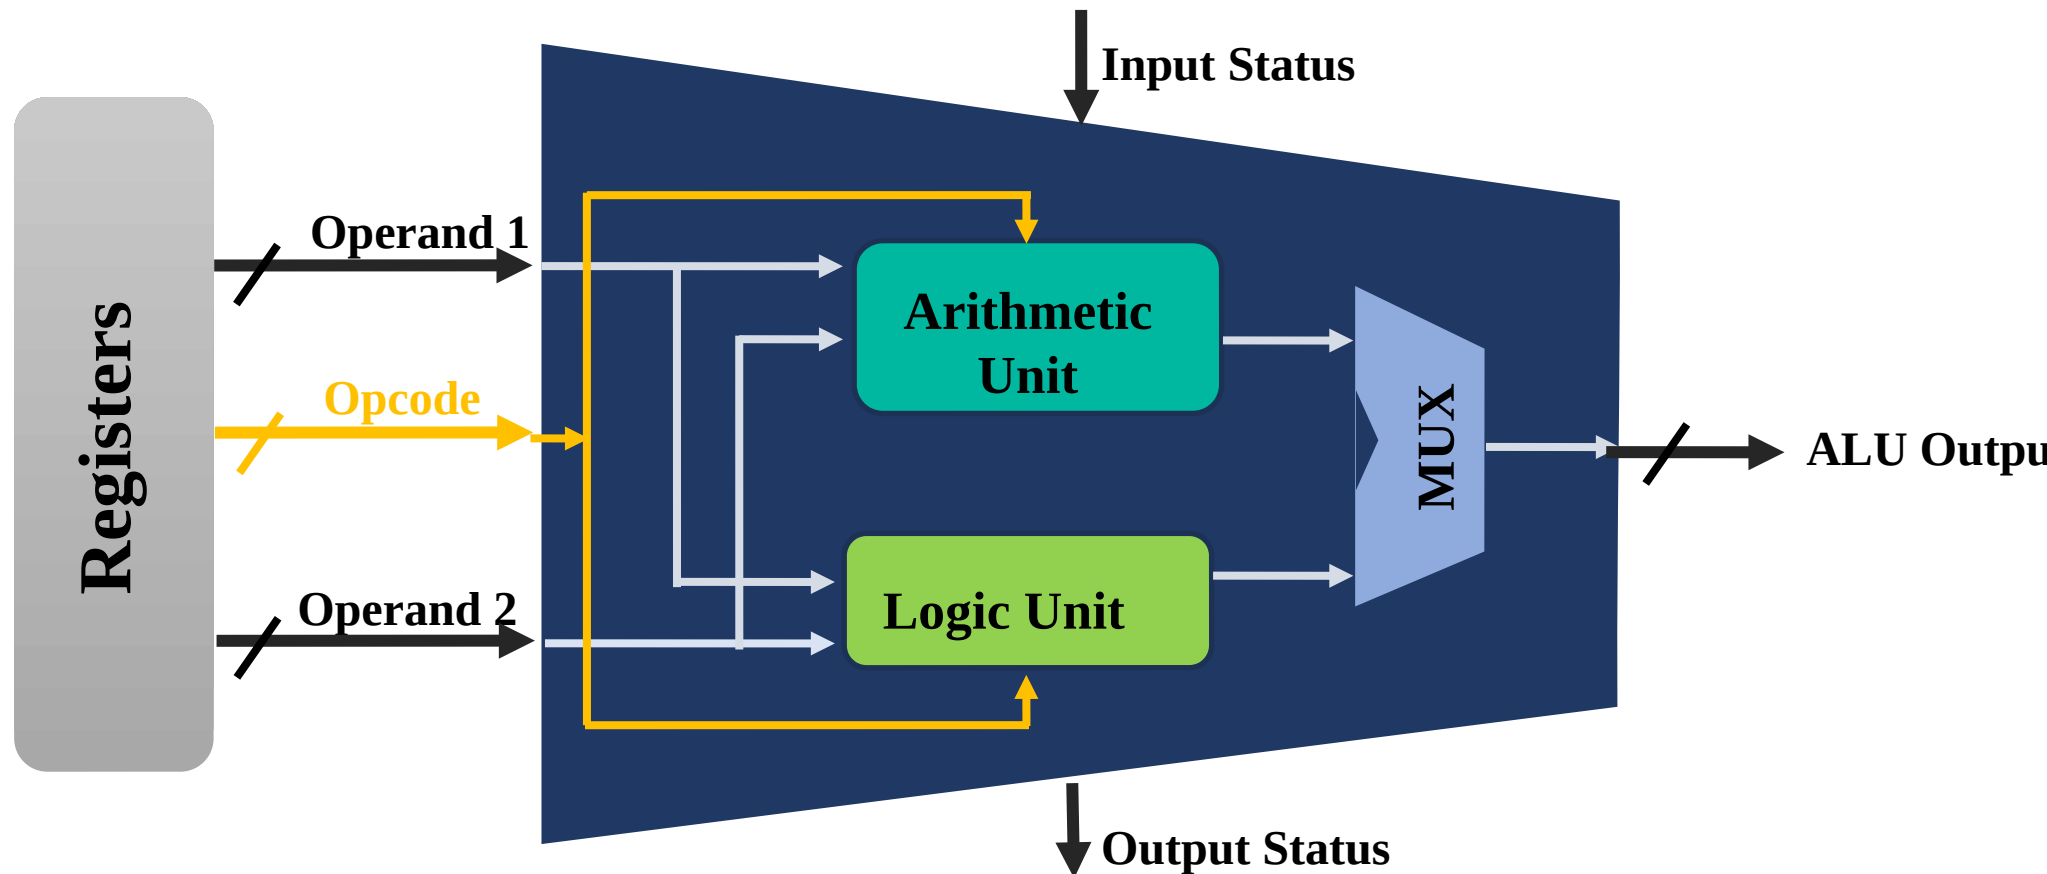
It is a fundamental building block of many types of computing circuits, including the central processing unit (CPU) of computers, FPUs, and graphics processing units (GPUs).

Functions of ALU:

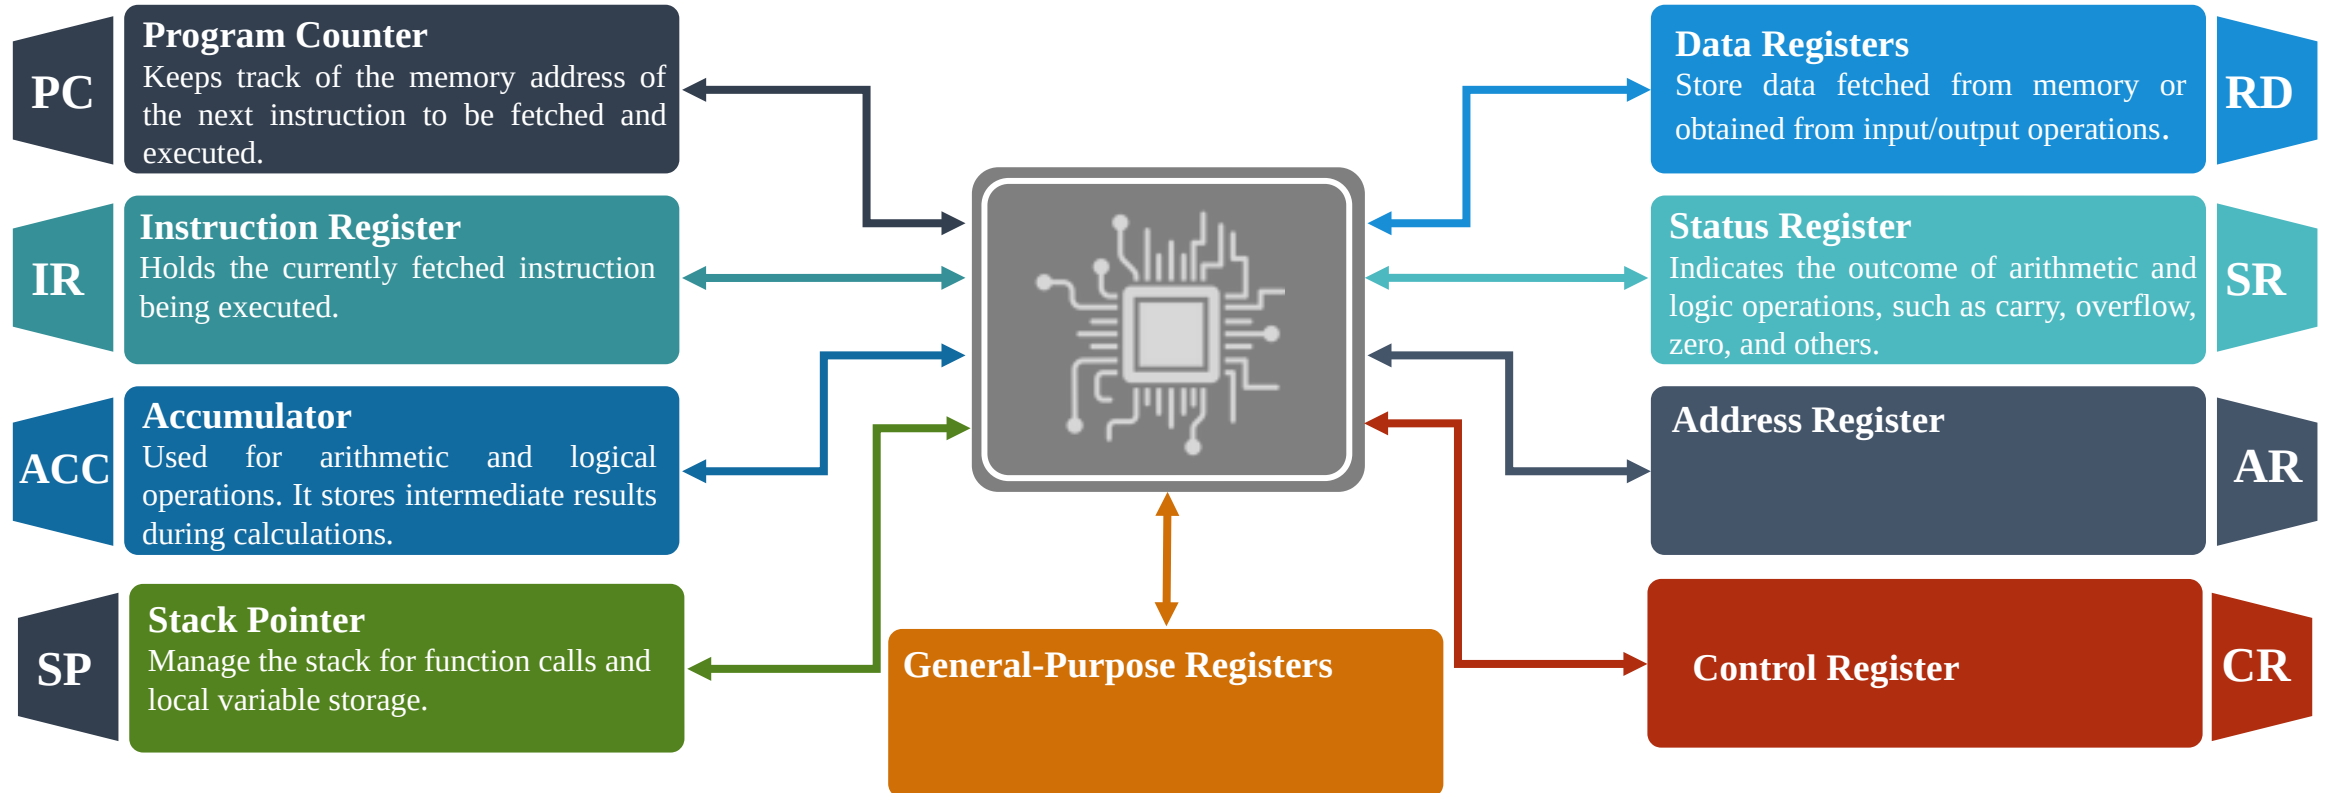
Basic Operations	Basic Instructions
Arithmetic operations	Addition, Subtraction, Multiplication, division
Logical operations	Logical Sum(OR), Logical Product(AND), Logical negation (NOT)
Comparison	Comparison Instruction (size compare)
Branch	Branch instructions to alter the instruction sequence based on conditions

Registers

- Registers are a type of computer memory built directly into the processor that is used to store and manipulate data during the execution of instructions.
- A register may hold an instruction, a storage address, or any kind of data (such as a bit sequence or individual characters).



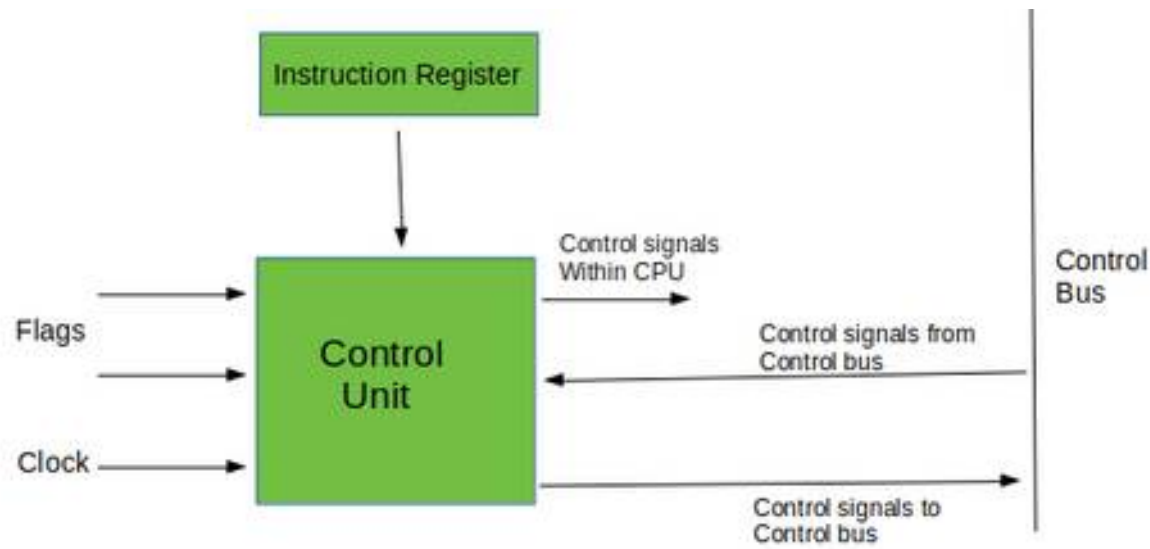
Registers in Processor Architecture



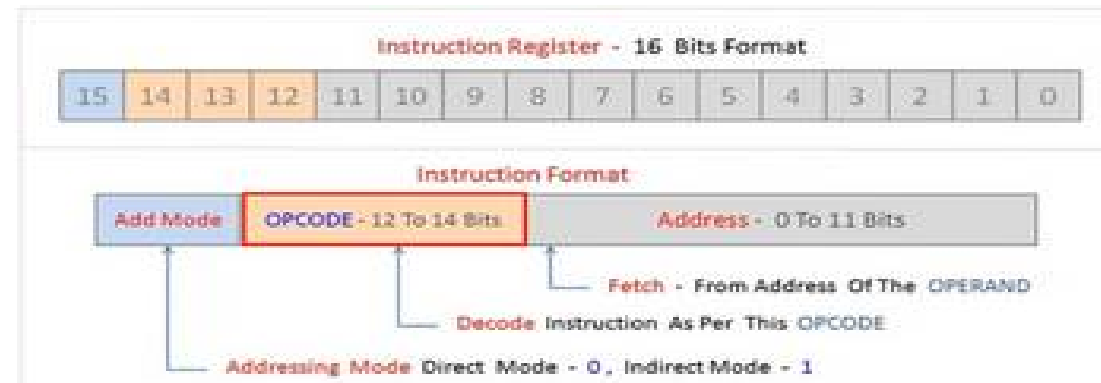
Control Unit:

The control unit controls all the operations of the processor. It retrieves, decodes and executes the code instructions one-by-one in the order they are stored in the main memory.

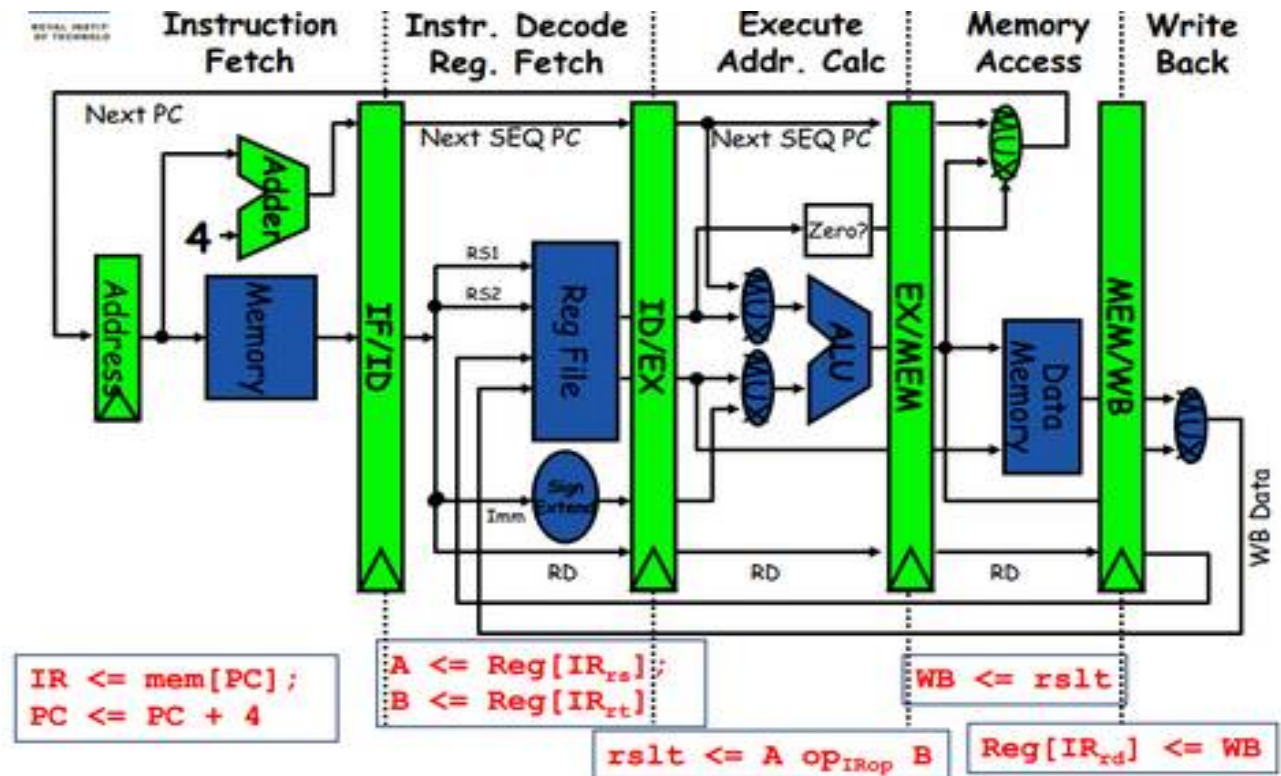
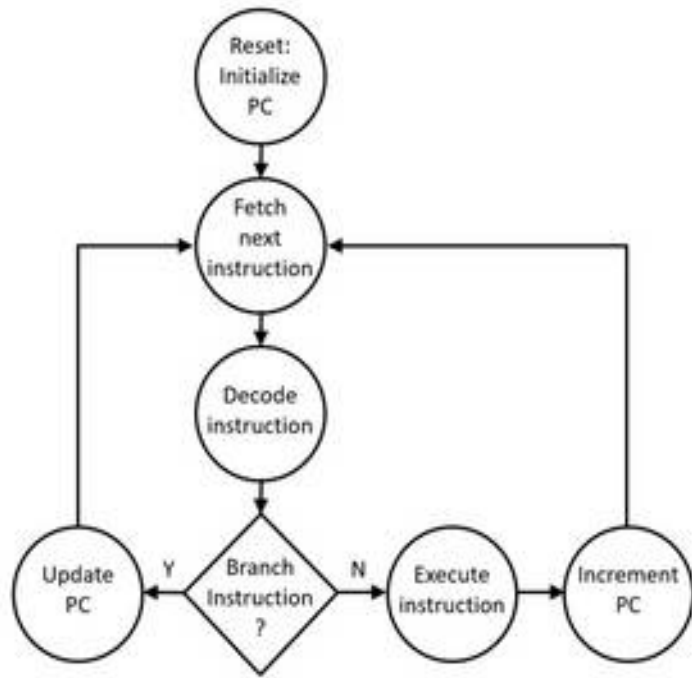
It instructs the arithmetic logic unit, memory, input/output devices how to respond to the instructions of the program.



Block Diagram of the Control Unit



Stages: Execution Clock Cycles



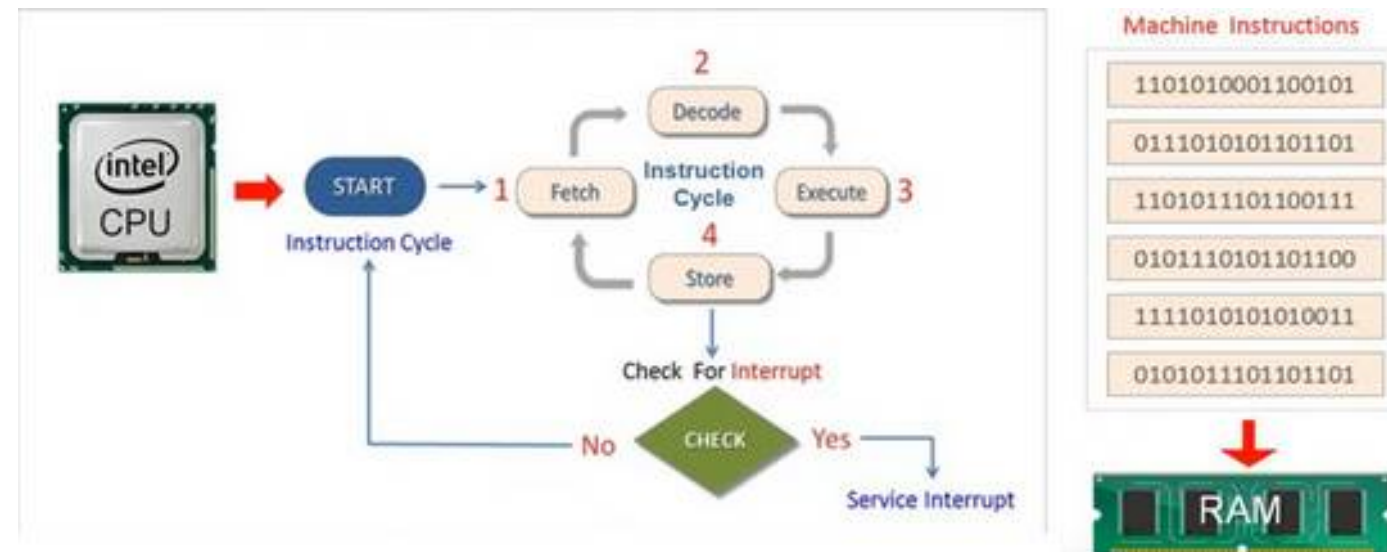
Instruction Set Architecture (ISA)

An **Instruction Set Architecture (ISA)** is part of the abstract model of a computer that defines how the CPU is controlled by the software.

- The ISA acts as an interface between the hardware and the software, specifying both what the processor is capable of doing as well as how it gets done.
- The ISA defines the supported data types, the registers, how the hardware manages main memory, key features (such as virtual memory), which instructions a microprocessor can execute, and the input/output model of multiple ISA implementations.

- Provides:

- } **Programmability**
- } **Flexibility**
- } **Reusability**
- } **Adaptability**
- } **Accessibility**



Instruction Set Format

31	27	26	25	24	20	19	15	14	12	11	7	6	0	
funct7				rs2		rs1		funct3		rd		opcode		R-type
imm[11:0]						rs1		funct3		rd		opcode		I-type
imm[11:5]				rs2		rs1		funct3		imm[4:0]		opcode		S-type
imm[12:10:5]				rs2		rs1		funct3		imm[4:1:11]		opcode		B-type
imm[31:12]										rd		opcode		U-type
imm[20:10:1:11:19:12]										rd		opcode		J-type

A form of representation of an instruction composed of fields of binary numbers.”

