



# Informatics



## Computersysteme

### Microarchitecture

---

Markus Bader

SS2024


# Introduction

---

DDCA Ch7 - Part 1: Microarchitecture Introduction <https://youtu.be/lrN-uBKooRY?si=QEiy6eyr5c32m31n>

## Chapter 7 :: Topics

- Introduction
- Performance Analysis
- Single-Cycle Processor
- Multicycle Processor
- Pipelined Processor
- Advanced Microarchitecture



Application Software	
Operating Systems	
Architecture	
Micro-architecture	
Logic	
Digital Circuits	
Analog Circuits	
Devices	
Physics	

## About these Notes

Digital Design and Computer Architecture Lecture Notes

© 2021 Sarah Harris and David Harris

These notes may be used and modified for educational and/or non-commercial purposes so long as the source is attributed.

# Microarchitecture

- Multiple implementations for a single architecture:
  - **Single-cycle**: Each instruction executes in a single cycle
  - **Multicycle**: Each instruction is broken up into series of shorter steps
  - **Pipelined**: Each instruction broken up into series of steps & multiple instructions execute at once

# Processor Performance

- Program execution time  
Execution Time = (#instructions)(cycles/instruction)(seconds/cycle)
- Definitions:
  - **CPI**: Cycles/instruction
  - **clock period**: seconds/cycle
  - **IPC**: instructions/cycle = IPC
- Challenge is to satisfy constraints of:
  - Cost
  - Power
  - Performance

# RISC-V Processor

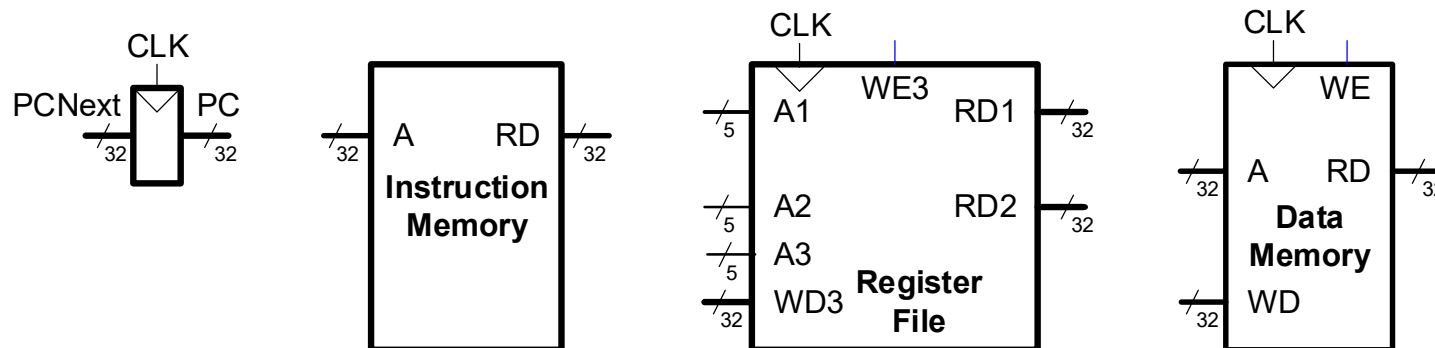
- Consider **subset** of RISC-V instructions:
- R-type ALU instructions:  
`add, sub, and, or, slt`
- Memory instructions:  
`lw, sw`
- Branch instructions:  
`beq`

# Architectural State Elements

Determines everything about a processor:

- **Architectural state:**

- 32 registers
- PC
- Memory





# Single-Cycle RISC-V Processor

---

DDCA Ch7 - Part 2: RISC-V Single-Cycle Processor Datapath: lw <https://youtu.be/AoBkibsIRBM?si=fgO1anrXwzMCdOrU>