Fields of instruction:

There are several fields of the instruction that serve a specific role in the format. Some common are fields are given below:

1. Opcode:

- Specifies the operation to be performed (e.g., add, subtract, load, store).
- Determines what action the CPU should take.

2. Operand:

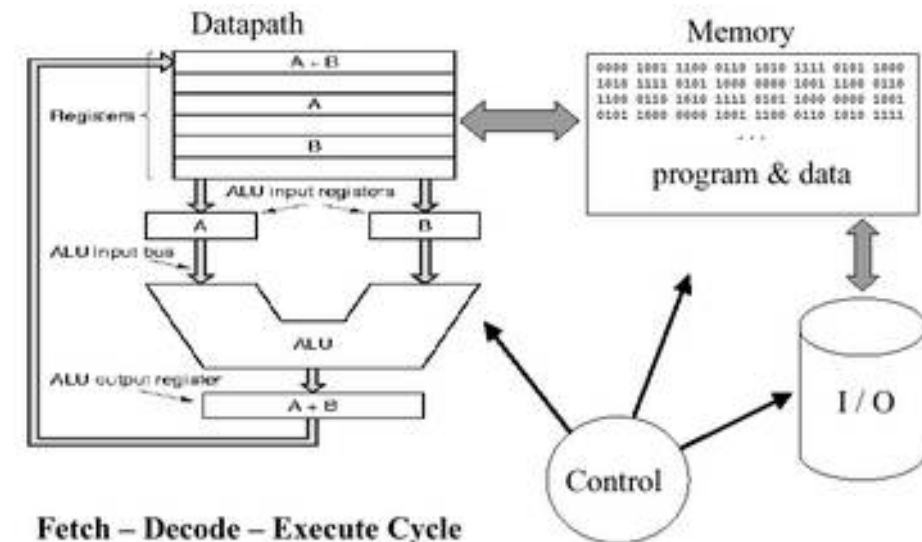
- The data or the addresses of the data on which the operation is to be performed.
- Can include immediate values, register addresses, or memory addresses.

3. Addressing Modes:

Processor uses different Addressing modes Common modes include: immediate, direct, indirect, register, and indexed addressing.

4. Registers:

Specifies which CPU registers are to be used in the operation.
Could include source and destination registers.



Instruction Types

A computer's instructions can be any length and have any number of addresses.

- The arrangement of a computer's registers determines the different address fields in the instruction format.
- The instruction can be classified as three, two, and one address instruction or zero address instruction, depending on the number of address fields.

Based on these differences the instructions are classified as

- 1) **Three Address Instruction**
- 2) **Two Address Instruction**
- 3) **One Address Instruction**
- 4) **Zero Address Instruction**

Three Address Instruction:

Three-address instruction is a format of machine instruction. It has one opcode and three address fields.

One address field is used for destination and two address fields for source.

OPCODE	DESTINATION	SOURCE 1	SOURCE 2
--------	-------------	----------	----------

Example:

ADD	R1, A, B	$R1 = M[A] + M[B]$
ADD	R2, C, D	$R2 = M[C] + M[D]$
MUL	X, R1, R2	$M[X] = R1 * R2$

Two Address Instruction:

Two-address instruction is a format of machine instruction. It has one opcode and two address fields which may be memory locations or registers..

One address field is used for destination and one address field for source.

For example, a two-address instruction might add the contents of two registers together and store the result in one of the registers.

OPCODE	DESTINATION	SOURCE
--------	-------------	--------

Example

:

MOV	R1, A	$R1 = M[A]$
ADD	R1, B	$R1 = R1 + M[B]$

One Address Instruction:

These instructions specify one operand or address, which typically refers to a memory location or register.

The instruction operates on the contents of that operand, and the result may be stored in the same or a different location.

For example, a one-address instruction might load the contents of a memory location into a register.

OPCODE	DESTINATION
--------	-------------

Example:

STORE	T	$M[T] = AC$
LOAD	C	$AC = M[C]$

Zero Address Instruction:

These instructions do not specify any operands or addresses. Instead, they operate on data stored in registers or memory locations implicitly defined by the instruction.

For example, a zero-address instruction might simply add the contents of two registers together without specifying the register names.

Types of Instructions and Addressing Modes

Implied Mode

Example: CLC ; Clear the carry flag, no operands needed

Immediate Mode

Example: ADDI x1, x2, 10 ; Add immediate value 10 to register x2 and store result in x1

Register Mode

Example: MOV r0, r1 ; Move the contents of register r1 to register r0

Register Indirect Mode

Example: LW \$t0, 0(\$t1) ; Load the word at the address in \$t1 into \$t0

Autodecrement Mode

Example: MOV -(R1), R0 ; Decrement R1 and then move the value at the new address in R1 to R0

Autoincrement Mode

Example: MOV (R1)+, R0 ; Move the value at the address in R1 to R0, then increment R1

Direct Address Mode

Example: LDA \$4000 ; Load the accumulator with the value at memory address \$4000

Indirect Address Mode

Example: JMP (\$1234) ; Jump to the address stored at memory location \$1234

Indexed Addressing Mode

Example: MOV AX, [BX+SI] ; Move the value at address (BX + SI) into AX

High-level
language
program
(in C)

```
swap(size_t v[], size_t k)
{
    size_t temp;
    temp = v[k];
    v[k] = v[k+1];
    v[k+1] = temp;
}
```

↓

Compiler

↓

Assembly
language
program
(for RISC-V)

```
swap:
    slli x6, x11, 3
    add x6, x10, x6
    lw x5, 0(x6)
    lw x7, 4(x6)
    sw x7, 0(x6)
    sw x5, 4(x6)
    jalr x0, 0(x1)
```

↓

Assembler

↓

Binary machine
language
program
(for RISC-V)

```
000000000001101011001001100010011
000000000011001010000001100110011
0000000000000000110011001010000011
000000000100000110011001110000011
00000000011100110011000000100011
00000000010100110011010000100011
0000000000000000100000001100111
```



Apple
338S00456 PMIC

Apple A12 APL1W81 +
Micron MT53D512M64D4SB-046 XT:E
4GB Mobile LPDDR4x SDRAM

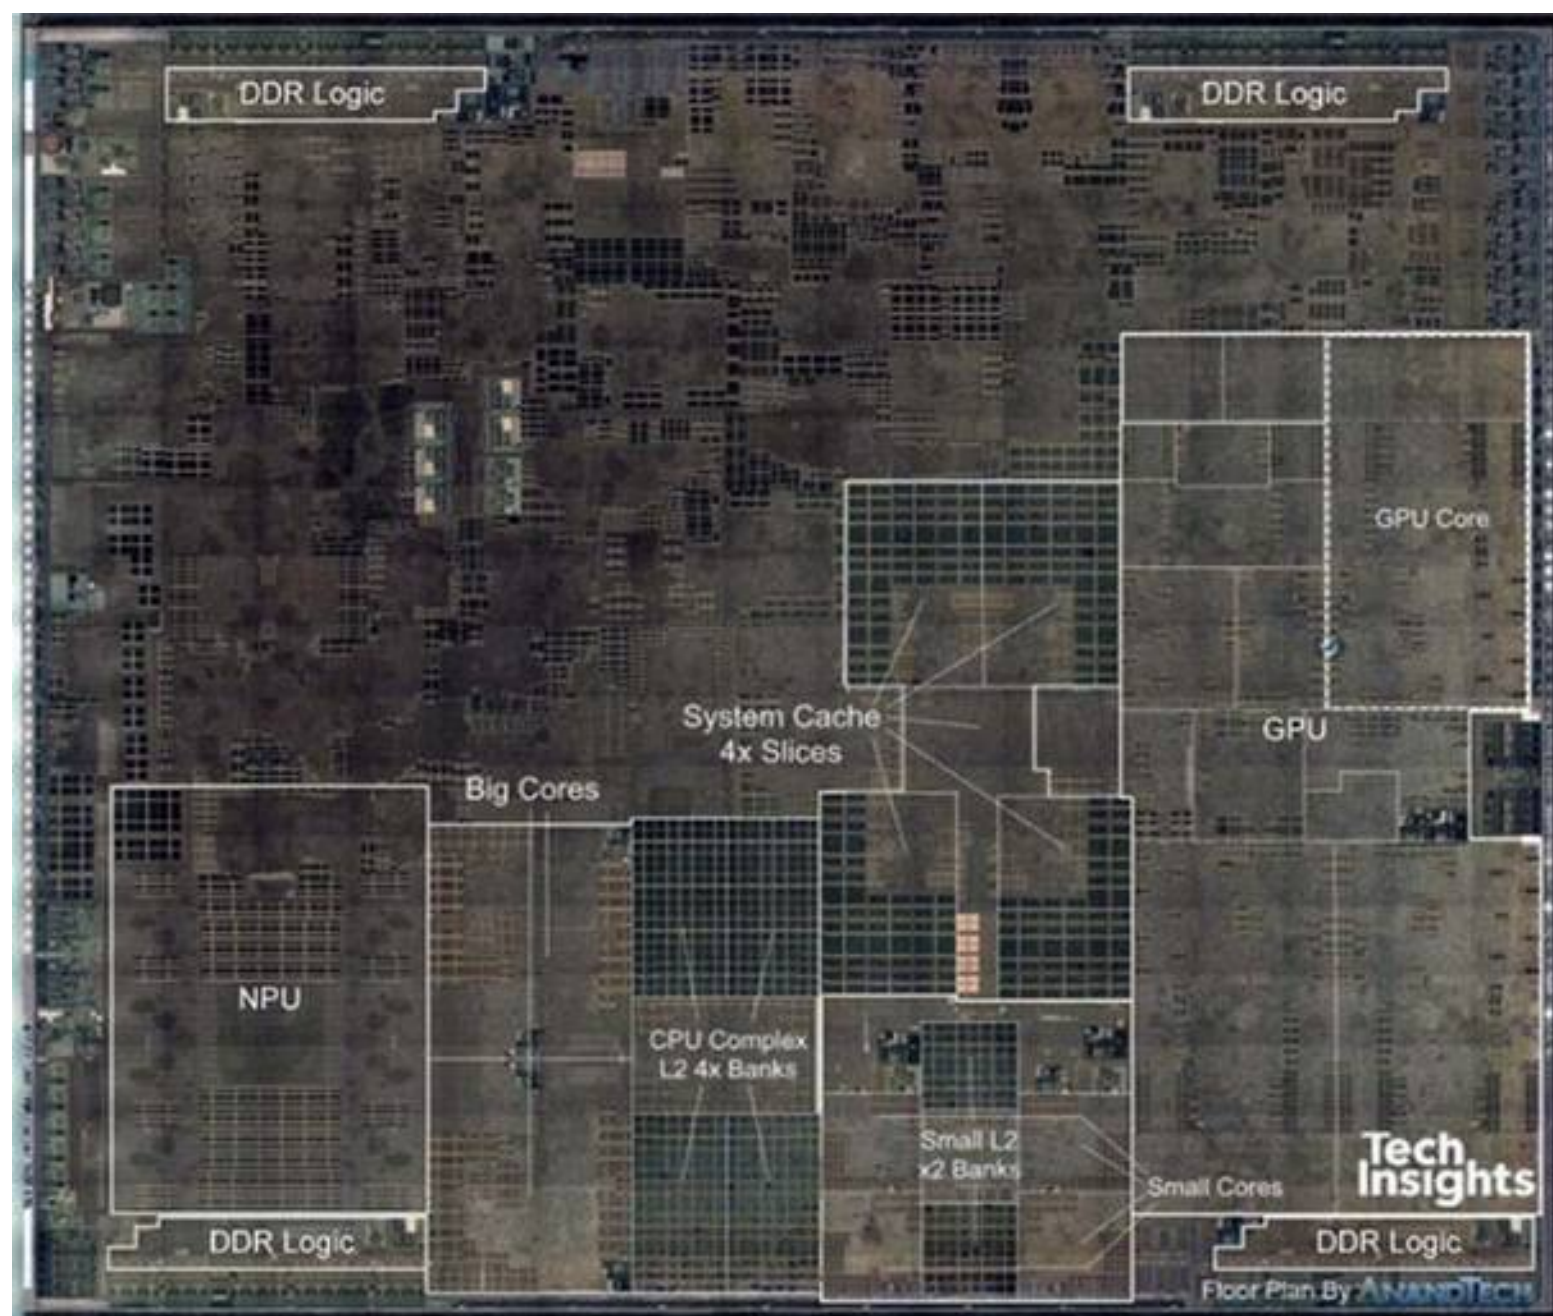
STMicroelectronics
STB601A0 PMIC

Apple
338S00375 PMIC

Texas Instruments
SN2600B1
Battery Charger

Apple 338S00411
Audio Amplifier

Tech
Insights



Instructions Types

R-type: Integer computation instructions on registers.

I-type: Integer computation instructions on registers and immediate values. Also includes JALR, Load instructions.

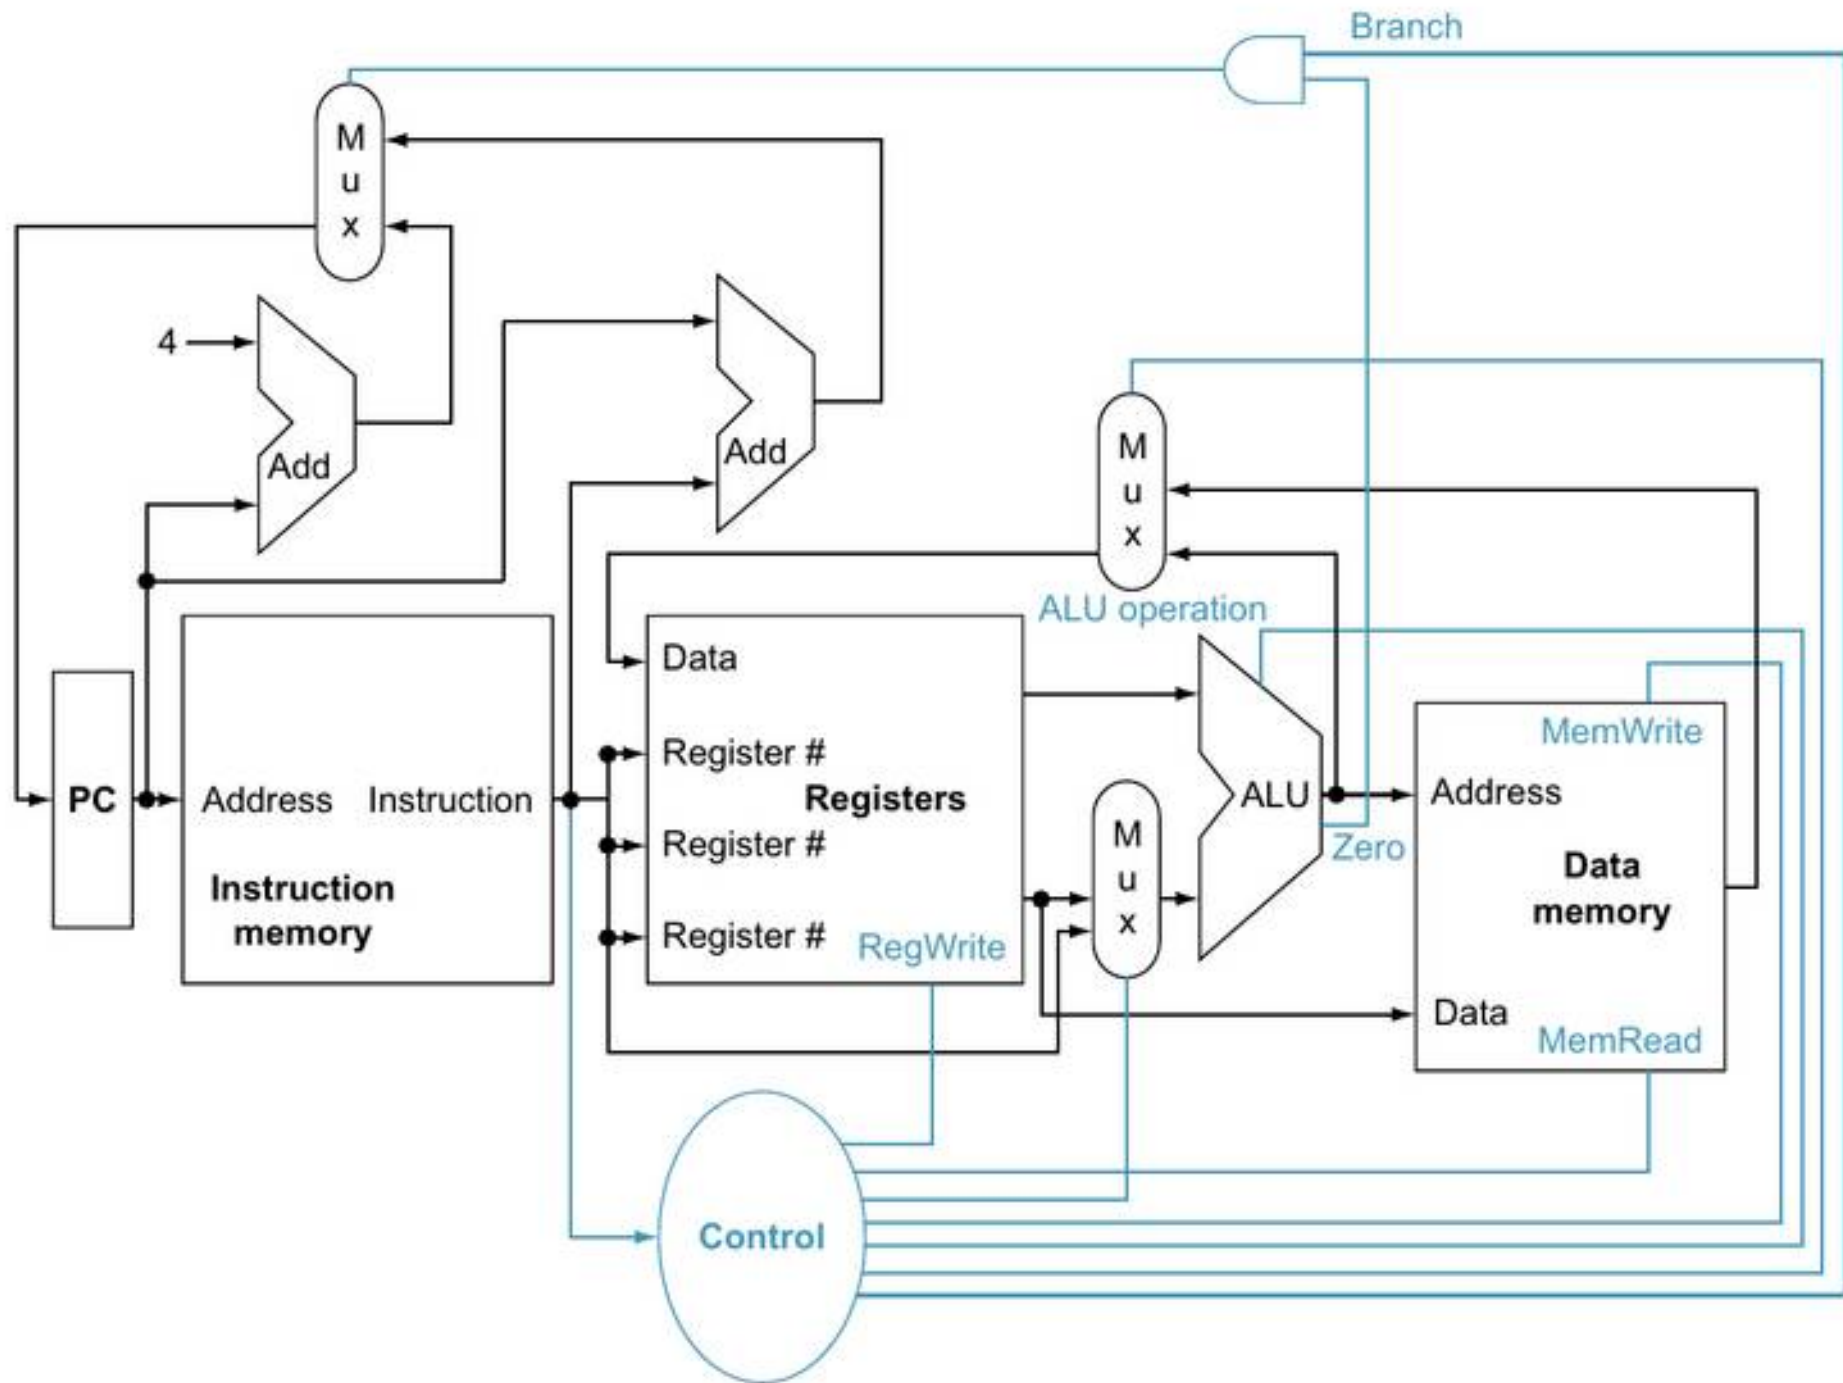
S-type: Store instructions.

B-type: Branch instructions.

U-type: Special instructions like LUI, AUIPC.

J-type: Jump instructions like JAL.

31	30	25	24	21	20	19	15	14	12	11	8	7	6	0				
funct7				rs2			rs1		funct3		rd			opcode		R-type		
imm[11:0]						rs1		funct3		rd			opcode			I-type		
imm[11:5]				rs2			rs1		funct3		imm[4:0]			opcode		S-type		
imm[12]		imm[10:5]			rs2			rs1		funct3		imm[4:1]		imm[11]		opcode		B-type
imm[31:12]										rd			opcode			U-type		
imm[20]		imm[10:1]				imm[11]		imm[19:12]			rd			opcode			J-type	



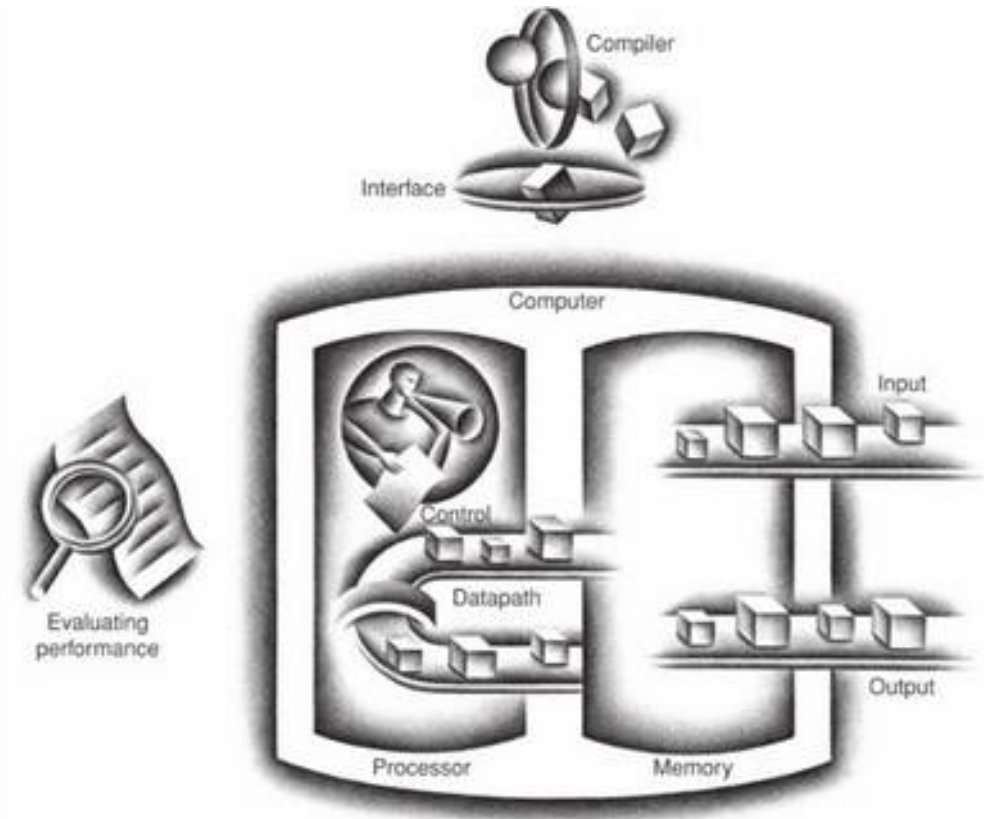
Important Parameters of a Processor

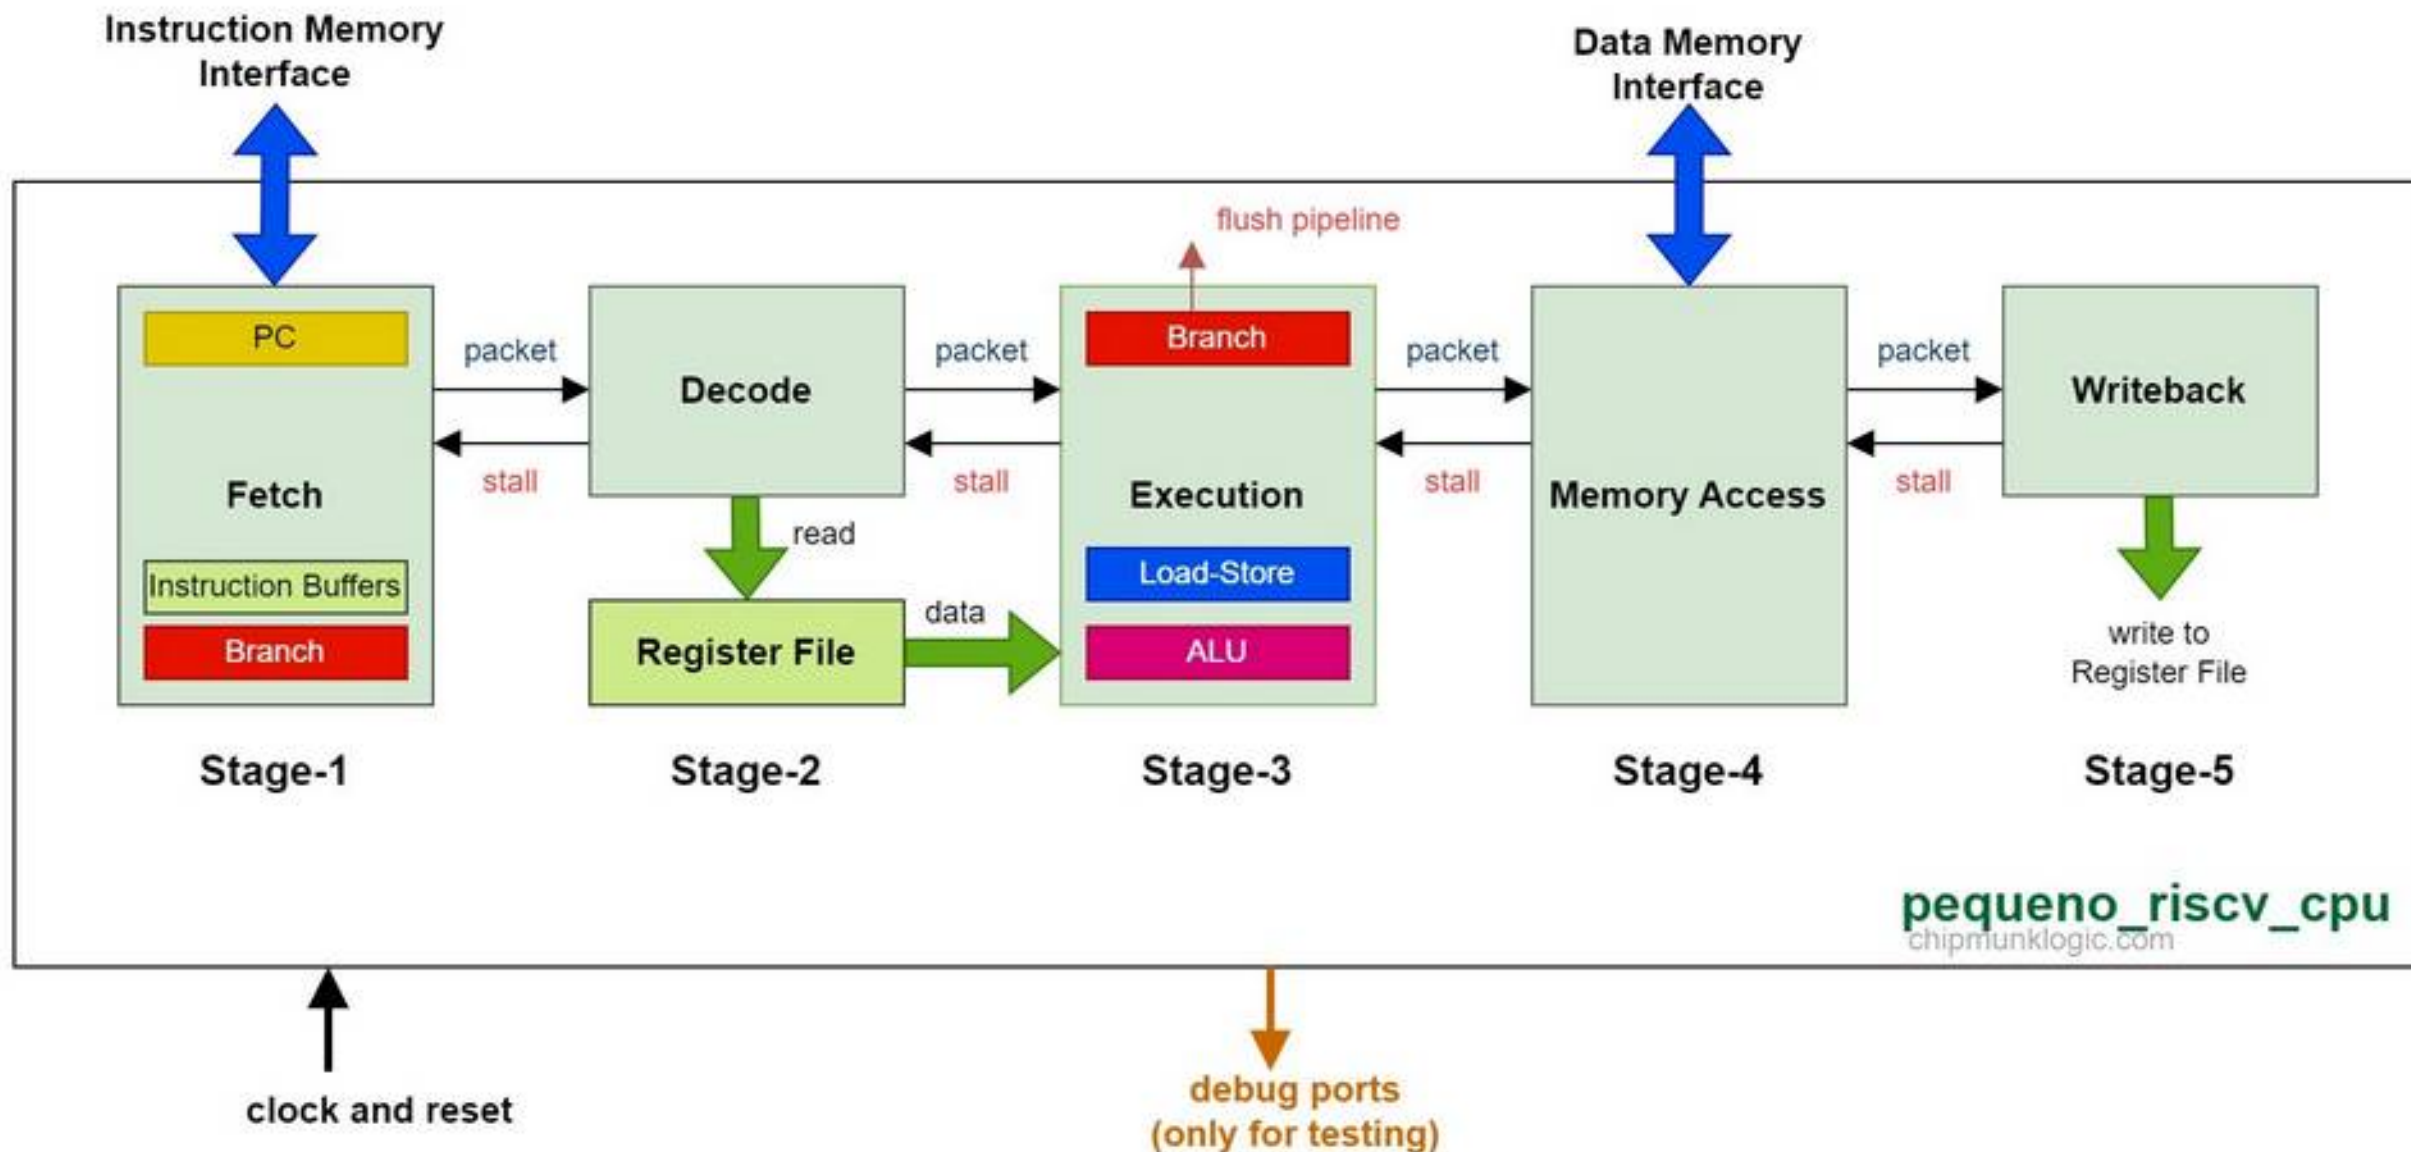
ISA →

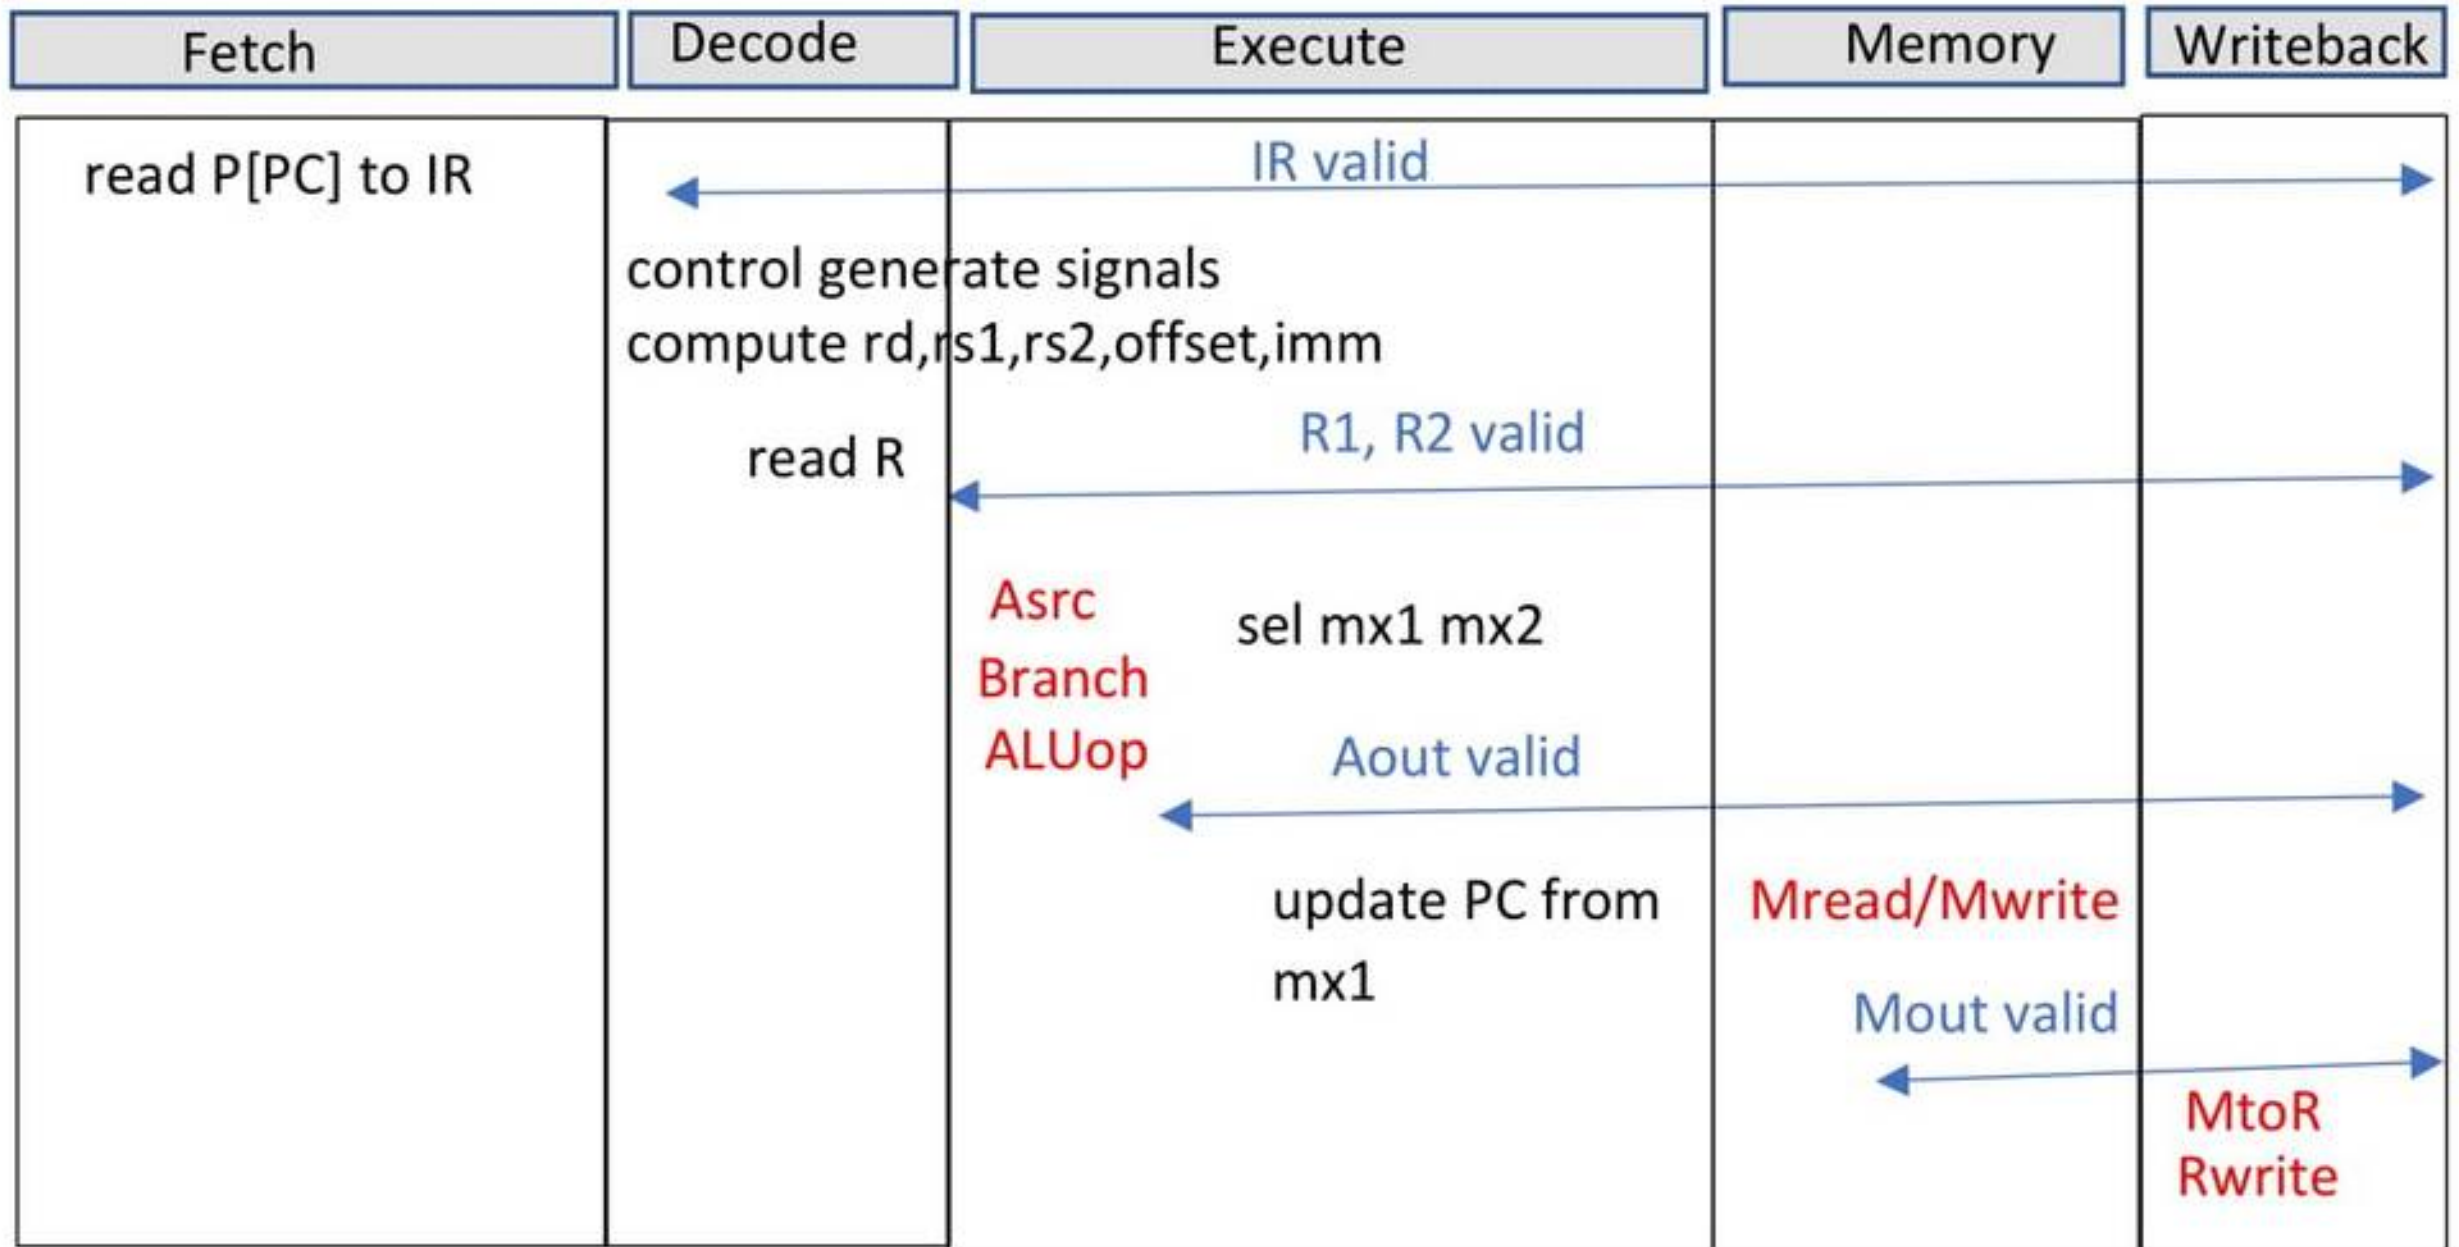
- Arithmetic Logic Unit (ALU): Performs arithmetic and logical operations.
- Floating Point Unit (FPU): Performs floating-point arithmetic operations (optional in some architectures).
- Registers: Small, fast storage locations within the CPU, used to store data and instructions temporarily.
- **Control Unit (CU): Directs operations of the processor, including instruction decoding and execution control.**
- **Program Counter (PC): Holds the address of the next instruction to be executed.**
- **Instruction Register (IR): Holds the current instruction being executed.**
- **Branch Predictor: Predicts the outcome of conditional branches to reduce instruction execution delays.**
- Bus Interface Unit (BIU): Manages data flow between the processor and external components like memory or peripherals.
- Pipeline: Allows overlapping execution of instructions to improve performance.
- Cache Memory: High-speed memory closer to the CPU, used to store frequently accessed data.
- Memory Management Unit (MMU): Manages memory access and translation between physical and virtual addresses.
- Input/Output (I/O) Unit: Handles communication with external devices.

Control Unit

- The Control Unit (CU) of a processor is responsible for directing the flow of data and the sequence of operations within the CPU. It coordinates the activities of the processor by interpreting and executing instructions. The CU can be broken down into several sub-parts, each handling specific tasks. Here are the main sub-parts:
- Instruction Decoder
- Sequencing Logic
- Control Logic Circuit
- Control Signal Generator
- Program Counter (PC) Control
- Status Flag Register
- Microprogram Control Storage
- Timing and Clock Control
- Branch and Jump Control
- Interrupt Control





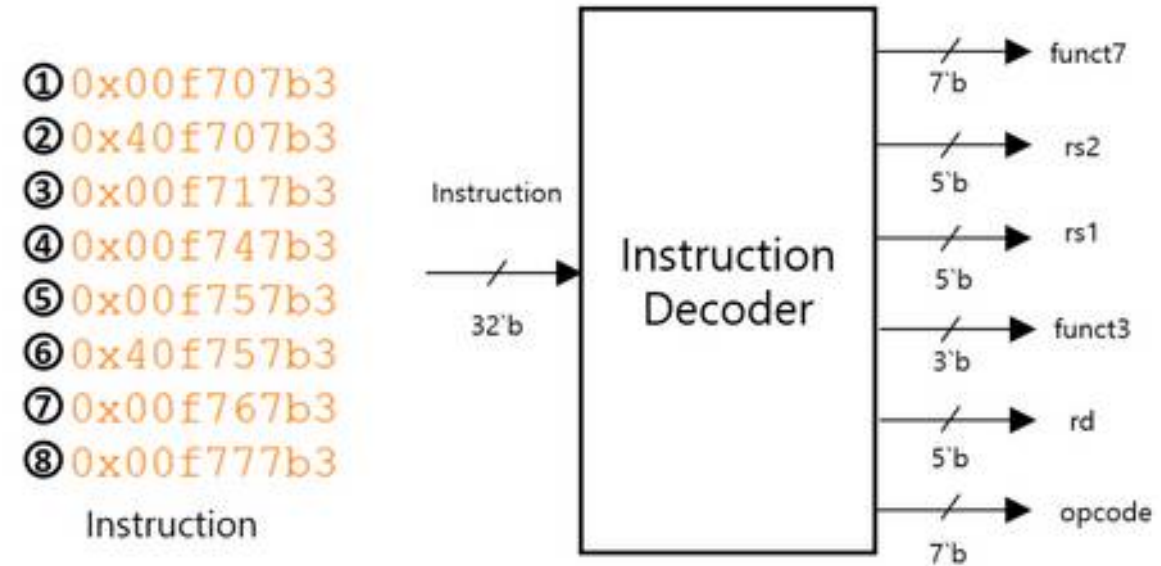


- **1. Instruction Decoder:**

- Decodes the fetched instruction from memory into signals that specify the operation to be performed.
Breaks down machine code into control signals that tell various parts of the processor what to do next.

- **2. Sequencing Logic:**

- Controls the order in which operations are carried out by determining the next instruction to be executed (through the Program Counter).
Manages the fetching, decoding, and execution cycle of instructions.
Synchronizes the processor's operation, often tied to the system clock.



clock cycle instruction	1	2	3	4	5
1	IF	ID	EX	MEM	WB
2		IF	ID	EX	MEM
3			IF	ID	EX
4				IF	ID
5					IF

- **3. Control Logic Circuit:**

- Contains the logic gates and combinational circuits that generate control signals based on the instruction decoded.