# Single-Cycle RISC-V Processor

- Datapath
- Control

## Example Program

- Design datapath
- View example program executing


Address	Instruction	Type	Fields					Machine Language	
0x1000	L7: lw x6, -4(x9)	I	<b>imm<sub>11:0</sub></b> 111111111100	<b>rs1</b> 01001	<b>f3</b> 010	<b>rd</b> 00110	<b>op</b> 0000011	FFC4A303	
0x1004	sw x6, 8(x9)	S	<b>imm<sub>11:5</sub></b> 0000000	<b>rs2</b> 00110	<b>rs1</b> 01001	<b>f3</b> 010	<b>imm<sub>4:0</sub></b> 01000	<b>op</b> 0100011	0064A423
0x1008	or x4, x5, x6	R	<b>funct7</b> 0000000	<b>rs2</b> 00110	<b>rs1</b> 00101	<b>f3</b> 110	<b>rd</b> 00100	<b>op</b> 0110011	0062E233
0x100C	beq x4, x4, L7	B	<b>imm<sub>12,10:5</sub></b> 1111111	<b>rs2</b> 00100	<b>rs1</b> 00100	<b>f3</b> 000	<b>imm<sub>4:1,11</sub></b> 10101	<b>op</b> 1100011	FE420AE3

# Single-Cycle RISC-V Processor

- **Datapath:** start with `lw` instruction


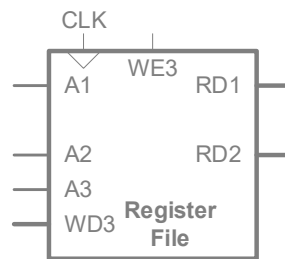
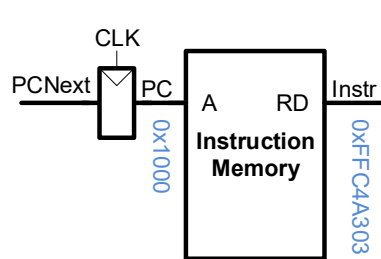
- Example:  
`lw x6, -4(x9)`  
`lw rd, imm(rs1)`

## I-Type



# Single-Cycle Datapath: $l_w$ fetch

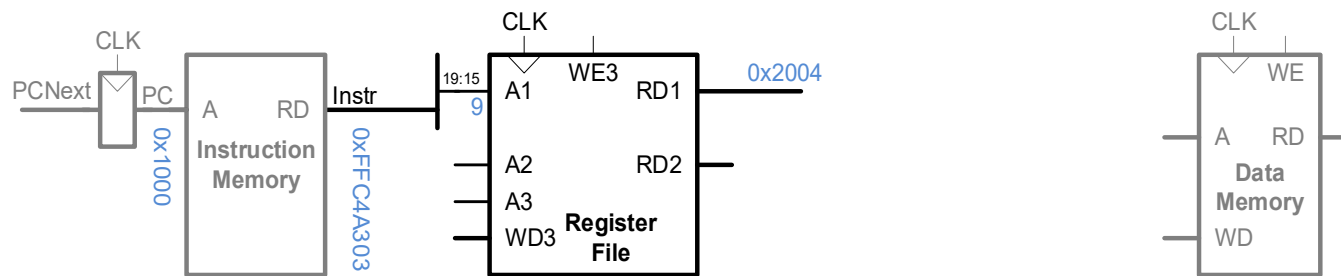
- STEP 1:** Fetch instruction



Address	Instruction	Type	Fields	Machine Language								
0x1000	L7: lw x6, -4(x9)	I	<table border="0"> <tr> <td><b>imm<sub>11:0</sub></b></td> <td><b>rs1</b></td> <td><b>f3</b></td> <td><b>rd</b></td> </tr> <tr> <td>111111111100</td> <td>01001</td> <td>010</td> <td>00110</td> </tr> </table>	<b>imm<sub>11:0</sub></b>	<b>rs1</b>	<b>f3</b>	<b>rd</b>	111111111100	01001	010	00110	0000011 FFC4A303
<b>imm<sub>11:0</sub></b>	<b>rs1</b>	<b>f3</b>	<b>rd</b>									
111111111100	01001	010	00110									