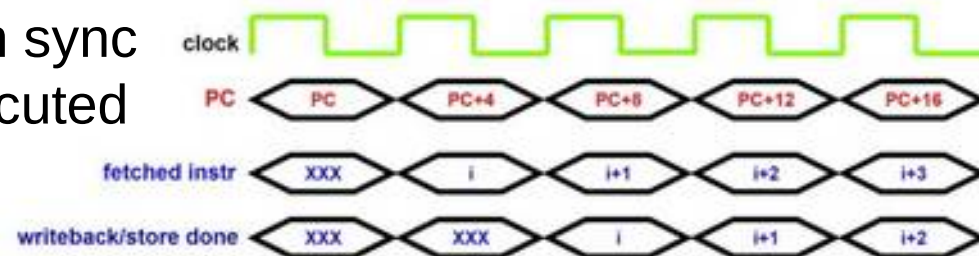
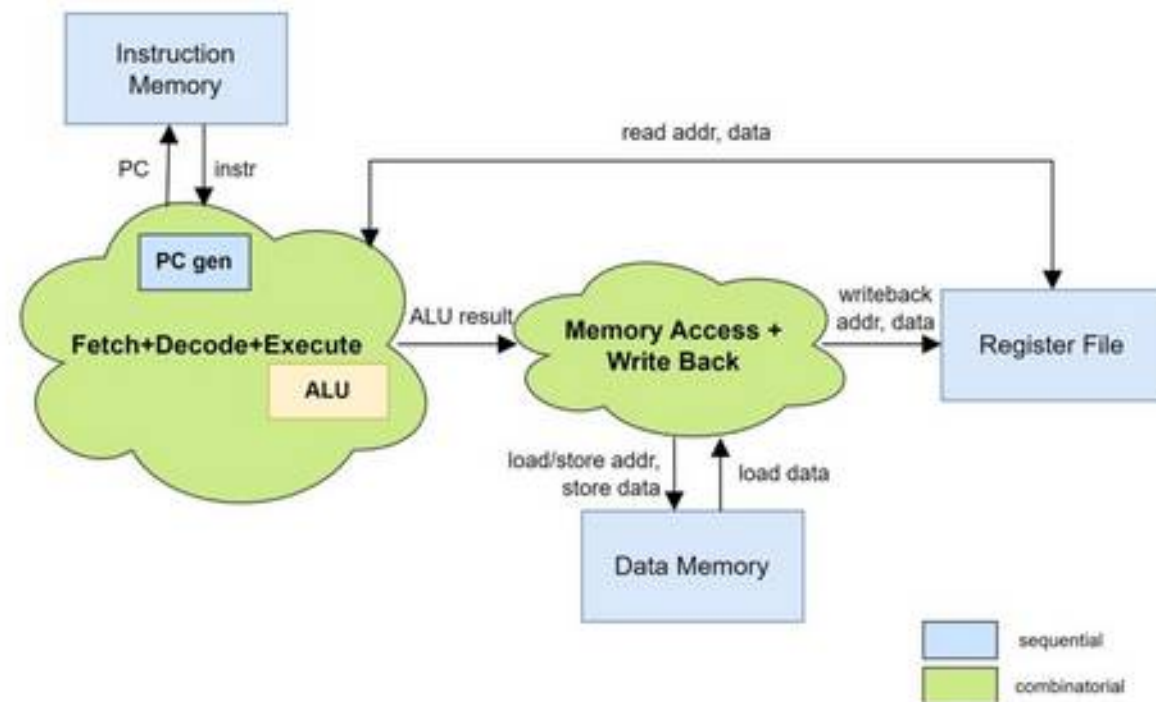
These control signals manage the internal data flow, timing, and operation of functional units (ALU, registers, etc.).

- **4. Control Signal Generator:**

- Generates the necessary control signals that dictate the actions of other parts of the CPU (ALU, memory interface, etc.). These signals direct data movement, ALU operation, register writes, and memory accesses.

- **ALU Control:** ALUOp, ALUSrc
- **Register Control:** RegWrite, RegRead
- **Memory Control:** MemRead, MemWrite, MemToReg
- **Branch/Jump Control:** Branch, Jump, PCSrc
- **Immediate Generation Control:** ImmSrc
- **Instruction Fetch Control:** PCWrite, IFIDWrite
- **Pipeline Control:** Stall, Flush, ForwardA, ForwardB
- **CSR Control:** CSRRead, CSRWrite

- **5. Control Instruction (PC):**
Manages the Program Counter, which holds the address of the next instruction to be executed.
Handles instruction sequencing, updating the PC after each instruction or adjusting it for branch and jump operations.
- **6. Status Flag Register:**
Contains flags that hold status information about the result of previous operations (e.g., Zero, Carry, Overflow, Sign flags).
These flags help the Control Unit make decisions regarding branching and conditional operations.
- **7. Timing and Clock Control:**
Coordinates the timing of operations across the CPU with the help of clock signals.
Ensures that all parts of the processor operate in sync and that each step of the instruction cycle is executed at the correct time.



- **8. Branch and Jump Control:**

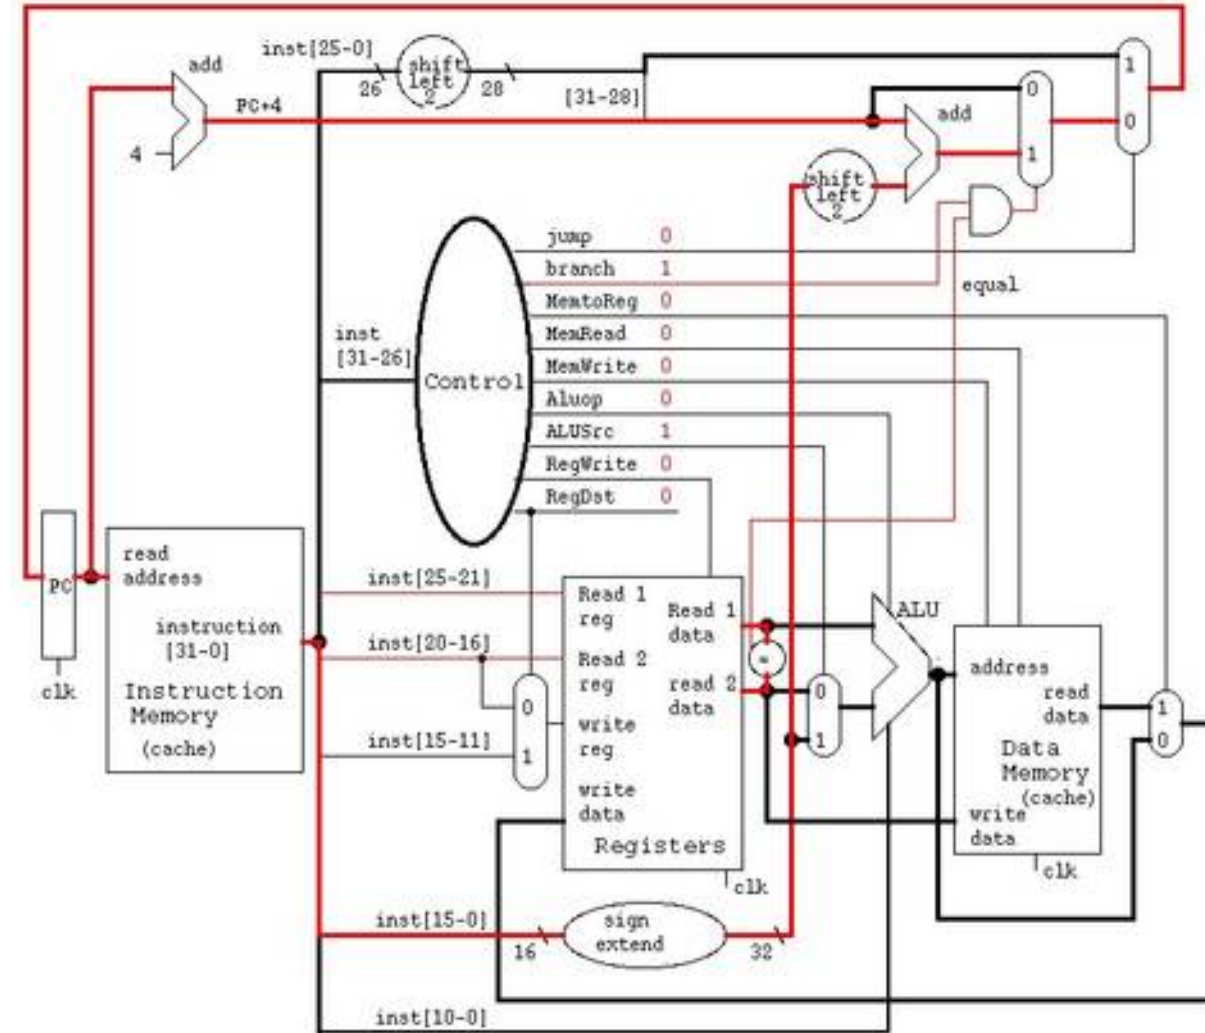
- Manages control transfer instructions, such as branches, jumps, and calls.
Works with the Branch Prediction Unit (in modern processors) to optimize branching and minimize delays caused by pipeline stalls.

- **9. Interrupt Control:**

- Handles interrupts by suspending the current execution and transferring control to the appropriate interrupt service routine.
Prioritizes interrupts and manages interrupt requests.

- **10. Microprogram Control Storage (in microprogrammed Cus):**

- In microprogrammed control units, the control signals are generated by executing a sequence of microinstructions stored in a microprogram memory (control memory).
Each instruction in the CPU is mapped to a set of microinstructions that control specific low-level operations.

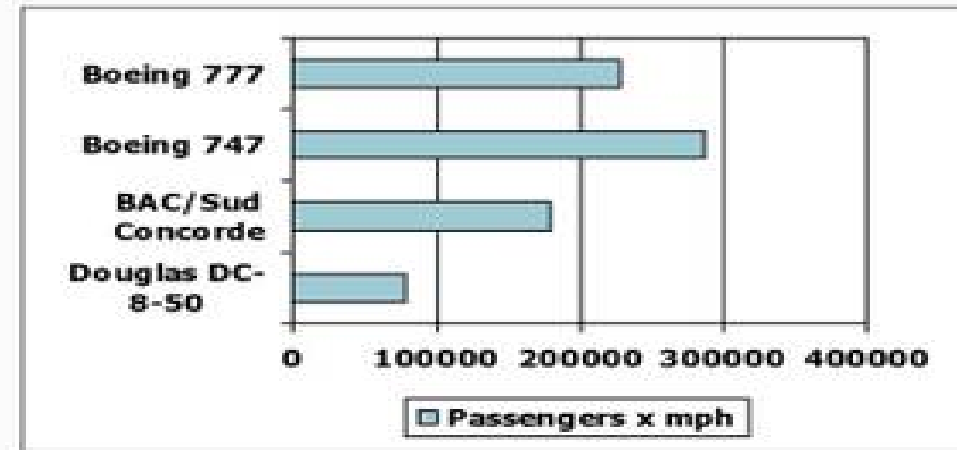
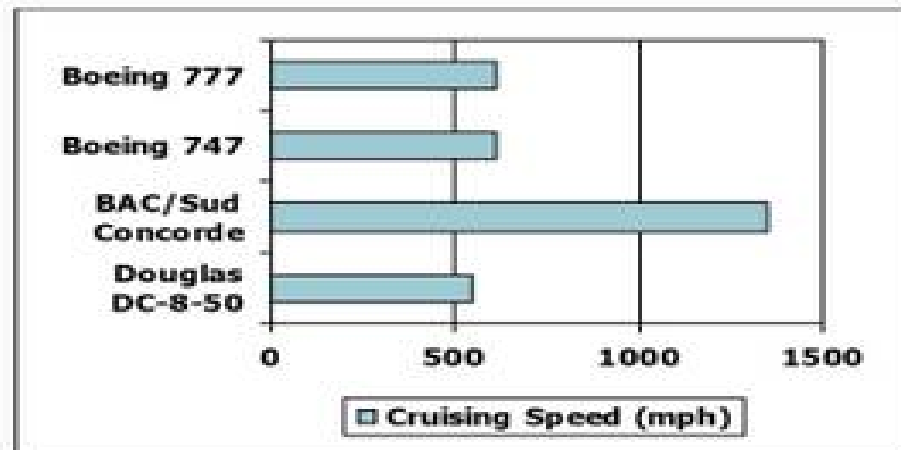
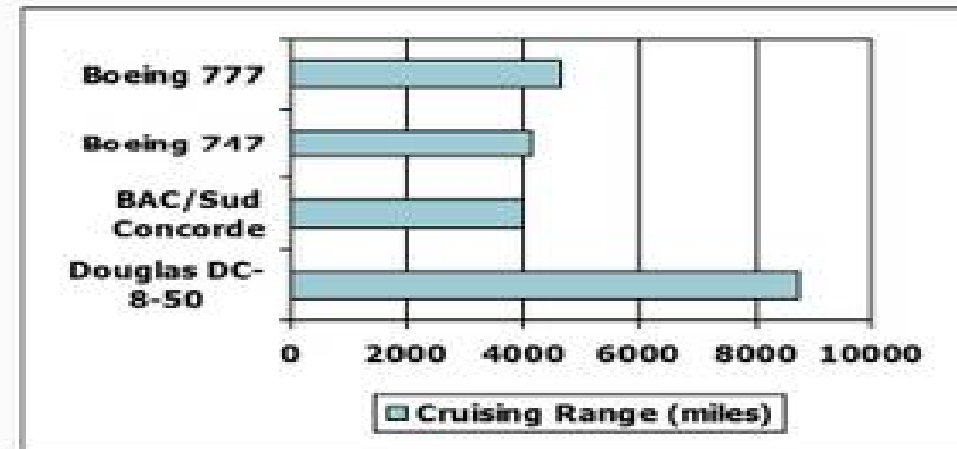
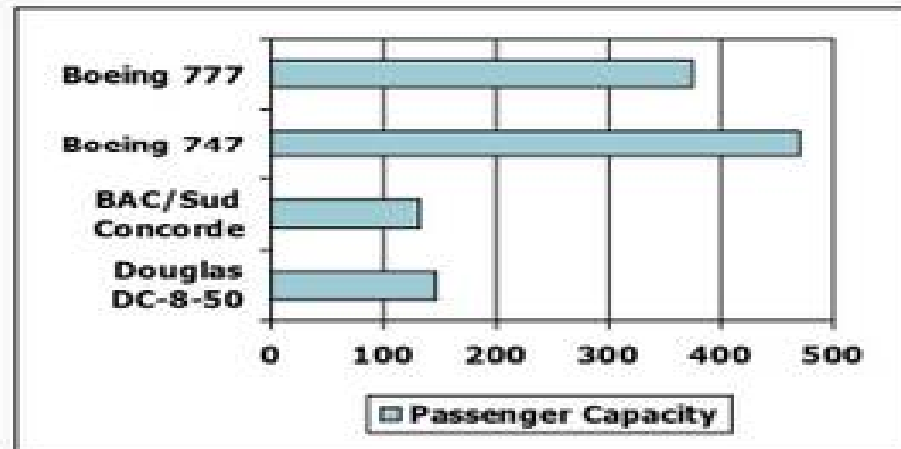




How do you define CPU performance?

Defining Performance

- Which airplane has the best performance?



Defining Performance

- Response time:
 - How long it takes to do a task.
 - It is also called execution time.
 - It includes disk access, memory access, I/O activities.
- Throughput :
 - Total amount of work done in a given time.
 - e.g., tasks/transactions/... per hour.
- We'll focus on response time for now...

Relative Performance

- Performance defined as:
- Then to evaluate two computers A & B .
- Can be phrased as “Processor X is n times faster than Processor Y”

Relative Performance

- Example: Assume a program runs in
 - 10s on Processor A.
 - 15s on Processor B.
 - How much is A faster than B.

$$\frac{\textit{Execution Time of B}}{\textit{Execution Time of A}} = \frac{15}{10} = 1.5$$

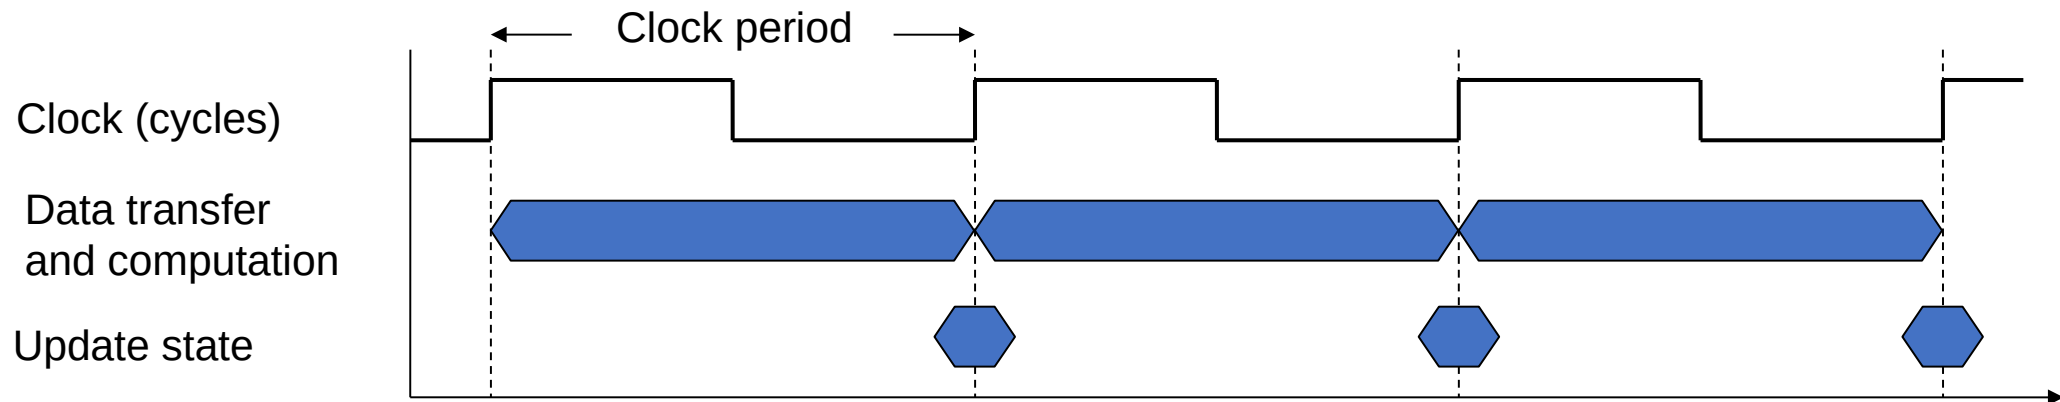
So, A is 1.5 times faster than B

Execution Time

- How do you measure execution time?
- Elapsed time
 - Total response time, including all aspects
 - Processing, I/O, OS overhead, idle time
 - Determines system performance
- CPU time
 - Time spent processing a given job
 - Discounts I/O time, other jobs' shares
- Different programs are affected differently by CPU and system performance

CPU Clocking

- Operation of digital hardware governed by a constant-rate clock



- Clock frequency (rate): cycles per second
 - e.g., $4.0\text{GHz} = 4000\text{MHz} = 4.0 \times 10^9\text{Hz}$
- Clock period: duration of a clock cycle
 - e.g., $250\text{ps} = 0.25\text{ns} = 250 \times 10^{-12}\text{s}$

CPU Time

$$\begin{aligned}\text{CPU Time} &= \text{CPU Clock Cycles} \times \text{Clock Cycle Time} \\ &= \frac{\text{CPU Clock Cycles}}{\text{Clock Rate}}\end{aligned}$$

- Performance improved by
 - Reducing number of clock cycles
 - Increasing clock rate
 - Hardware designer must often trade off clock rate against cycle count

CPU Time

- Computer A run a program in 10 seconds with a 2 GHz clock. We have to design a computer B such that it can run the same program within 6 seconds. Determine the clock rate for computer B. Assume that due to increase in clock cycle rate , CPU design of computer B is affected, and it requires 1.2 times as many clock cycles as computer A for execution this program.

CPU Time Example

- Computer A: 2GHz clock, 10s CPU time
- Designing Computer B such that:
 - Aim for 6s CPU time
 - Can do faster clock, but causes $1.2 \times$ clock cycles
- How fast must Computer B clock be?

$$\text{Clock Rate}_B = \frac{\text{Clock Cycles}_B}{\text{CPU Time}_B} = \frac{1.2 \times \text{Clock Cycles}_A}{6s}$$

$$\begin{aligned}\text{Clock Cycles}_A &= \text{CPU Time}_A \times \text{Clock Rate}_A \\ &= 10s \times 2\text{GHz} = 20 \times 10^9\end{aligned}$$

$$\text{Clock Rate}_B = \frac{1.2 \times 20 \times 10^9}{6s} = \frac{24 \times 10^9}{6s} = 4\text{GHz}$$

Instruction Performance

- The computer had to execute the instructions to run the program.
- The execution time must depend on the number of instructions in a program.

Instruction Count and CPI

$\text{Clock Cycles} = \text{Instruction Count} \times \text{Cycles per Instruction}$

$\text{CPU Time} = \text{Instruction Count} \times \text{CPI} \times \text{Clock Cycle Time}$

$$= \frac{\text{Instruction Count} \times \text{CPI}}{\text{Clock Rate}}$$

- Instruction Count for a program
 - Determined by program, ISA and compiler
- Average cycles per instruction
 - Determined by CPU hardware

CPI Example

- Computer A – Cycle Time = 250ps, CPI = 2.0
- Computer B – Cycle Time = 500ps, CPI = 1.2
- Same ISA
- Which is faster? By how much?

CPI Example

- Computer A: Cycle Time = 250ps, CPI = 2.0
- Computer B: Cycle Time = 500ps, CPI = 1.2
- Same ISA
- Which is faster, and by how much?

$$\begin{aligned}\text{CPU Time}_A &= \text{Instruction Count} \times \text{CPI}_A \times \text{Cycle Time}_A \\ &= 1 \times 2.0 \times 250\text{ps} = 1 \times 500\text{ps} \end{aligned}$$

A is faster...

$$\begin{aligned}\text{CPU Time}_B &= \text{Instruction Count} \times \text{CPI}_B \times \text{Cycle Time}_B \\ &= 1 \times 1.2 \times 500\text{ps} = 1 \times 600\text{ps} \end{aligned}$$

$$\frac{\text{CPU Time}_B}{\text{CPU Time}_A} = \frac{1 \times 600\text{ps}}{1 \times 500\text{ps}} = 1.2$$

...by this much

CPI in More Detail

- If different instruction classes take different numbers of cycles

$$\text{Clock Cycles} = \sum_{i=1}^n (\text{CPI}_i \times \text{Instruction Count}_i)$$

- Weighted average CPI

$$\text{CPI} = \frac{\text{Clock Cycles}}{\text{Instruction Count}} = \sum_{i=1}^n \left(\text{CPI}_i \times \underbrace{\frac{\text{Instruction Count}_i}{\text{Instruction Count}}}_{\text{Relative frequency}} \right)$$

CPI Example

Comparing Code Segments

A compiler designer is trying to decide between two code sequences for a computer. The hardware designers have supplied the following facts:

	CPI for each instruction class		
	A	B	C
CPI	1	2	3

For a particular high-level language statement, the compiler writer is considering two code sequences that require the following instruction counts:

Code sequence	Instruction counts for each instruction class		
	A	B	C
1	2	1	2
2	4	1	1

Which code sequence executes the most instructions? Which will be faster? What is the CPI for each sequence?

CPI Example

- Alternative compiled code sequences using instructions in classes A, B, C

Class	A	B	C
CPI for class	1	2	3
IC in sequence 1	2	1	2
IC in sequence 2	4	1	1

- Sequence 1: IC = 5

- Clock Cycles
 $= 2 \times 1 + 1 \times 2 + 2 \times 3$
 $= 10$
- Avg. CPI = $10/5 = 2.0$

- Sequence 2: IC = 6

- Clock Cycles
 $= 4 \times 1 + 1 \times 2 + 1 \times 3$
 $= 9$
- Avg. CPI = $9/6 = 1.5$

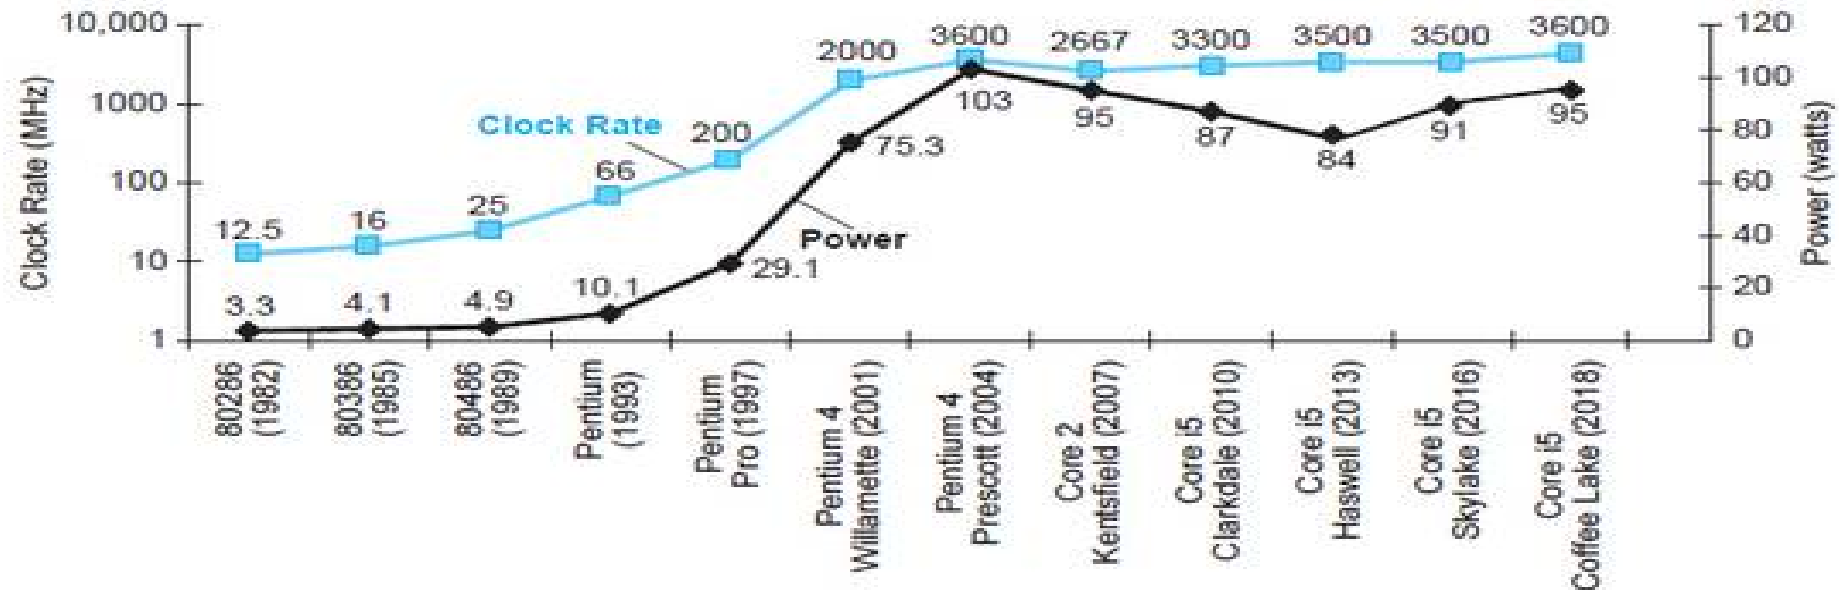
Performance Summary

The BIG Picture

$$\text{CPU Time} = \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Clock cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Clock cycle}}$$

- Performance depends on
 - Algorithm: affects IC, possibly CPI
 - Programming language: affects IC, CPI
 - Compiler: affects IC, CPI
 - Instruction set architecture: affects IC, CPI

Power Trends



- In CMOS IC technology

$$\text{Power} = \text{Capacitive load} \times \text{Voltage}^2 \times \text{Frequency}$$

×30

5V → 1V

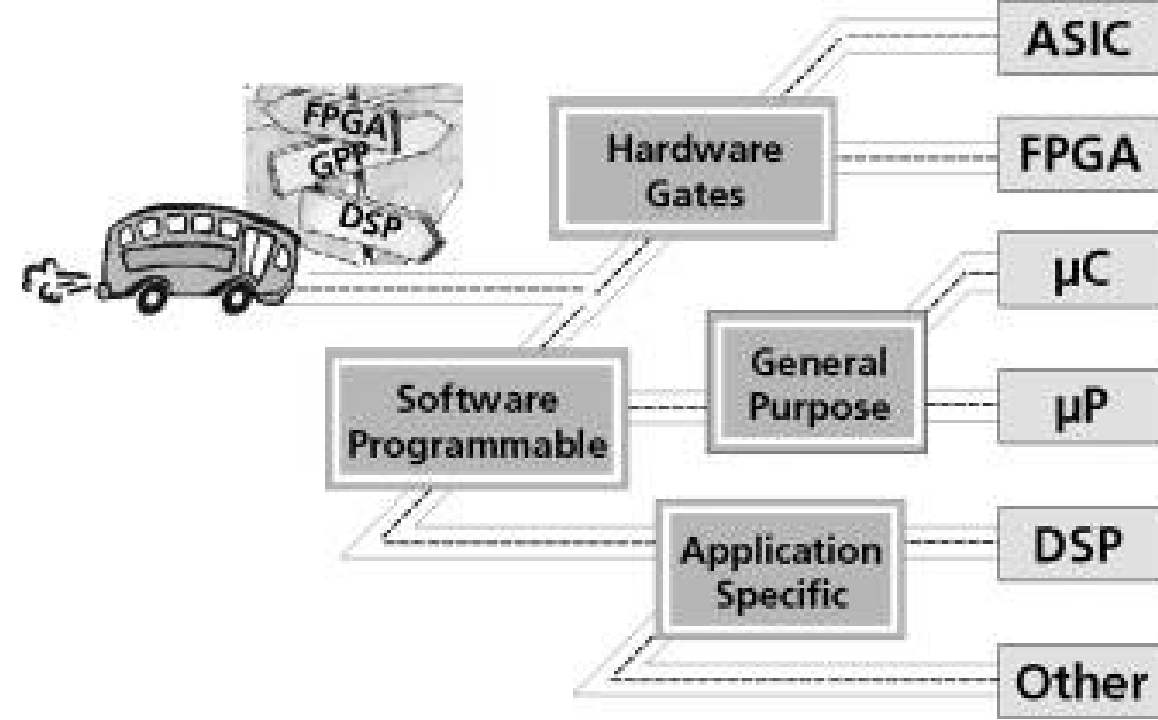
×1000

Topics

1. Basic Processor Architecture
- 2. Different Types of Processor Architectures**
3. RISC-V Processor Architecture
4. RISC-V Instruction Set Architecture
5. Programming RISC-V using assembly language

Processors Types

- General Purpose Processor
- Digital Signal Processor
- Vector Processor
- Application specific Processor



Flynn Taxonomy

- The matrix below defines the 4 possible classifications according to Flynn

S I S D Single Instruction, Single Data	S I M D Single Instruction, Multiple Data
M I S D Multiple Instruction, Single Data	M I M D Multiple Instruction, Multiple Data

Types of Processor ISA

Reduced Instruction Set Computing (RISC) vs Complex Instruction Set Computing (CISC)

Aspect	RISC	CISC
Instructions Per Cycle	Small and fixed length	Large and variable length
Instruction Complexity	Simple and standardised	Complex and versatile
Instruction Execution	Single clock cycle	Several clock cycles
RAM Usage	Heavy use of RAM	More efficient use of RAM
Memory	Increased memory usage to store instructions	Memory efficient coding
Cost	Cheaper than CISC	Higher

RISC vs CISC

The RISC approach has several advantages over CISC:

- **Simplifies Hardware Implementation:** It simplifies the hardware implementation of the processor, as fewer instructions need to be decoded and executed. This can lead to faster execution times and lower power consumption.
- **Higher Instruction Level Parallelism:** RISC processors typically have a higher instruction-level parallelism, allowing them to execute multiple instructions simultaneously, which can further improve performance.
- **Simplicity:** The simplicity of the RISC instruction set makes it easier to develop compilers and other software tools that can generate efficient code for the processor.

RISC vs CISC

RISC is a processor design philosophy that emphasizes simplicity and efficiency by using a small set of simple and general-purpose instructions.

- The ***complex instruction set computing*** (CISC), employs a larger set of more complex instructions that can perform multiple operations in a single instruction.
- RISC architectures prioritize simplicity and execute one instruction per clock cycle, resulting in streamlined designs and efficient decoding.
- CISC architectures, on the other hand, employ complex instructions capable of performing multiple actions but may require several clock cycles for execution. Both the CPUs aim to enhance CPU performance.

Single-purpose processors

Digital circuit designed to execute exactly one program
a.k.a. coprocessor, accelerator or peripheral

Features

Contains only the components needed to execute a single program

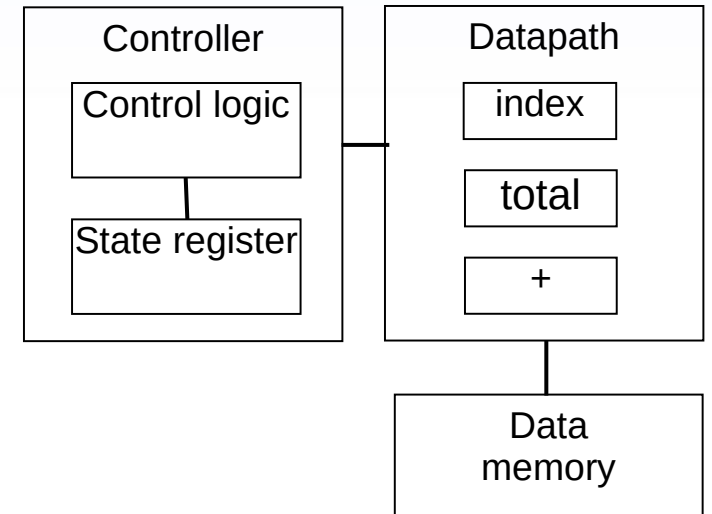
No program memory

Benefits

Fast

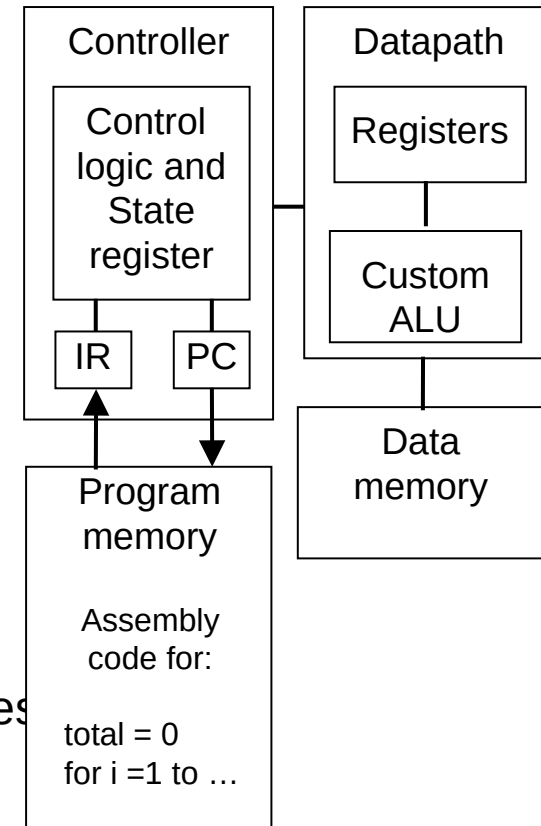
Low power

Small size



Embedded System Processor Architecture

- **Reduced Instruction Set Computing (RISC):**
 - › Common architectures: ARM, RISC-V.
 - › Simple, efficient instruction set optimized for low power and high performance.
- **System on Chip (SoC):**
 - › Frequently used in embedded systems.
 - › Integrates CPU, memory, peripherals, and other components on a single chip.
- **Microcontroller Units (MCUs):**
 - › Often used in simpler embedded applications.
 - › Includes integrated peripherals like ADCs, DACs, timers, and communication interfaces.
- **Real-Time Capabilities:**
 - › Designed for deterministic performance and real-time operating system (RTOS) support.
- **Low Power Consumption:**
 - › Architectures and components optimized for minimal power usage.
- **Integrated Analog and Digital Peripherals:**
 - › Features like GPIOs, serial communication interfaces, and specialized hardware accelerators.



Digital Signal Processor

Specialized Instruction Set:

Optimized for mathematical operations like multiply-accumulate (MAC).

Single-cycle multiply and MAC instructions.

Harvard Architecture:

Separate program and data memories to allow simultaneous access and increase throughput.

Specialized Data Path:

Multiple data buses and address buses.

Dedicated hardware for specific tasks such as FFT (Fast Fourier Transform) and filters.

High-Performance ALUs:

Multiple arithmetic logic units (ALUs) to perform parallel operations.

Support for fixed-point and floating-point arithmetic.

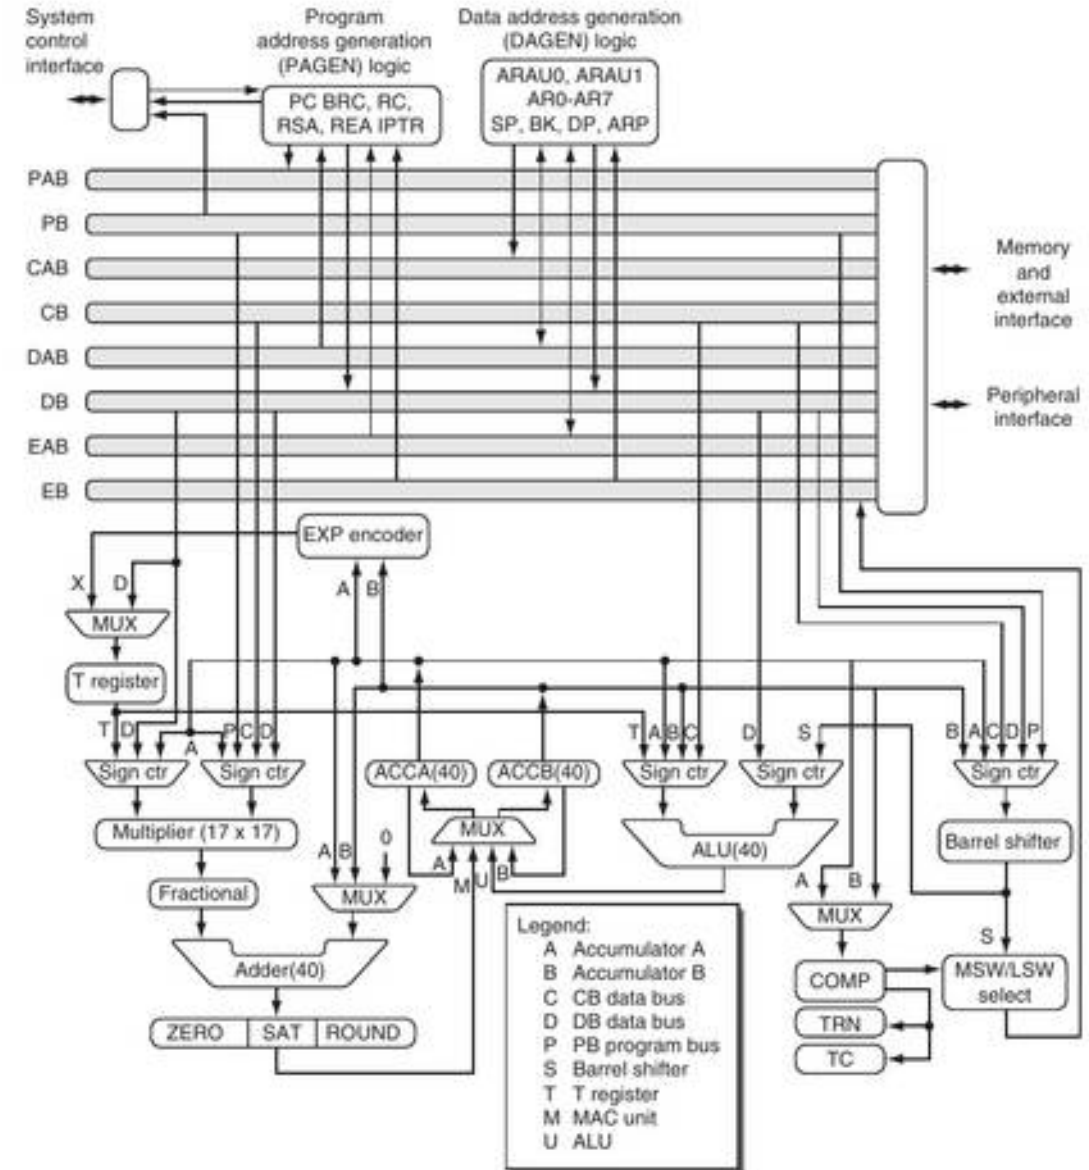
Circular and Bit-Reversed Addressing:

Efficiently manage circular buffers and data structures used in signal processing.

Low-Latency Memory Access:

On-chip RAM with very low access latency.

Multi-level cache hierarchy optimized for predictable access patterns.



General-Purpose Processor (GPP) Architecture

- **Complex Instruction Set Computing (CISC):**

- › Common architecture: x86.
- › Rich instruction set with complex instructions.
- › Often integrates many features directly in hardware.

- **Multi-Core and Hyper-Threading:**

- › Multiple cores for parallel processing.
- › Hyper-threading for improved performance through parallel execution within each core.

- **Large Cache Hierarchy:**

- › Multiple levels of cache (L1, L2, L3) to reduce latency and increase speed.
- › Advanced Branch Prediction and Speculative Execution:
- › Techniques to predict instruction paths and execute ahead to improve performance.

- **Integrated Memory Management Unit (MMU):**

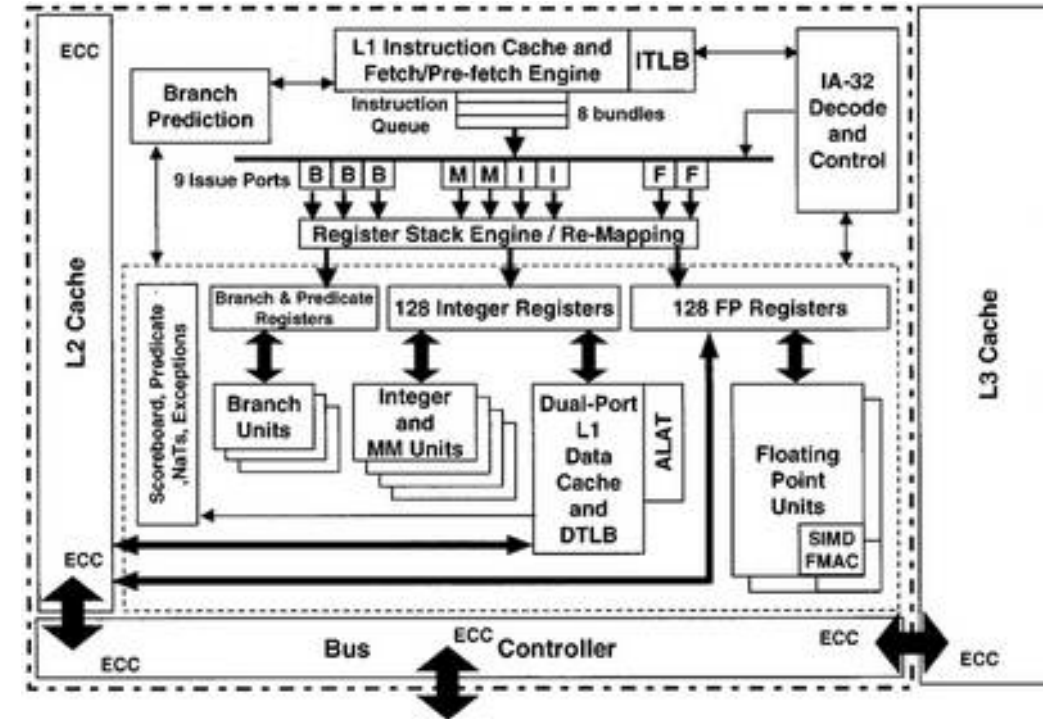
- › Manages virtual memory, enabling sophisticated operating system features.

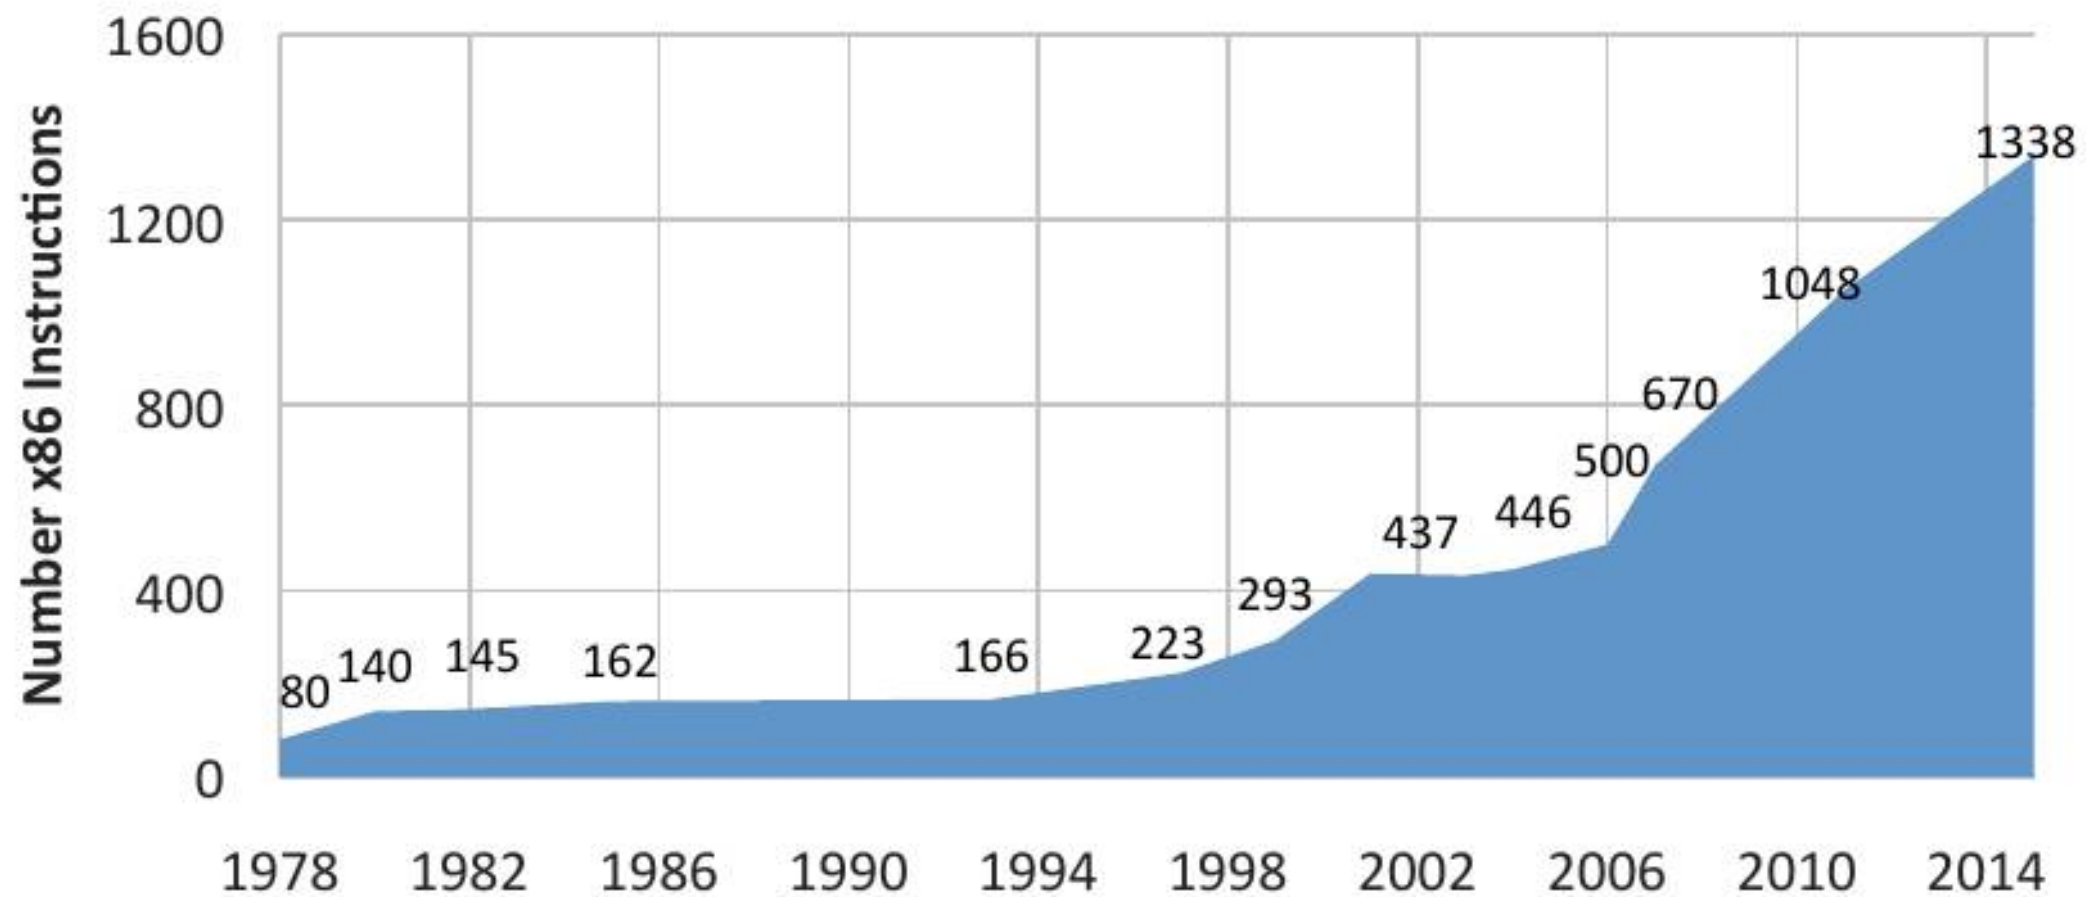
- **High-Speed Interconnects:**

- › Fast communication between CPU, memory, and peripherals.