# Single-Cycle Datapath: $lw$ Reg Read

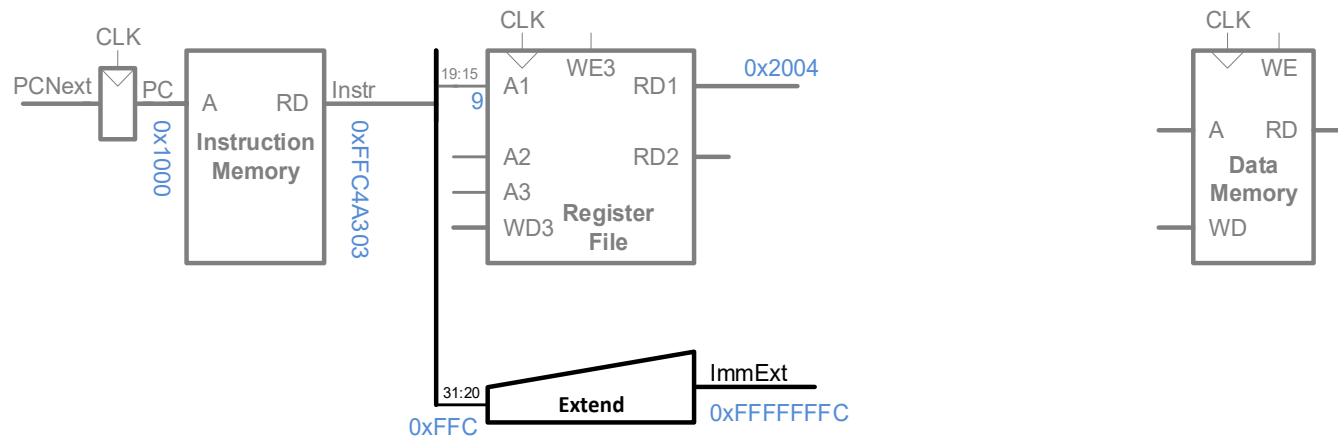
- **STEP 2:** Read source operand (**rs1**) from RF



Address	Instruction	Type	Fields	Machine Language								
0x1000	L7: $lw$ x6, -4(x9)	I	<table border="0"> <tr> <td><math>imm_{11:0}</math></td> <td><b>rs1</b></td> <td><b>f3</b></td> <td><b>rd</b></td> </tr> <tr> <td>111111111100</td> <td>01001</td> <td>010</td> <td>00110</td> </tr> </table>	$imm_{11:0}$	<b>rs1</b>	<b>f3</b>	<b>rd</b>	111111111100	01001	010	00110	0000011 FFC4A303
$imm_{11:0}$	<b>rs1</b>	<b>f3</b>	<b>rd</b>									
111111111100	01001	010	00110									