- **Graphics Processing Unit (GPU) Integration:**

- › Some GPPs include integrated GPUs for handling graphics processing tasks.





Topics

1. Basic Processor Architecture
2. Different Types of Processor Architectures
- 3. RISC-V Processor Architecture**
4. RISC-V Instruction Set Architecture
5. Programming RISC-V using assembly language



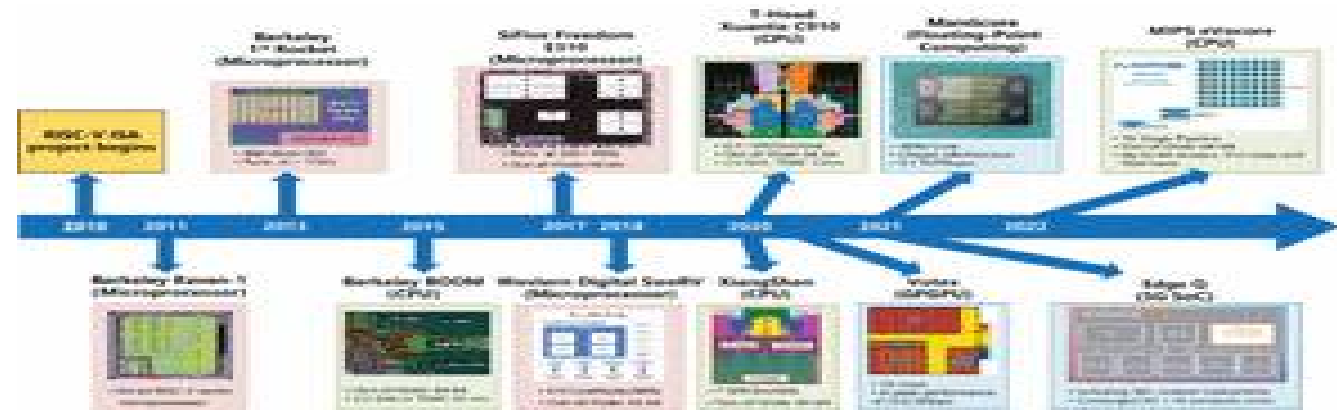
RISC-V Processor Architecture

The RISC-V (pronounced as risk-five) architecture is an open-source instruction set architecture (ISA) implementation of reduced instruction set computing RISC.

RISC-V is open-hardware architecture, its open source allows anyone to utilize the ISA.

History of RISC-V

- Prof. Krste Asanović and graduate students Yunsup Lee and Andrew Waterman started the RISC-V instruction set in May 2010 as part of the [Parallel Computing Laboratory](#) (Par Lab) at UC Berkeley, of which Prof. David Patterson was Director.
- No patents were filed related to RISC-V in any of these projects, as the RISC-V ISA itself does not represent any new technology.
- RISC processor implementations—including some based on other open ISA standards—are widely available from various vendors worldwide.



Processor Architecture

Base Instruction Set

RV32I	Base Integer Instruction Set, 32-bit
RV32E	Base Integer Instruction Set (embedded), 32-bit
RV64I	Base Integer Instruction Set, 64-bit

Extension:

Name	Description
M	Standard Extension for Integer Multiplication and Division
A	Standard Extension for Atomic Instructions
F	Standard Extension for Single-Precision Floating-Point
D	Standard Extension for Double-Precision Floating-Point
Zicsr	Control and Status Register (CSR) Instructions
Zifencei	Instruction-Fetch Fence
G	Shorthand for the IMAFDZicsr_Zifencei base and extensions
C	Standard Extension for Compressed Instructions

Base and Extension of RISC-V

- Four base integer ISAs
 - RV32E, RV32I, RV64I, RV128I
 - RV32E is 16-register subset of RV32I
 - Only <50 hardware instructions needed for base
- Standard extensions
 - M: Integer multiply/divide
 - A: Atomic memory operations (AMOs + LR/SC)
 - F: Single-precision floating-point
 - D: Double-precision floating-point
 - G = IMAFD, "General-purpose" ISA
 - Q: Quad-precision floating-point
- All the above are a fairly standard RISC encoding in a fixed 32-bit instruction format
- Above user-level ISA components frozen in 2014
 - Supported forever after



RISCV: Registers and Mapping

RISC-V uses a memory-mapped I/O architecture, which means that input and output operations, memory access, and peripheral access are all performed using the same load and store instructions.

This unified approach simplifies the instruction set and enhances the flexibility and efficiency of the architecture.

There are two basic types of instructions:

- Instructions that either load memory into registers or store data from registers into memory
- Instructions that perform arithmetical or logical operations between two registers

Why RISC-V

Open Hardware: Allowing anyone to design, implement, and customize processors without restrictions, fostering innovation and collaboration within the community.

Royalty-Free: There are no licensing fees, reducing costs for developers and manufacturers.

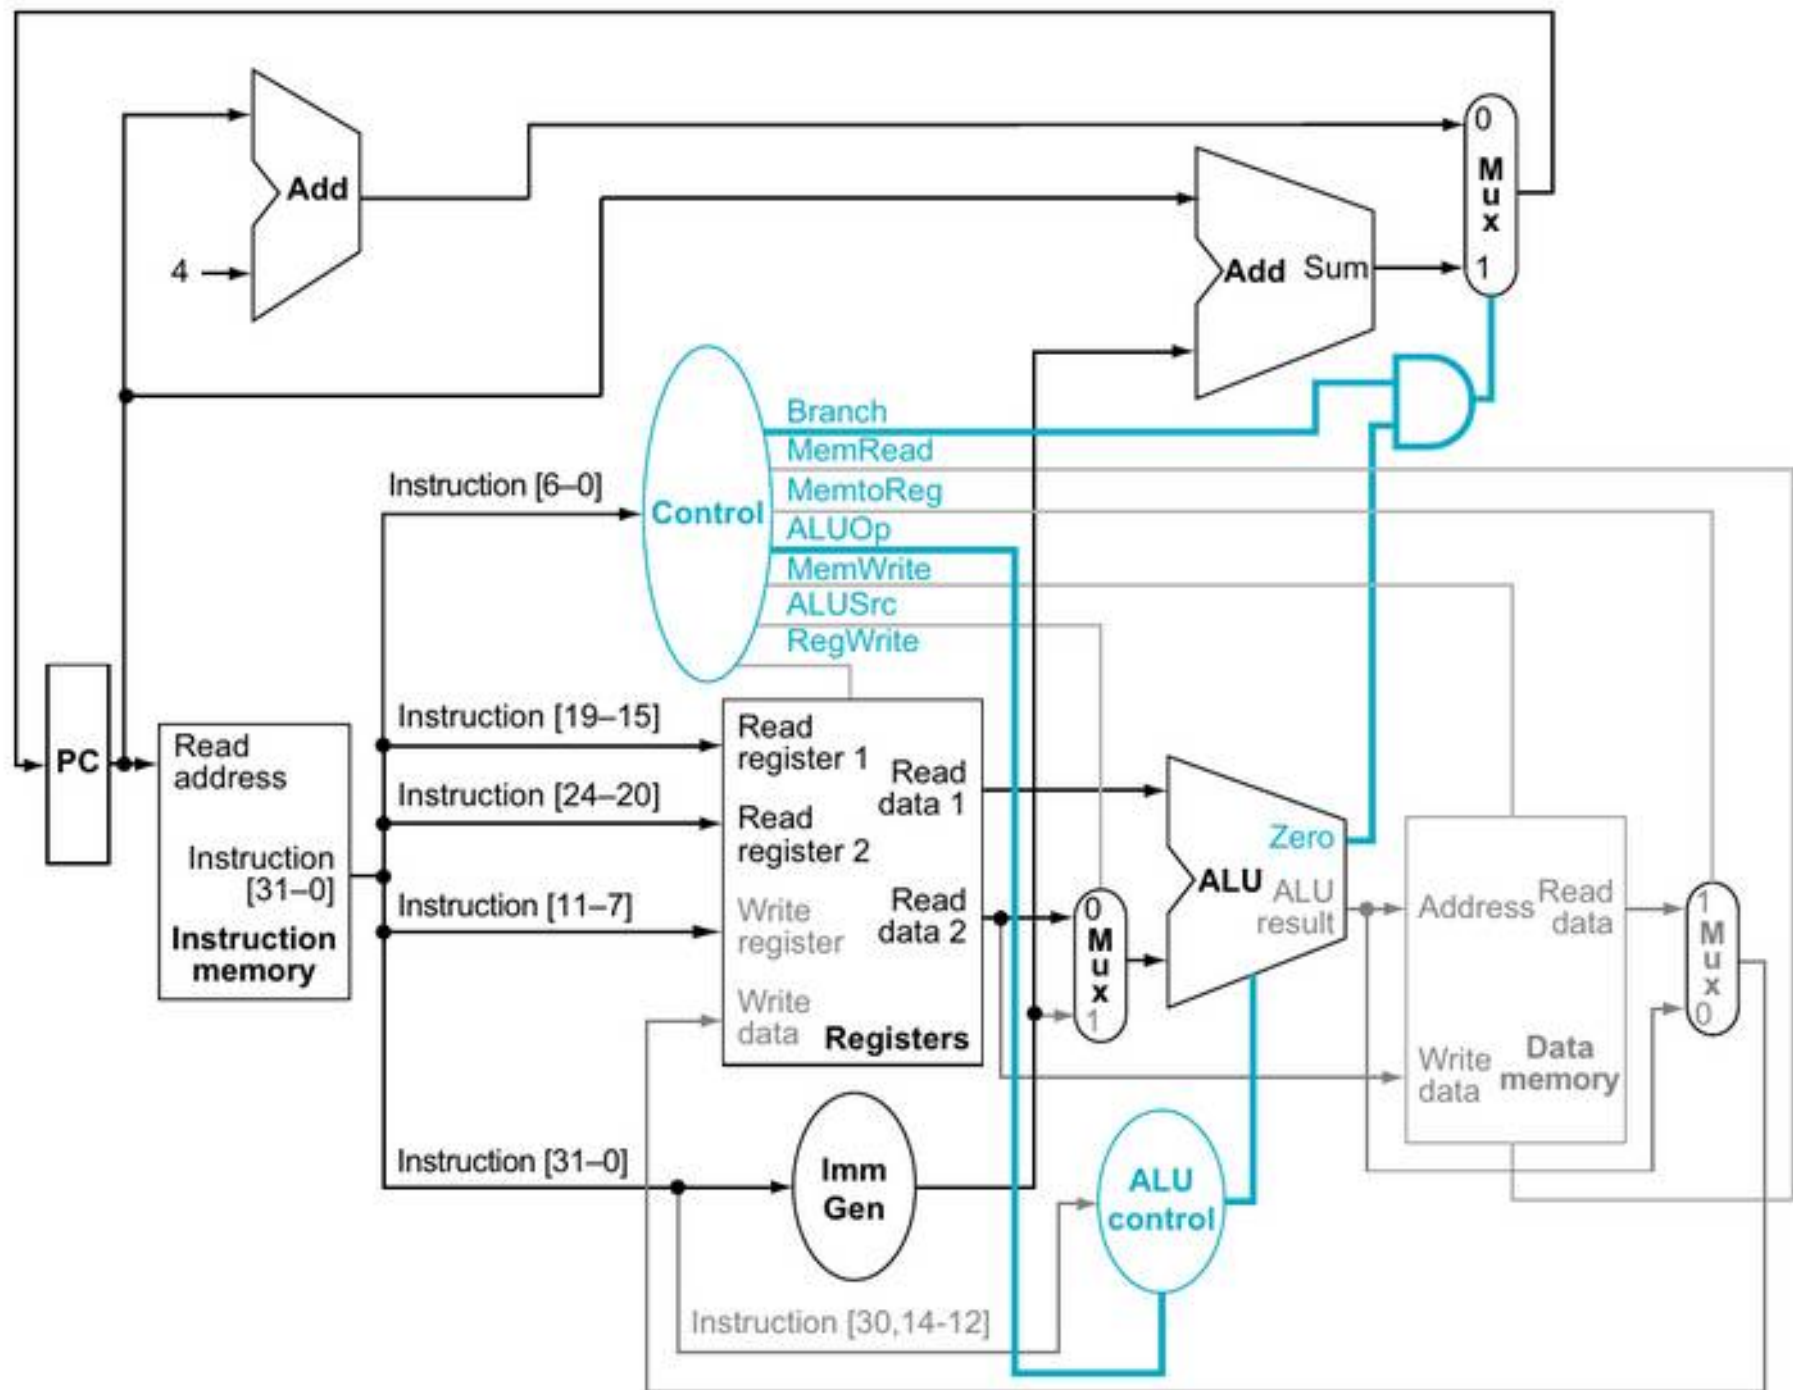
Security: Rigorous security analysis and the implementation of custom security features, enhancing trustworthiness.

ISA	Pages	Words	Hours to read	Weeks to read
RISC-V	236	76,702	6	0.2
ARM-32	2736	895,032	79	1.9
x86-32	2198	2,186,259	182	4.5

	ARM-32 (1986)	MIPS-32 (1986)	x86-32 (1978)	Lessons learned RV32I (2011)
Cost	Integer multiply mandatory	Integer multiply and divide mandatory	8-bit and 16-bit operations. Integer multiply and divide mandatory	No 8-bit and 16-bit operations. Integer multiply and divide optional (RV32M)
Simplicity	No zero register. Conditional instruction execution. Complex data address modes. Stack instructions (push/pop). Shift-option for arithmetic/logic instructions	Zero- and sign-extended immediates. Some arithmetic instructions can cause overflow traps	No zero register. Complex procedure call/return instructions (enter/leave). Stack instructions (push/pop). Complex data address modes. Loop instructions	Register x0 dedicated to 0. Immediates only sign-extended. One data addressing mode. No conditional execution. No complex call/return or stack instructions. No traps for arithmetic overflow. Separate shift instructions
Performance	Condition codes for branches. Source and destination registers vary in instruction format. Load multiple. Computed immediates. PC a general purpose register	Source and destination registers vary in instruction format.	Condition codes for branches. At most 2 registers per instruction	Compare and branch instructions (no condition codes). 3 registers per instruction. No load multiple. Source and destination registers fixed in instruction format. Constant immediates. PC not a general purpose register
Isolate architecture from implementation	Exposes the pipeline length when writing the PC as a general purpose register	Delayed branch. Delayed load. HI and LO registers just for multiply and divide	Registers not general purpose (AX, CX, DX, DI, SI have unique uses)	No delayed branch. No delayed load. General purpose registers
Room for growth	Limited available opcode space	Limited available opcode space		Generous available opcode space
Program size	Only 32-bit instructions (+Thumb-2 as separate ISA)	Only 32-bit instructions (+microMIPS as separate ISA)	Byte-variable instructions, but poor choices	32-bit instructions + 16-bit RV32C extension
Ease of programming / compiling / linking	Only 15 registers. Aligned data in memory. Irregular data address modes. Inconsistent performance counters	Aligned data in memory. Inconsistent performance counters	Only 8 registers. No PC-relative data addressing. Inconsistent performance counters	31 registers. Data can be unaligned. PC-relative data addressing. Symmetric data address mode. Performance counters defined in architecture

Types of RISC-V Processor Architectures

- RISC-V provides a detailed, open Instruction Set Architecture (ISA), which serves as a blueprint for designing processors architecture.
- Single-Cycle Architecture:
- Multi-Cycle Architecture:
- **Pipelined Architecture:**
- Superscalar Architecture:
- Out-of-Order Execution:
- Very Long Instruction Word (VLIW) Architecture:
- Vector Processing Architecture:
- Custom Instruction Set Extensions:



Defining/Designing RISC-V Processor Architecture

- Fetch: Retrieve the instruction from memory.
- Decode: Interpret the instruction and prepare operands.
- Execute: Perform the computation or operation (ALU operations, branches).
- Memory: Access memory for load/store operations.
- Writeback: Write the result to the register file or memory.

5 Stages of Processor Arch

- **Fetch Unit**

Function: Retrieves instructions from memory.

PC Usage: The PC holds the address of the next instruction to be fetched. After fetching an instruction, the PC is typically incremented to point to the next instruction address.

Example: If the starting address of the first instruction is 0x8000000, the Fetch Unit will fetch the instruction from address 0x8000000 initially.

- **Decode Unit**

Function: Interprets the fetched instruction to determine its operation and operands.

Memory Access: Decodes memory addresses and identifies whether they are for RAM, ROM, or I/O devices. It also decodes which registers are involved.

ALU: Determines the type of ALU operation required (e.g., addition, subtraction) and prepares operands for execution.

Example: Decodes an instruction to add two registers and prepare the operands for the ALU.

Execute Unit

Function: Performs the arithmetic or logical operations as specified by the instruction.

ALU: Executes ALU operations (e.g., addition, subtraction) using the operands provided by the Decode Unit.

Memory Access: Computes effective addresses for load/store operations.

Example: Executes an addition operation on two registers or calculates the address for a load instruction.

Memory Unit

Function: Accesses memory or I/O based on the address computed in the Execute stage.

Memory Access: Performs read/write operations to RAM or memory-mapped I/O devices based on the effective address.

Example: Reads data from address 0x00002000 in RAM or writes data to a memory-mapped I/O device at 0x20000000.

Write Back Unit

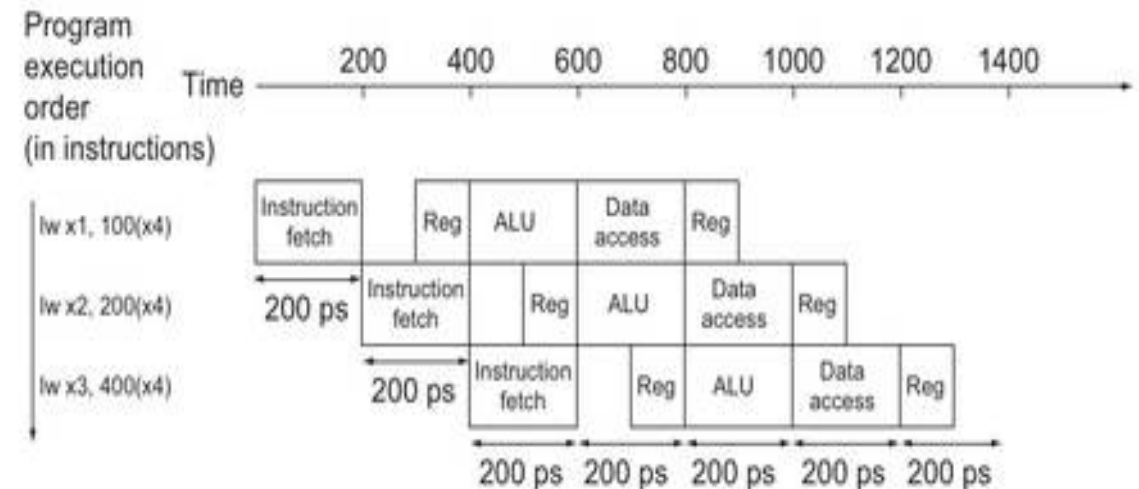
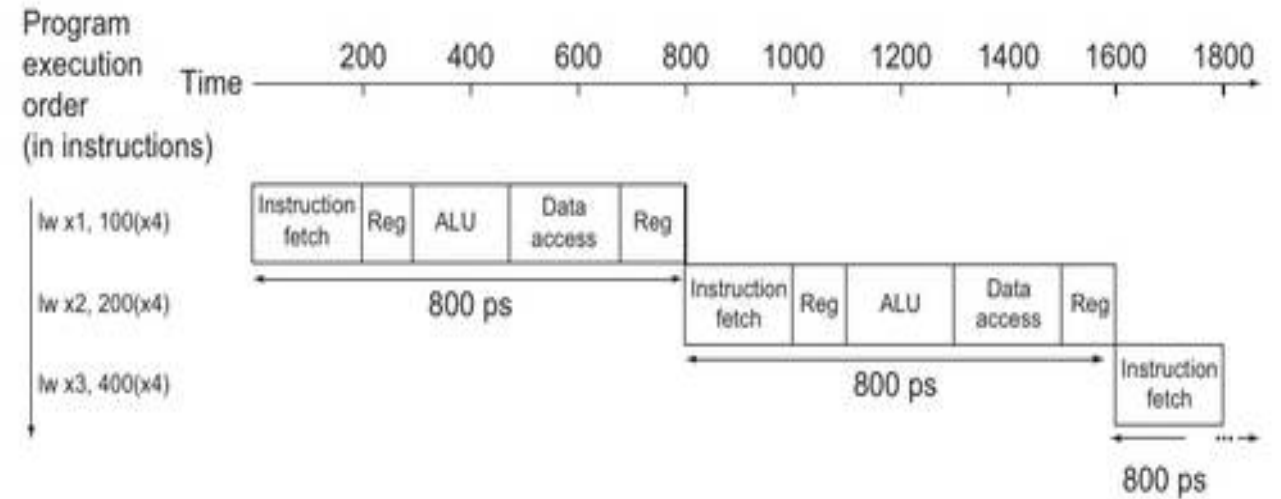
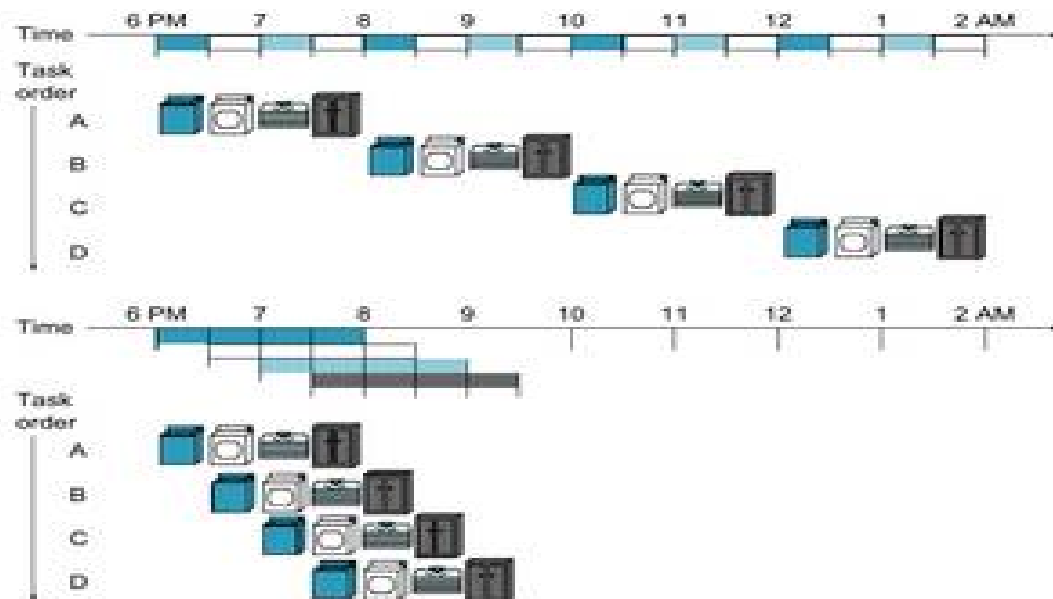
Function: Writes the result of computations or memory accesses back to the register file or memory.

Memory Access: Updates the register file with results from the Memory Unit or ALU operations.

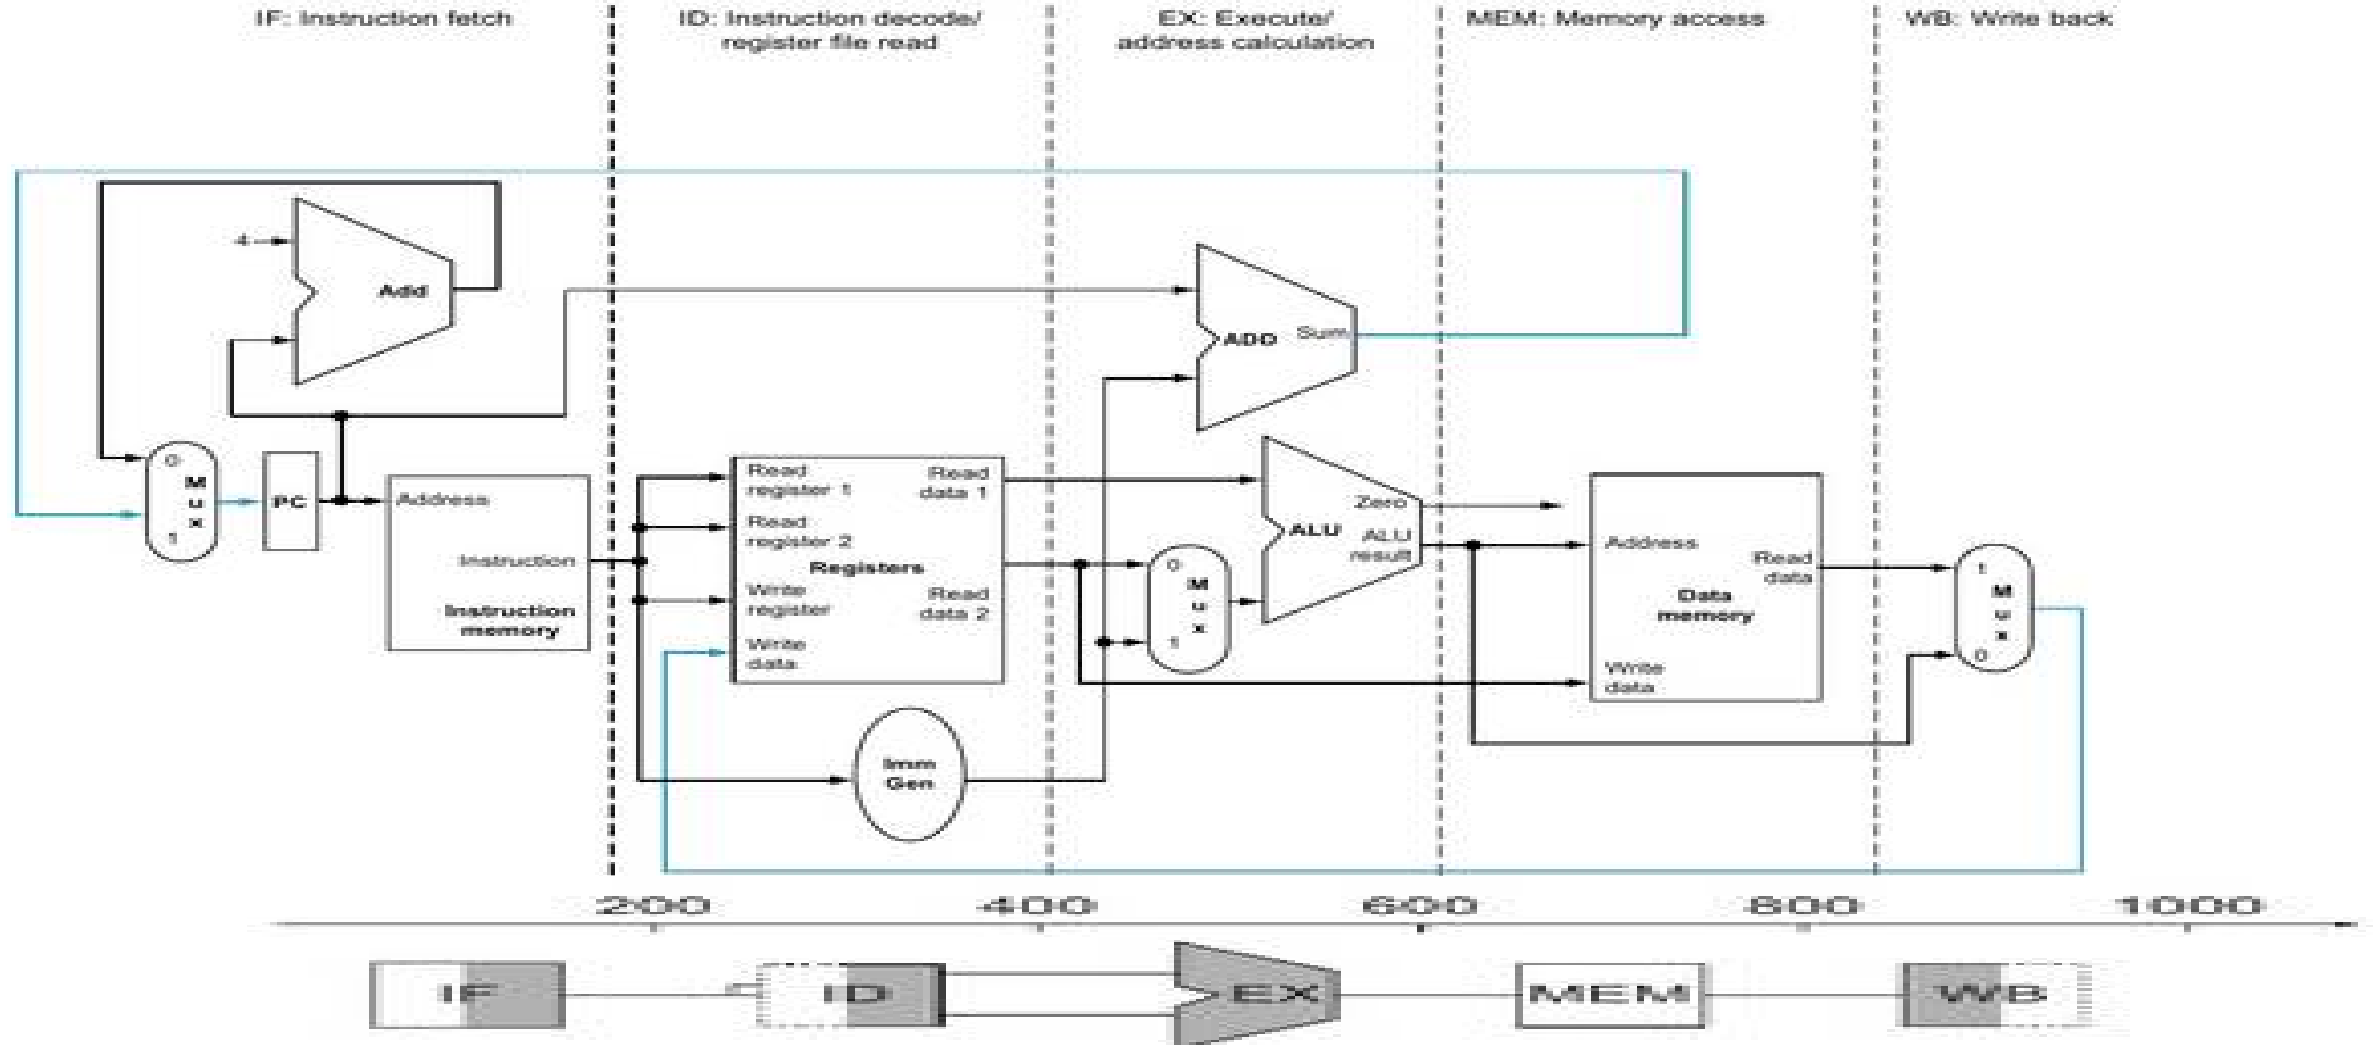
Example: Writes the result of an addition operation back to a register or stores data retrieved from memory to a register.

Pipe-lined VS Single Cycle Processor Architecture

1. Fetch instruction from memory.
2. Read registers and decode the instruction.
3. Execute the operation or calculate an address.
4. Access an operand in data memory (if necessary).
5. Write the result into a register (if necessary).



5 Stage Pipelined Processor Architecture



Topics

1. Basic Processor Architecture
2. Different Types of Processor Architectures
3. RISC-V Processor Architecture
- 4. RISC-V Instruction Set Architecture**
5. Programming RISC-V using assembly language



Instruction Set Architecture is like the language that a computer's hardware understands. I

- ISA is a set of instructions that tells the computer's processor what to do. These instructions are basic operations like adding two numbers, moving data from one place to another, or jumping to a different part of a program.
- The ISA defines the "rules" of computer architecture or the "language" that the CPU uses to execute instructions.
- It determines what the processor can do. If a processor supports a particular ISA, it can execute any program written in that ISA.
- It also ensures that different programs can run on different computers as long as those computers understand the same ISA.
- **It's like setting the vocabulary and grammar for a conversation between software and hardware.**
- **It is is the language that a computer's hardware understands. I**

Commercial ISA

- Commercial ISAs are proprietary.
- Commercial ISAs are only popular in certain market domains.
- Commercial ISAs come and go.
- Popular commercial ISAs are complex.
- Commercial ISAs alone are not enough to bring up applications.
- Popular commercial ISAs were not designed for extensibility.
- A modified commercial ISA is a new ISA.

OpenRISC

- *OpenRISC has condition codes and branch delay slots, which complicate higher performance implementations.*
- *OpenRISC uses a fixed 32-bit encoding and 16-bit immediates, which precludes a denser instruction encoding and limits space for later expansion of the ISA.*
- *OpenRISC does not support the 2008 revision to the IEEE 754 floating-point standard.*
- *The OpenRISC 64-bit design had not been completed when we began.*

RISCV Instructions Set

- RISC-V (Reduced Instruction Set Computing V) is an open standard instruction set architecture (ISA) that is designed to be scalable and extensible. The number of instructions in RISC-V can vary based on the specific subset or extensions of the ISA being used. Here's a breakdown of the primary RISC-V instruction sets and their respective instruction counts:
- Base ISA:
 - › RV32I (32-bit Integer): The base integer instruction set for 32-bit processors includes approximately 47 instructions.
 - › RV64I (64-bit Integer): The base integer instruction set for 64-bit processors extends RV32I and includes a few additional instructions specific to 64 bit operations.

Instruction Extensions

- Standard Extensions:

- M (Multiply/Divide): Adds multiply and divide instructions.
- A (Atomic): Adds atomic instructions for synchronization.
- F (Single-Precision Floating-Point): Adds single-precision floating-point instructions.
- D (Double-Precision Floating-Point): Adds double-precision floating-point instructions.
- Q (Quad-Precision Floating-Point): Adds quad-precision floating-point instructions.
- C (Compressed): Adds compressed instructions to reduce code size.

- Other Extensions:

- B (Bit-Manipulation): Adds instructions for bit manipulation.
- V (Vector): Adds vector processing instructions.
- P (Packed-SIMD): Adds packed SIMD instructions.
- Z (Various small extensions): These include specific sets of instructions like Zifencei for instruction-fence or Zicsr for control and status registers.

Basic RISC-V Processor

- The 47 standard instructions in RV32I include:

- Arithmetic Instructions: ADD, SUB, MUL, etc.
- Logical Instructions: AND, OR, XOR, etc.
- Immediate Instructions: ADDI, ORI, XORI, etc.
- Load Instructions: LB, LH, LW, etc.
- Store Instructions: SB, SH, SW, etc.
- Branch Instructions: BEQ, BNE, BLT, etc.
- Jumps: JAL, JALR
- System Instructions: ECALL, EBREAK
- Other Instructions: NOP, AUIPC, LUI, etc.

31	25	24	19	15	14	12	11	7	6	0	
imm[31:12]							rd	0110111		U lui	
imm[31:12]							rd	0010111		U auipc	
imm[20 10:1 11 19:12]							rd	1101111		J jal	
imm[11:0]				rs1	000		rd	1100111		I jalr	
imm[12 10:5]		rs2		rs1	000		imm[4:1 11]	1100011		B beq	
imm[12 10:5]		rs2		rs1	001		imm[4:1 11]	1100011		B bne	
imm[12 10:5]		rs2		rs1	100		imm[4:1 11]	1100011		B blt	
imm[12 10:5]		rs2		rs1	101		imm[4:1 11]	1100011		B bge	
imm[12 10:5]		rs2		rs1	110		imm[4:1 11]	1100011		B bltu	
imm[12 10:5]		rs2		rs1	111		imm[4:1 11]	1100011		B bgeu	
imm[11:0]				rs1	000		rd	0000011		I lb	
imm[11:0]				rs1	001		rd	0000011		I lh	
imm[11:0]				rs1	010		rd	0000011		I lw	
imm[11:0]				rs1	100		rd	0000011		I lbu	
imm[11:0]				rs1	101		rd	0000011		I lhu	
imm[11:5]		rs2		rs1	000		imm[4:0]	0100011		S sb	
imm[11:5]		rs2		rs1	001		imm[4:0]	0100011		S sh	
imm[11:5]		rs2		rs1	010		imm[4:0]	0100011		S sw	
imm[11:0]				rs1	000		rd	0010011		I addi	
imm[11:0]				rs1	010		rd	0010011		I slti	
imm[11:0]				rs1	011		rd	0010011		I sltiu	
imm[11:0]				rs1	100		rd	0010011		I xori	
imm[11:0]				rs1	110		rd	0010011		I ori	
imm[11:0]				rs1	111		rd	0010011		I andi	
0000000		shamt		rs1	001		rd	0010011		I slli	
0000000		shamt		rs1	101		rd	0010011		I srli	
0100000		shamt		rs1	101		rd	0010011		I srai	
0000000		rs2		rs1	000		rd	0110011		R add	
0100000		rs2		rs1	000		rd	0110011		R sub	
0000000		rs2		rs1	001		rd	0110011		R sll	
0000000		rs2		rs1	010		rd	0110011		R slt	
0000000		rs2		rs1	011		rd	0110011		R sltu	
0000000		rs2		rs1	100		rd	0110011		R xor	
0000000		rs2		rs1	101		rd	0110011		R srl	
0100000		rs2		rs1	101		rd	0110011		R sra	
0000000		rs2		rs1	110		rd	0110011		R or	
0000000		rs2		rs1	111		rd	0110011		R and	
0000		pred		succ		00000	000	00000	0001111	I fence	
0000		0000		0000		00000	001	00000	0001111	I fence.i	
0000000000000				00000		000	00000	1110011		I ecall	
00000000000001				00000		000	00000	1110011		I ebREAK	
csr				rs1		001	rd	1110011		I csrrw	
csr				rs1		010	rd	1110011		I csrrs	
csr				rs1		011	rd	1110011		I csrrc	
csr				zimm		101	rd	1110011		I csrrwi	
csr				zimm		110	rd	1110011		I csrrsi	
csr				zimm		111	rd	1110011		I csrrci	

Types of RISC-V ISA

RISC-V Instruction Set:

The RISC-V instruction set is a collection of instructions that define the operations a RISC-V processor can perform.

These instructions are designed to be simple, efficient, and easily extensible, allowing for a high degree of customization and optimization.

Instruction Types:

1. **R-Type (Register Type):** Used for register-register arithmetic and logical operations.
2. **I- Type (Instruction Type):** Used for immediate arithmetic, load instructions, and register-immediate operations.
3. **S-Type (Store Type):** Used for store instructions.
4. **U-Type (Upper Immediate Type) :** Used for upper immediate instructions
5. **B-Type (Branch Type) :** Used for conditional branch instructions.
6. **J-Type (Jump Type) :** Used for jump instructions like JAL.
7. **F-Type (Floating-Point) Instructions**
8. **A-Type (Atomic) Instructions**
9. **C-Type (Compressed) Instructions**

Registers

- Total Registers: 32 general-purpose registers, additional special-purpose and control registers.
 - › General Purpose Registers: x0 to x31, with specific roles for some registers.
 - › Special Purpose Registers: Includes PC, SP, GP, TP.
 - Program Counter (PC): Holds the address of the current instruction being executed.
 - Instruction Register (IR): Holds the current instruction being executed (in some implementations).
 - Stack Pointer (SP): Points to the top of the stack.
 - Global Pointer (GP): Points to the global data region.
 - Thread Pointer (TP): Points to the thread-local storage.
 - › Control and Status Registers: Includes MSR, MEPC, MCAUSE, MSTATUS, MTVEC.
 - Machine Status Register (MSR): Controls machine-level status and configuration.
 - Machine Exception Program Counter (MEPC): Holds the address of the instruction where an exception occurred.
 - Machine Cause Register (MCAUSE): Contains information about the cause of the last exception.
 - Machine Status Register (MSTATUS): Holds the status of the machine, including interrupts and mode.
 - Machine Trap Vector Base Address Register (MTVEC): Base address for the trap vector.
 - › Floating-Point Registers: If included, f0 to f31 for floating-point operations.

Register name	Symbolic name	Description	Saved by
32 integer registers			
x0	Zero	Always zero	
x1	ra	Return address	Caller
x2	sp	Stack pointer	Callee
x3	gp	Global pointer	
x4	tp	Thread pointer	
x5	t0	Temporary / alternate return address	Caller
x6-7	t1-2	Temporary	Caller
x8	s0/fp	Saved register / frame pointer	Callee
x9	s1	Saved register	Callee
x10-11	a0-1	Function argument / return value	Caller
x12-17	a2-7	Function argument	Caller
x18-27	s2-11	Saved register	Callee
x28-31	t3-6	Temporary	Caller

RISC-V Instruction Format

31	27	26	25	24	20	19	15	14	12	11	7	6	0	
funct7				rs2		rs1		funct3		rd		opcode		R-type
imm[11:0]						rs1		funct3		rd		opcode		I-type
imm[11:5]				rs2		rs1		funct3		imm[4:0]		opcode		S-type
imm[12 10:5]				rs2		rs1		funct3		imm[4:1 11]		opcode		B-type
imm[31:12]										rd		opcode		U-type
imm[20 10:1 11 19:12]										rd		opcode		J-type

RISC-V Instructions and Formats

RV32I Base Integer Instructions

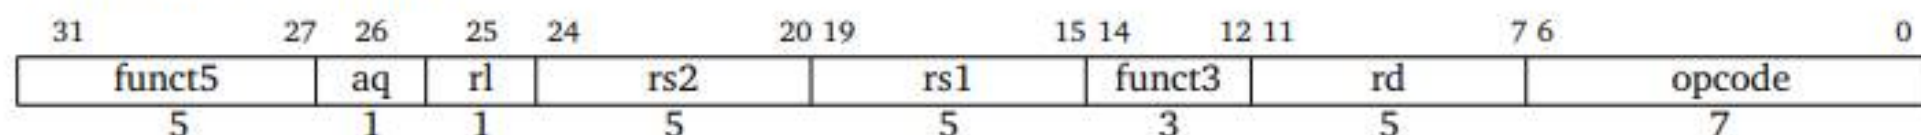
Inst	Name	FMT	Opcode	funct3	funct7	Description (C)	Note
add	ADD	R	0110011	0x0	0x00	$rd = rs1 + rs2$	msb-extends zero-extends
sub	SUB	R	0110011	0x0	0x20	$rd = rs1 - rs2$	
xor	XOR	R	0110011	0x4	0x00	$rd = rs1 \wedge rs2$	
or	OR	R	0110011	0x6	0x00	$rd = rs1 rs2$	
and	AND	R	0110011	0x7	0x00	$rd = rs1 \& rs2$	
sll	Shift Left Logical	R	0110011	0x1	0x00	$rd = rs1 \ll rs2$	
srl	Shift Right Logical	R	0110011	0x5	0x00	$rd = rs1 \gg rs2$	
sra	Shift Right Arith*	R	0110011	0x5	0x20	$rd = rs1 \gg rs2$	
slt	Set Less Than	R	0110011	0x2	0x00	$rd = (rs1 < rs2)?1:0$	
sltu	Set Less Than (U)	R	0110011	0x3	0x00	$rd = (rs1 < rs2)?1:0$	
addi	ADD Immediate	I	0010011	0x0	imm[5:11]=0x00 imm[5:11]=0x00 imm[5:11]=0x20	$rd = rs1 + imm$	msb-extends zero-extends
xori	XOR Immediate	I	0010011	0x4		$rd = rs1 \wedge imm$	
ori	OR Immediate	I	0010011	0x6		$rd = rs1 imm$	
andi	AND Immediate	I	0010011	0x7		$rd = rs1 \& imm$	
slli	Shift Left Logical Imm	I	0010011	0x1		$rd = rs1 \ll imm[0:4]$	
srli	Shift Right Logical Imm	I	0010011	0x5		$rd = rs1 \gg imm[0:4]$	
srai	Shift Right Arith Imm	I	0010011	0x5		$rd = rs1 \gg imm[0:4]$	
slti	Set Less Than Imm	I	0010011	0x2		$rd = (rs1 < imm)?1:0$	
sltiu	Set Less Than Imm (U)	I	0010011	0x3		$rd = (rs1 < imm)?1:0$	

sb	Store Byte	S	0100011	0x0		$M[rs1+imm][0:7] = rs2[0:7]$	
sh	Store Half	S	0100011	0x1		$M[rs1+imm][0:15] = rs2[0:15]$	
sw	Store Word	S	0100011	0x2		$M[rs1+imm][0:31] = rs2[0:31]$	
beq	Branch ==	B	1100011	0x0		if(rs1 == rs2) PC += imm	
bne	Branch !=	B	1100011	0x1		if(rs1 != rs2) PC += imm	
blt	Branch <	B	1100011	0x4		if(rs1 < rs2) PC += imm	
bge	Branch ≥	B	1100011	0x5		if(rs1 ≥ rs2) PC += imm	
bltu	Branch < (U)	B	1100011	0x6		if(rs1 < rs2) PC += imm	zero-extends
bgeu	Branch ≥ (U)	B	1100011	0x7		if(rs1 ≥ rs2) PC += imm	zero-extends
jal	Jump And Link	J	1101111			rd = PC+4; PC += imm	
jalr	Jump And Link Reg	I	1100111	0x0		rd = PC+4; PC = rs1 + imm	
lui	Load Upper Imm	U	0110111			rd = imm << 12	
auipc	Add Upper Imm to PC	U	0010111			rd = PC + (imm << 12)	
ecall	Environment Call	I	1110011	0x0	imm=0x0	Transfer control to OS	
ebreak	Environment Break	I	1110011	0x0	imm=0x1	Transfer control to debugger	

RV32M Multiply Extension

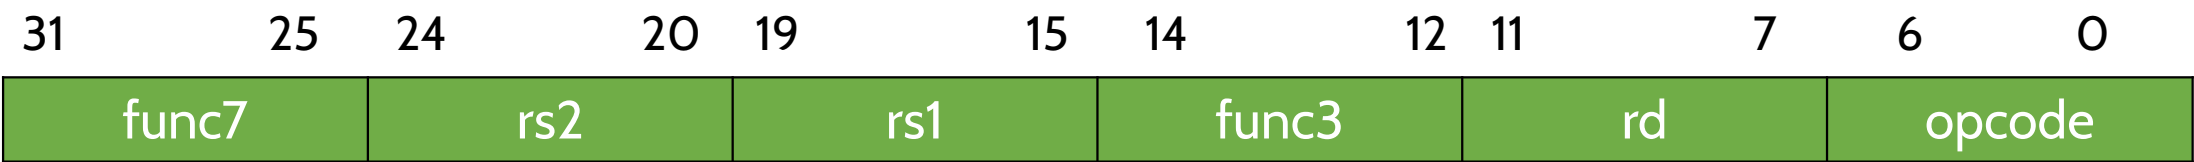
Inst	Name	FMT	Opcode	funct3	funct7	Description (C)
mul	MUL	R	0110011	0x0	0x01	$rd = (rs1 * rs2)[31:0]$
mulh	MUL High	R	0110011	0x1	0x01	$rd = (rs1 * rs2)[63:32]$
mulsu	MUL High (S) (U)	R	0110011	0x2	0x01	$rd = (rs1 * rs2)[63:32]$
mulu	MUL High (U)	R	0110011	0x3	0x01	$rd = (rs1 * rs2)[63:32]$
div	DIV	R	0110011	0x4	0x01	$rd = rs1 / rs2$
divu	DIV (U)	R	0110011	0x5	0x01	$rd = rs1 / rs2$
rem	Remainder	R	0110011	0x6	0x01	$rd = rs1 \% rs2$
remu	Remainder (U)	R	0110011	0x7	0x01	$rd = rs1 \% rs2$