# Single-Cycle Datapath: $l_w$ Immediate

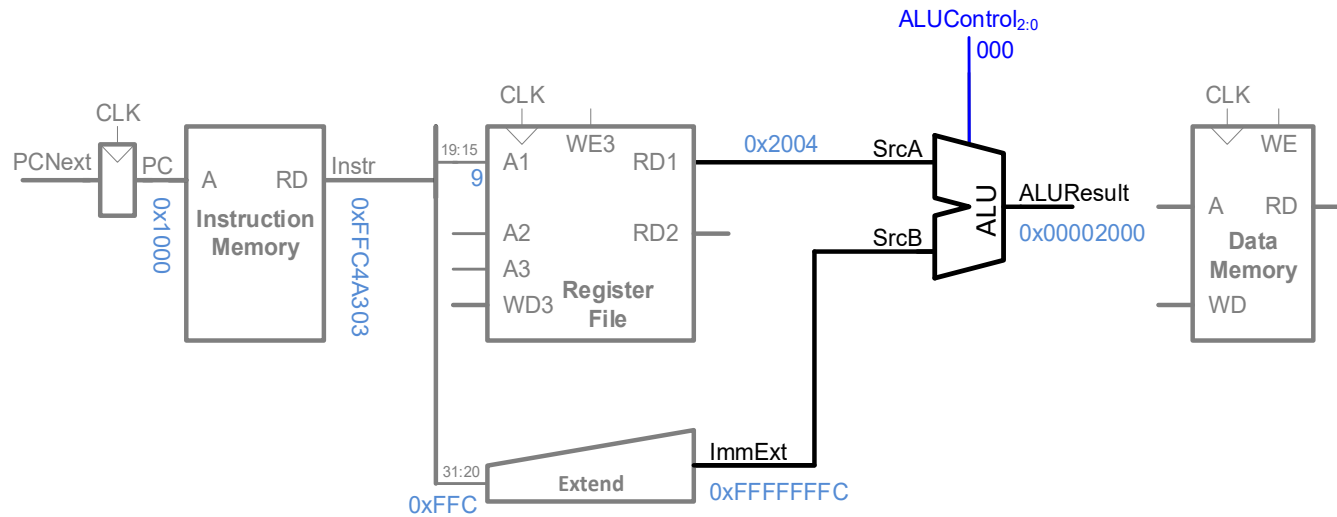
- **STEP 3:** Extend the immediate



Address	Instruction	Type	Fields	Machine Language										
0x1000	L7: $l_w$ x6, -4(x9)	I	<table border="0"> <tr> <td><b>imm<sub>11:0</sub></b></td> <td><b>rs1</b></td> <td><b>f3</b></td> <td><b>rd</b></td> <td><b>op</b></td> </tr> <tr> <td>1111111111100</td> <td>01001</td> <td>010</td> <td>00110</td> <td>0000011</td> </tr> </table>	<b>imm<sub>11:0</sub></b>	<b>rs1</b>	<b>f3</b>	<b>rd</b>	<b>op</b>	1111111111100	01001	010	00110	0000011	FFC4A303
<b>imm<sub>11:0</sub></b>	<b>rs1</b>	<b>f3</b>	<b>rd</b>	<b>op</b>										
1111111111100	01001	010	00110	0000011										

# Single-Cycle Datapath: $l_w$ Address

- **STEP 4:** Compute the memory address



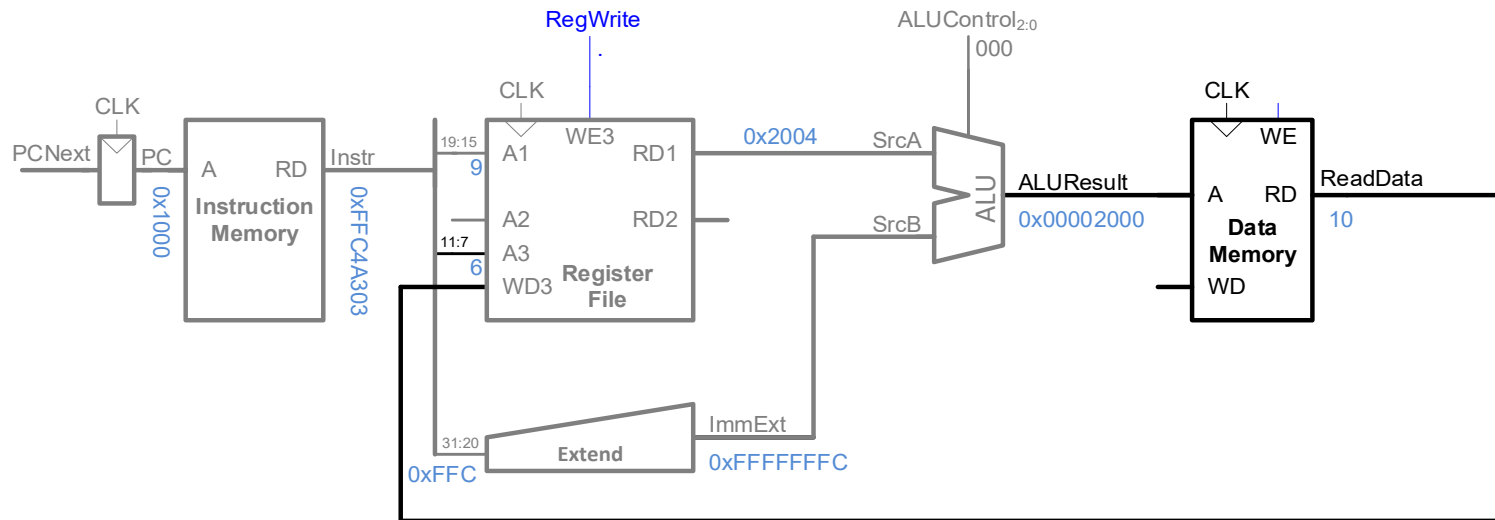
ALUControl <sub>2:0</sub>	Function
000	add
001	subtract
010	and
011	or
101	SLT