RV32A Atomic Extension



Inst	Name	FMT	Opcode	funct3	funct5	Description (C)
lr.w	Load Reserved	R	0101111	0x2	0x02	$rd = M[rs1]$, reserve $M[rs1]$
sc.w	Store Conditional	R	0101111	0x2	0x03	if (reserved) { $M[rs1] = rs2$; $rd = 0$ } else { $rd = 1$ }
amoswap.w	Atomic Swap	R	0101111	0x2	0x01	$rd = M[rs1]$; $swap(rd, rs2)$; $M[rs1] = rd$
amoadd.w	Atomic ADD	R	0101111	0x2	0x00	$rd = M[rs1] + rs2$; $M[rs1] = rd$
amoand.w	Atomic AND	R	0101111	0x2	0x0C	$rd = M[rs1] \& rs2$; $M[rs1] = rd$
amoor.w	Atomic OR	R	0101111	0x2	0x0A	$rd = M[rs1] rs2$; $M[rs1] = rd$
amoxor.w	Atomix XOR	R	0101111	0x2	0x04	$rd = M[rs1] ^ rs2$; $M[rs1] = rd$
amomax.w	Atomic MAX	R	0101111	0x2	0x14	$rd = \max(M[rs1], rs2)$; $M[rs1] = rd$
amomin.w	Atomic MIN	R	0101111	0x2	0x10	$rd = \min(M[rs1], rs2)$; $M[rs1] = rd$

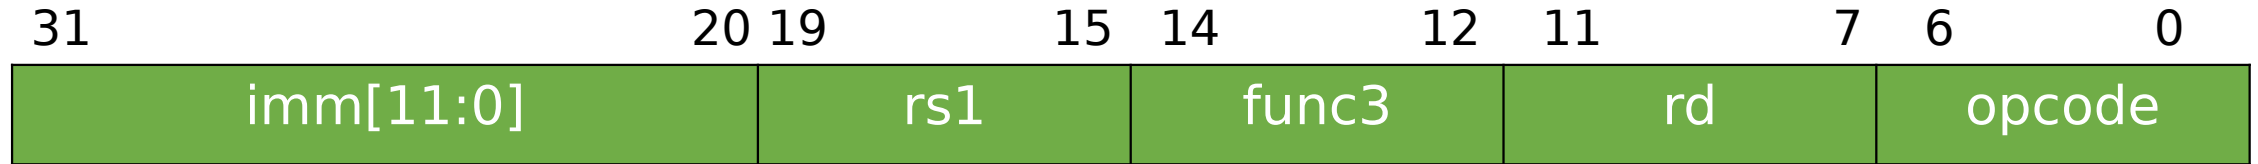
RISC-V Instruction Format: 1: R-Type



Field	No. O f bits	Function
opcode:	R1, A, B	Basic operation of the instruction, and this abbreviation is its traditional name.
operand:	R2, C, D	The register destination operand. It gets the result of the operation
funct3:	X, R1, R2	An additional opcode field.
rs1:		The first register source operand.
rs2:		The second register source operand.
func7		An additional opcode field.

Assembly	Field Values						Machine Code						
	func7	rs2	rs1	funct3	rd	op	func7	rs2	rs1	funct3	rd	op	
add s2, s3, s4	0	20	19	0	18	51	0000,000	10100	10011	000,	10010	011,0011	(0x01498933)
add x18,x19,x20													
sub t0, t1, t2	32	7	6	0	5	51	0100,000	00111	00110	000,	00101	011,0011	(0x407302B3)
sub x5, x6, x7													
	7 bits	5 bits	5 bits	3 bits	5 bits	7 bits	7 bits	5 bits	5 bits	3 bits	5 bits	7 bits	

Immediate: I-Type



Field	No. Of bits	Function
opcode:	R1, A, B	Basic operation of the instruction, and this abbreviation is its traditional name.
rd:	R2, C, D	The register destination operand. It gets the result of the operation
funct3:	X, R1, R2	An additional opcode field.
rs1:		The first register source operand.
imm		The second register source operand.

Assembly

```
addi s0, s1, 12
addi x8, x9, 12
addi s2, t1, -14
addi x18, x6, -14
lw t2, -6(s3)
lw x7, -6(x19)
lh s1, 27(zero)
lh x9, 27(x0)
lb s4, 0x1F(s4)
lb x20, 0x1F(x20)
```

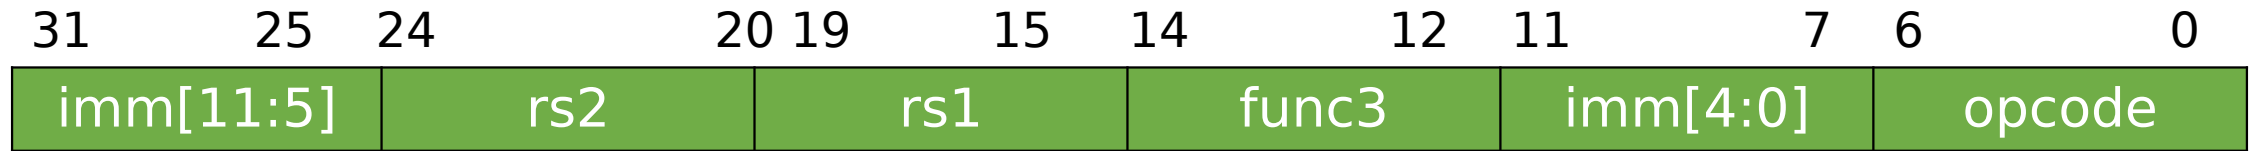
Field Values

imm _{11:0}	rs1	funct3	rd	op
12	9	0	8	19
-14	6	0	18	19
-6	19	2	7	3
27	0	1	9	3
0x1F	20	0	20	3
12 bits	5 bits	3 bits	5 bits	7 bits

Machine Code

imm _{11:0}	rs1	funct3	rd	op	
0000 0000 1100	01001	000	01000	001 0011	(0x00C48413)
1111 1111 0010	00110	000	10010	001 0011	(0xFF230913)
1111 1111 1010	10011	010	00111	000 0011	(0xFFA9A383)
0000 0001 1011	00000	001	01001	000 0011	(0x01B01483)
0000 0001 1111	10100	000	10100	000 0011	(0x01FA0A03)
12 bits	5 bits	3 bits	5 bits	7 bits	

Store: S-Type



i	No. O f bits	Function
opcode:	R1, A, B	Basic operation of the instruction, and this abbreviation is its traditional name.
rd:	R2, C, D	The register destination operand. It gets the result of the operation
imm[4:0]:	X, R1, R2	An additional opcode field.
rs1:		The first register source operand.
rs2:		The second register source operand.
imm[11:5]		An additional opcode field.

#x1 based address
 SW x2, 0(x1) # Memory[x1 + 0] = x2

Branch: B-Type

31	25	24	20	19	15	14	12	11	7	6	0
imm[11:5]	rs2			rs1		func3		imm[4:0]		opcode	

i	No. O f bits	Function
opcode:	R1, A, B	Basic operation of the instruction, and this abbreviation is its traditional name.
rd:	R2, C, D	The register destination operand. It gets the result of the operation
imm[4:0]:	X, R1, R2	An additional opcode field.
rs1:		The first register source operand.
rs2:		The second register source operand.
imm[11:5]		An additional opcode field.

BEQ x1, x2, equal # If x1 == x2, branch to label 'equal'

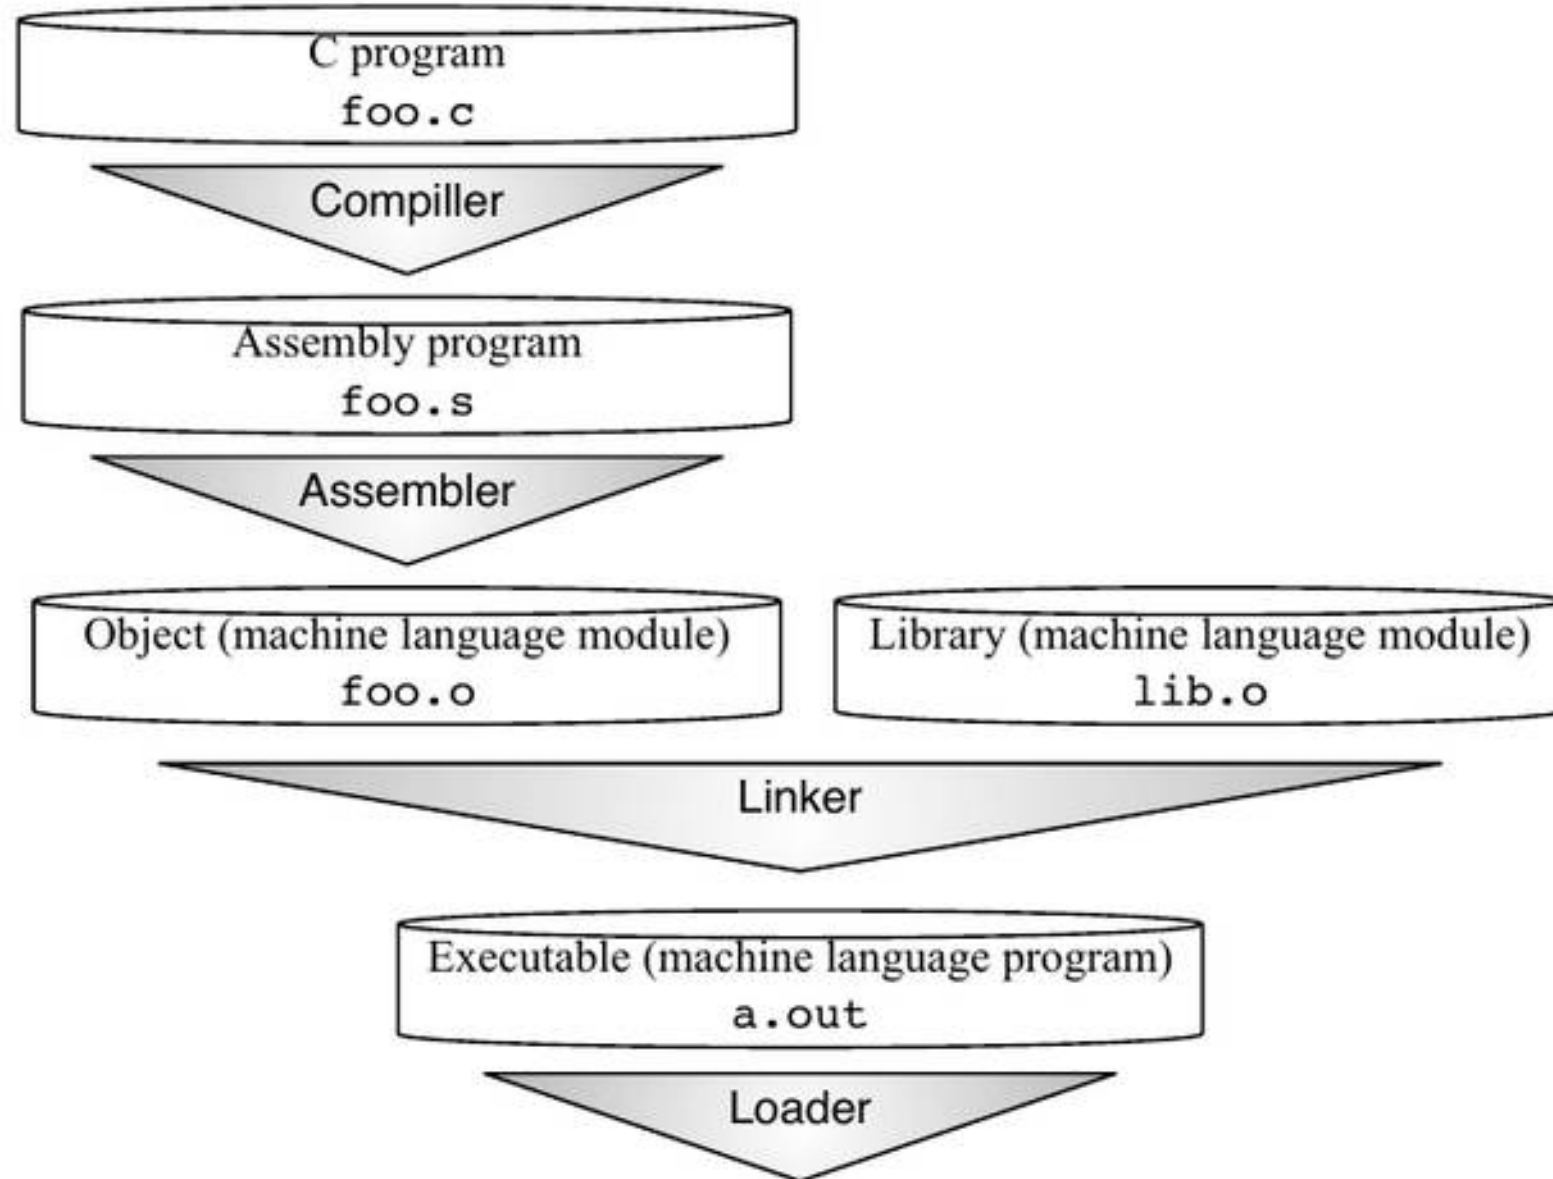
Jump: J-Type

- Unconditional jump with an optional link to store the return address.
 - } opcode (7 bits): Operation code that specifies the jump instruction (e.g., JAL).
 - } rd (5 bits): Destination register for the return address.
 - } immediate (20 bits): Jump target offset, which is used to calculate the jump address relative to the current program counter (PC).

Topics

1. Basic Processor Architecture
2. Different Types of Processor Architectures
3. RISC-V Processor Architecture
4. RISC-V Instruction Set Architecture
5. **Programming RISC-V using assembly language**





RISCV GCC Assembler

`_start:`

`ld s3 0x001121`

`ld rs2 0x0022233`

`add rd, rs1, rs2`

`st rd 0x00000001`

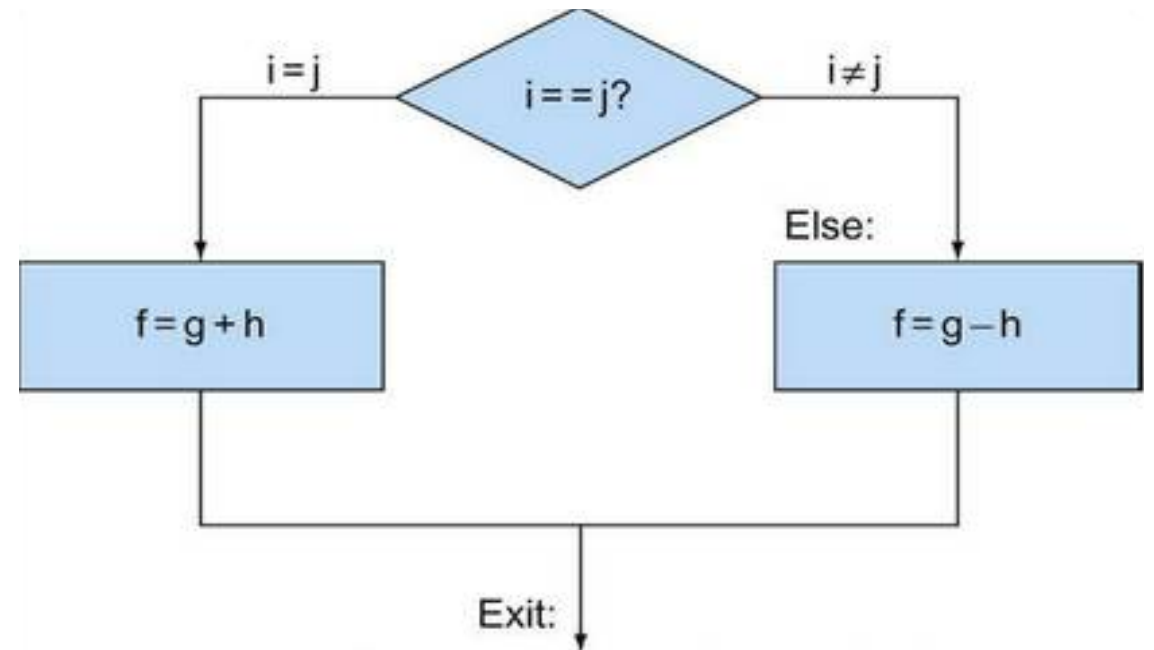
Assembly or C/C++

- Write Efficient Code
- Secure Application
- Multi-Threaded and Complex Program to run multiple devices (OS)
- Real-Time Applications for Real world Problems

Programming RISC-V

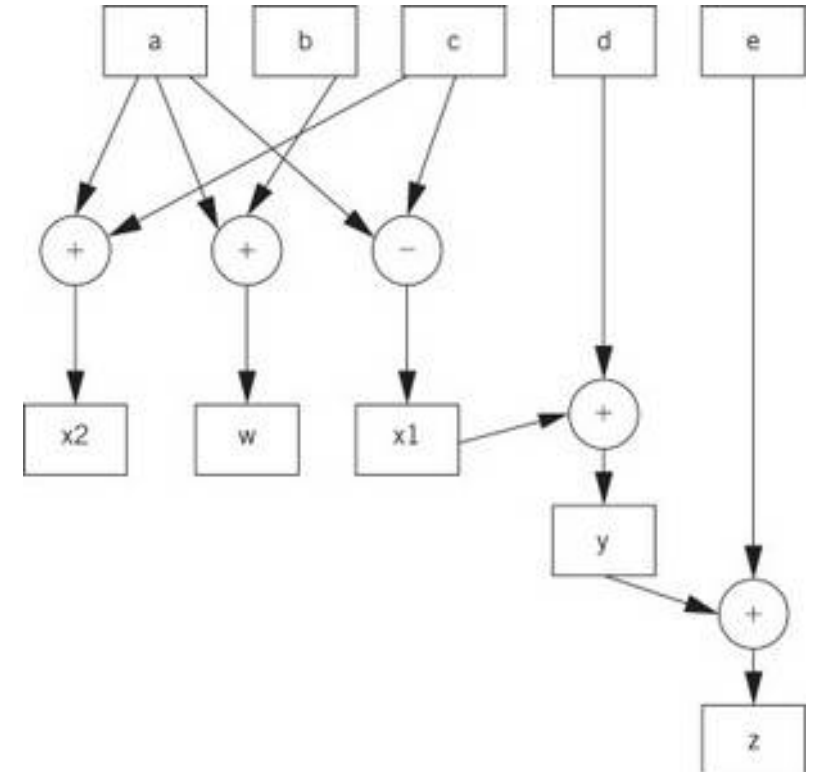
- Problem
- Write it in your own words
- Make Pseudo Code
- Create Control and Data-flow Graph
- Program (C/C++, ASM)
- Debug
- Profile
- Optimize/Fine Tune
- Execute
- Test

Flowchar



Hazards

- Data Hazards: Instructions are waiting for data from other instructions.
- Control Hazards: Changes in instruction flow cause delays.
- Structural Hazards: Limited hardware resources cause delays.



```
// example.c
int global_var = 10;
int main() {
    int local_var = 5;
    int result = global_var +
local_var;
    return result;
}
```

```
riscv32-unknown-elf-gcc example.o -o example
```

- The compiler generates an object file in ELF format. This object file contains machine code, data, and metadata, organized into different sections like .text (code), .data (initialized data), and .bss (uninitialized data).

- Instruction Section: Contains the compiled machine code instructions (text section).
- Data Section: Contains initialized data (data section).
- The linker combines the code and data sections, resolves symbols, and sets up memory addresses.
- The linker script defines how different sections are mapped into the memory of the microcontroller.
- It specifies memory regions and assigns addresses to different sections of the code and data.

```

• MEMORY
• {
  •   ROM (rx) : ORIGIN = 0x08000000, LENGTH = 512K
  •   RAM (rwx) : ORIGIN = 0x20000000, LENGTH = 64K
  • }
•
• SECTIONS
• {
  •   .text : {
  •       *(.text)
  •   } > ROM
  •
  •   .data : {
  •       *(.data)
  •   } > RAM
  • }

```

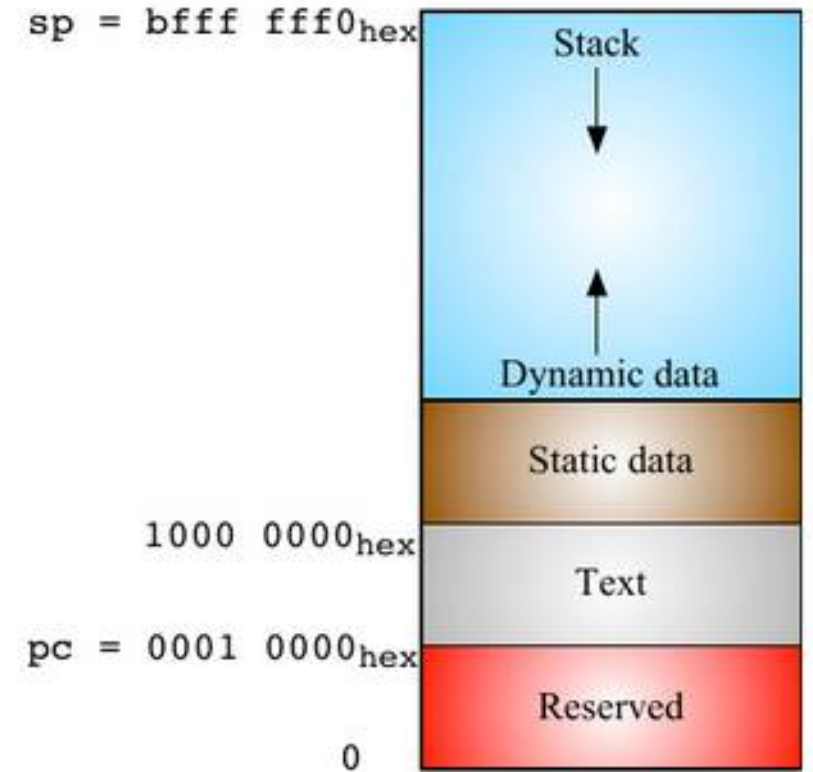
- Next step involves using a programmer or debugger tool to flash the firmware into the RISC-V System.
 - } Instruction Memory: The code from the .text section is loaded into the system instruction memory.
 - } Data Memory: The initialized data from the .data section is loaded into the system data memory.

Linker Script: Program and Data Memory Allocation

The high addresses are the top of the figure and the low addresses are the bottom.

The stack pointer (sp) starts at BFFF FFF0 hex and grows down toward the Static data. The text (program code) starts at 0001 0000hex and includes the statically-linked libraries.

The Static data starts immediately above the text region; in this example, we assume that address is 1000 0000hex . Dynamic data, allocated in C by malloc(), is just above the Static data. Called the heap, it grows upward toward the stack. It includes the dynamically-linked libraries.



Testing and Executing the Code

RIPES

<https://ripes.me/>

https://github.com/mortbopet/Ripes/releases/download/v2.2.6/Ripes-v2.2.6-win-x86_64.zip

Next:

RISCV Micro Controller

RISCV Simulator and Emulators

RISCV Single Board Computer



Center of Excellence: Supercomputing for RISC-V Processor Hardware and Instruction Set Architecture

by: Tassadaq Hussain

Director Centre for AI and BigData

Professor Department of Electrical Engineering

Namal University Mianwali

Collaborations:

Barcelona Supercomputing Center, Spain

European Network on High Performance and Embedded Architecture and Compilation

Pakistan Supercomputing Center