Address	Instruction	Type	Fields	Machine Language
0x1000	L7: lw x6, -4(x9)	I	imm <sub>11:0</sub> 1111111111100 rs1 01001 f3 010 rd 00110 op 0000011	FFC4A303



# Single-Cycle Datapath: $l_w$ Mem Read

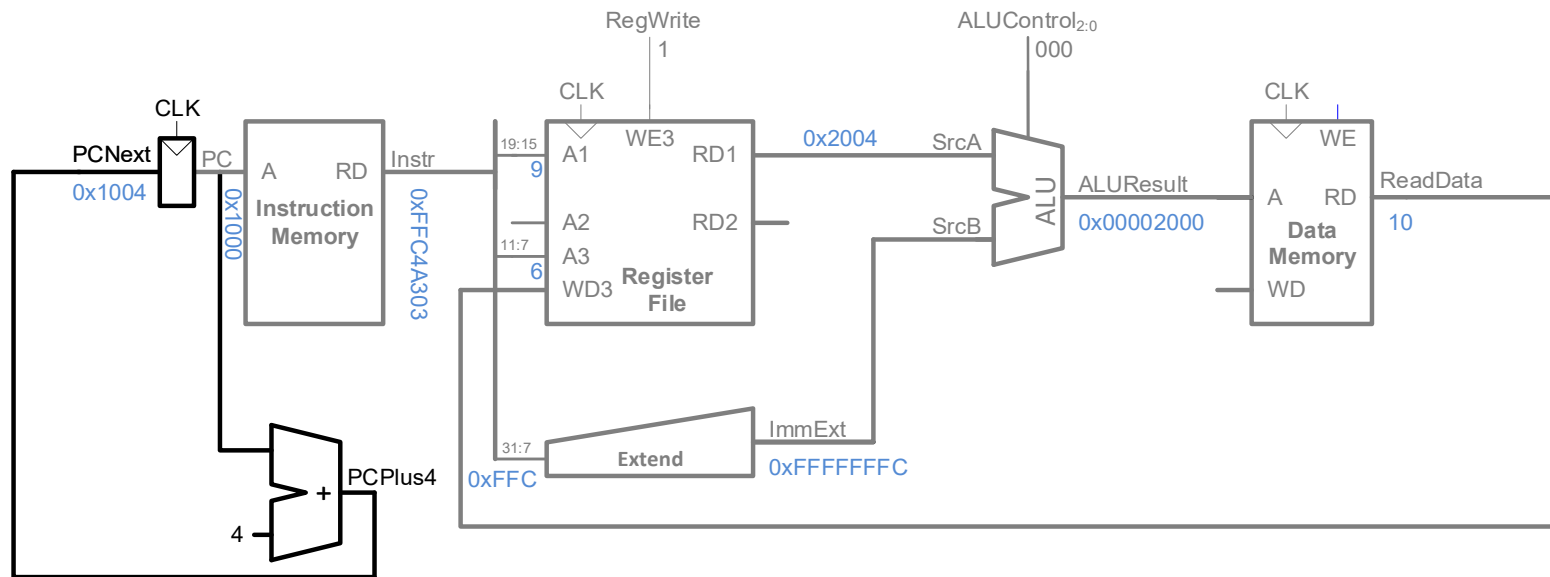
- **STEP 5:** Read data from memory and write it back to register file



Address	Instruction	Type	Fields	Machine Language
0x1000	L7: $l_w$ x6, -4(x9)	I	<b>imm<sub>11:0</sub></b> 1111111111100 <b>rs1</b> 01001 <b>f3</b> 010 <b>rd</b> 00110 <b>op</b> 0000011	FFC4A303

# Single-Cycle Datapath: PC Increment

- **STEP 6:** Determine address of next instruction



Address	Instruction	Type	Fields	Machine Language
0x1000	L7: lw x6, -4(x9)	I	<b>imm<sub>11:0</sub></b> 111111111100 <b>rs1</b> 01001 <b>f3</b> 010 <b>rd</b> 00110 <b>op</b> 0000011	FFC4A303

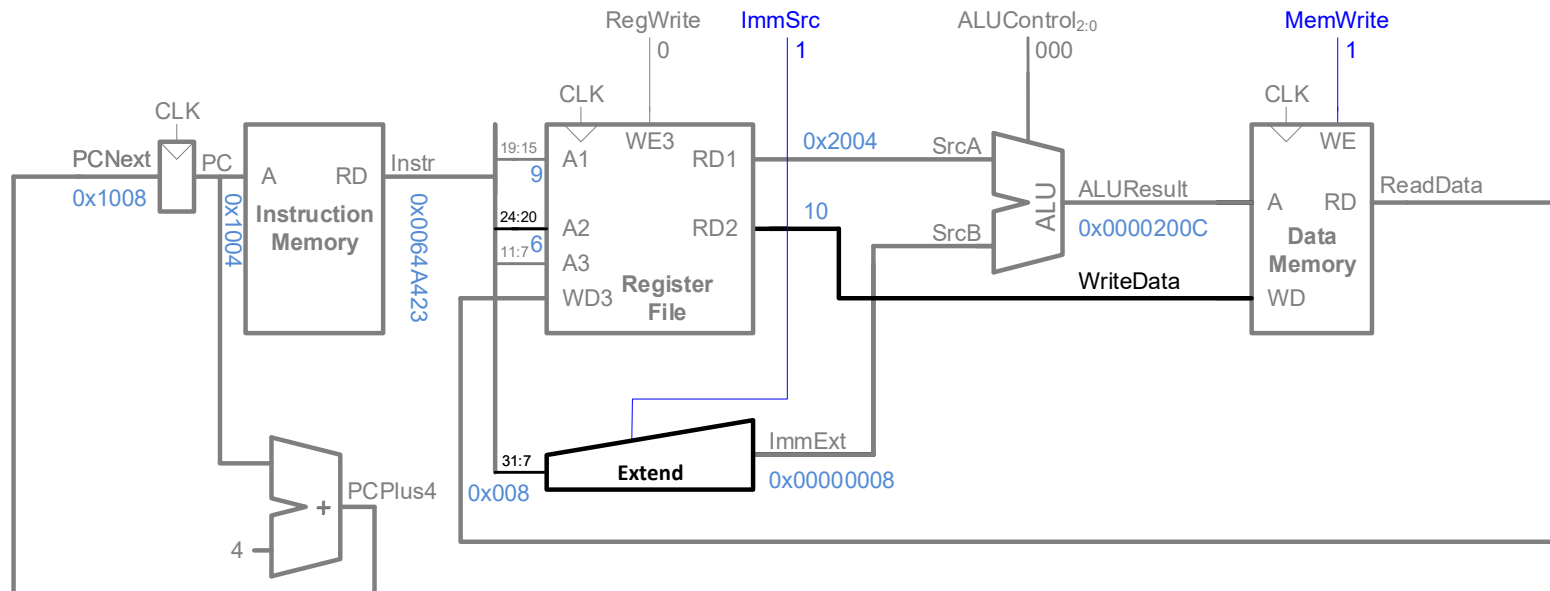
# Single-Cycle Datapath: Other Instructions

---

DDCA Ch7 - Part 3: RISC-V Single-Cycle Processor Datapath [https://youtu.be/sVZmqLRkbVk?si=SE\\_gnswCekVKNuwe](https://youtu.be/sVZmqLRkbVk?si=SE_gnswCekVKNuwe)

# Single-Cycle Datapath: $sw$

- **Immediate:** now in  $\{instr[31:25], instr[11:7]\}$
- **Add control signals:** ImmSrc, MemWrite




Address	Instruction	Type	Fields	Machine Language
0x1004	sw x6, 8(x9)	S	imm <sub>11:5</sub> 0000000 rs2 00110 rs1 01001 f3 010 imm <sub>4:0</sub> 01000 op 0100011	0064A423


## Single-Cycle Datapath: Immediate

ImmSrc	ImmExt	Instruction Type
0	{{20{instr[31]}}, <b>instr[31:20]</b> }	I-Type
1	{{20{instr[31]}}, <b>instr[31:25]</b> , <b>instr[11:7]</b> }	S-Type

### I-Type

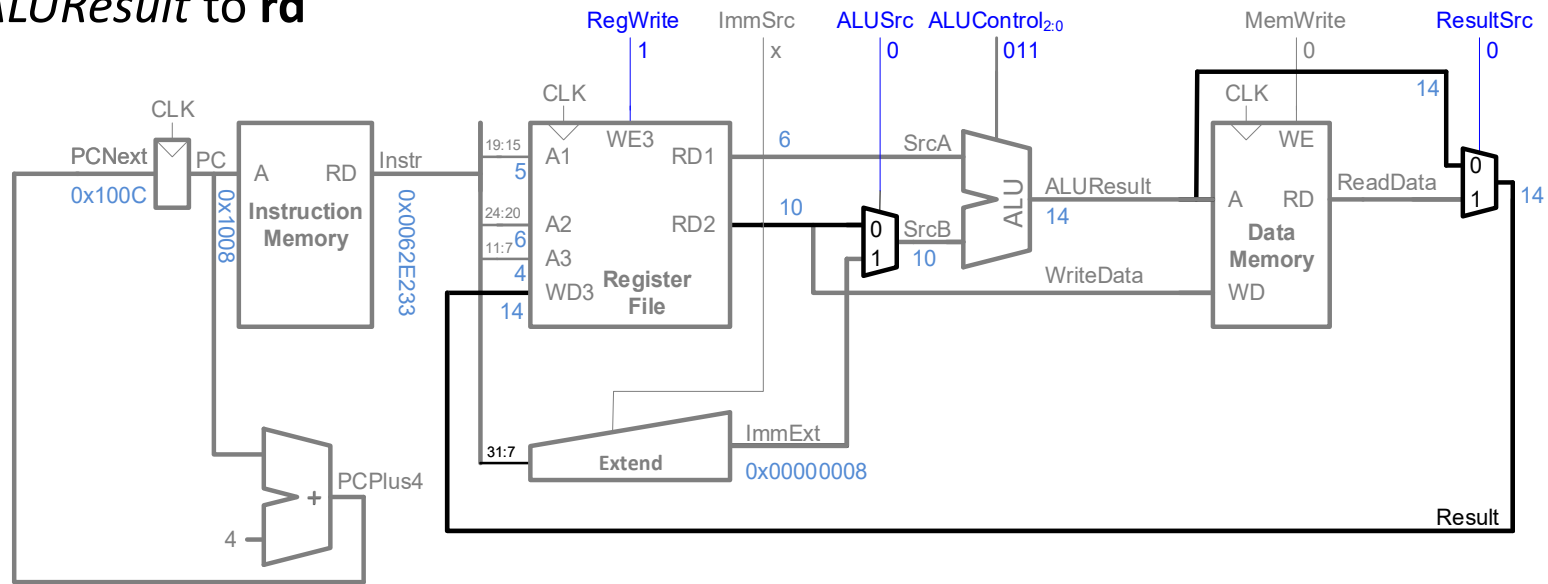


### S-Type



# Single-Cycle Datapath: R-type

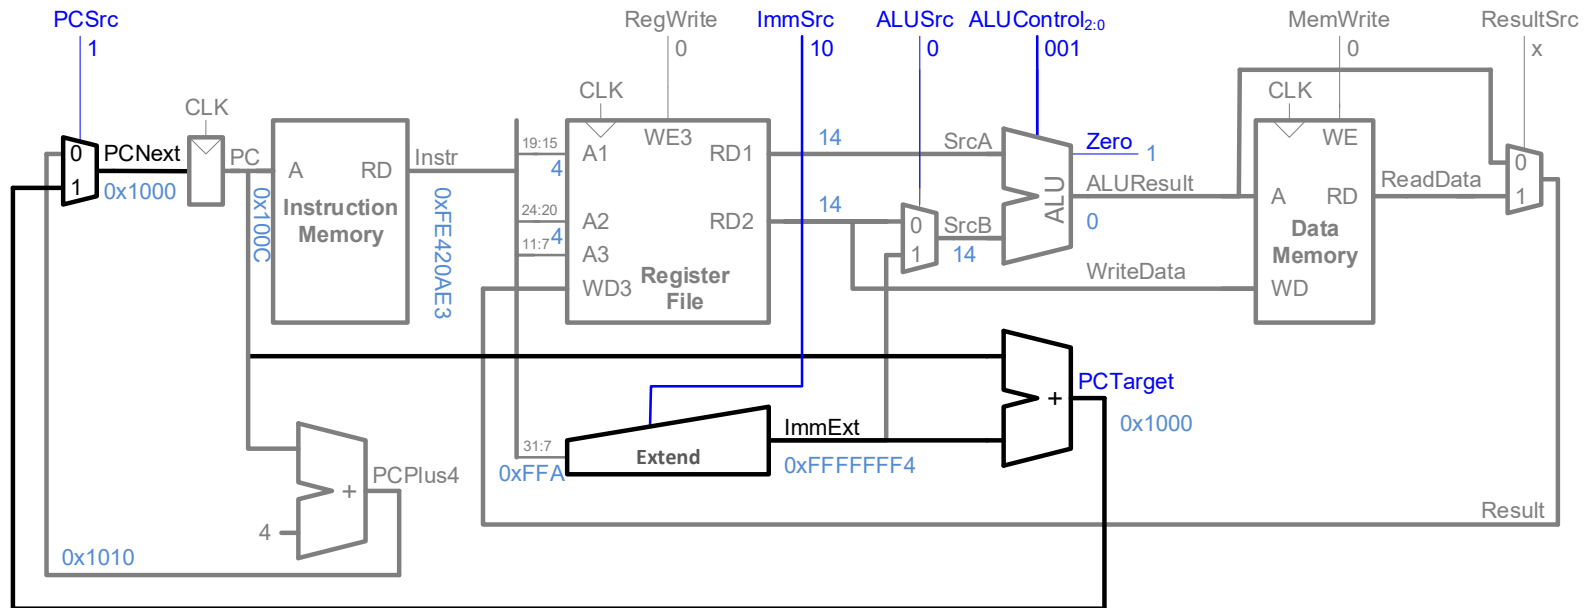
- Read from **rs1** and **rs2** (instead of **imm**)
- Write **ALUResult** to **rd**



Address	Instruction	Type	Fields					Machine Language	
			funct7	rs2	rs1	f3	rd	op	
0x1008	or x4, x5, x6	R	0000000	00110	00101	110	00100	0110011	0062E233

# Single-Cycle Datapath: beq

- Calculate **target address**:  $PCTarget = PC + imm$




Address	Instruction	Type	Fields					Machine Language
0x100C	beq x4, x4, L7	B	$imm_{12,10:5}$ 1111111	$rs2$ 00100	$rs1$ 00100	$f3$ 000	$imm_{4:1,11}$ 10101	1100011 FE420AE3


## Single-Cycle Datapath: ImmExt

ImmSrc <sub>1:0</sub>	ImmExt	Instruction Type
00	{{20{instr[31]}}, <b>instr[31:20]</b> }	I-Type
01	{{20{instr[31]}}, <b>instr[31:25], instr[11:7]</b> }	S-Type
10	{{19{instr[31]}}, <b>instr[31], instr[7], instr[30:25], instr[11:8], 1'b0</b> }	B-Type


### I-Type



### S-Type

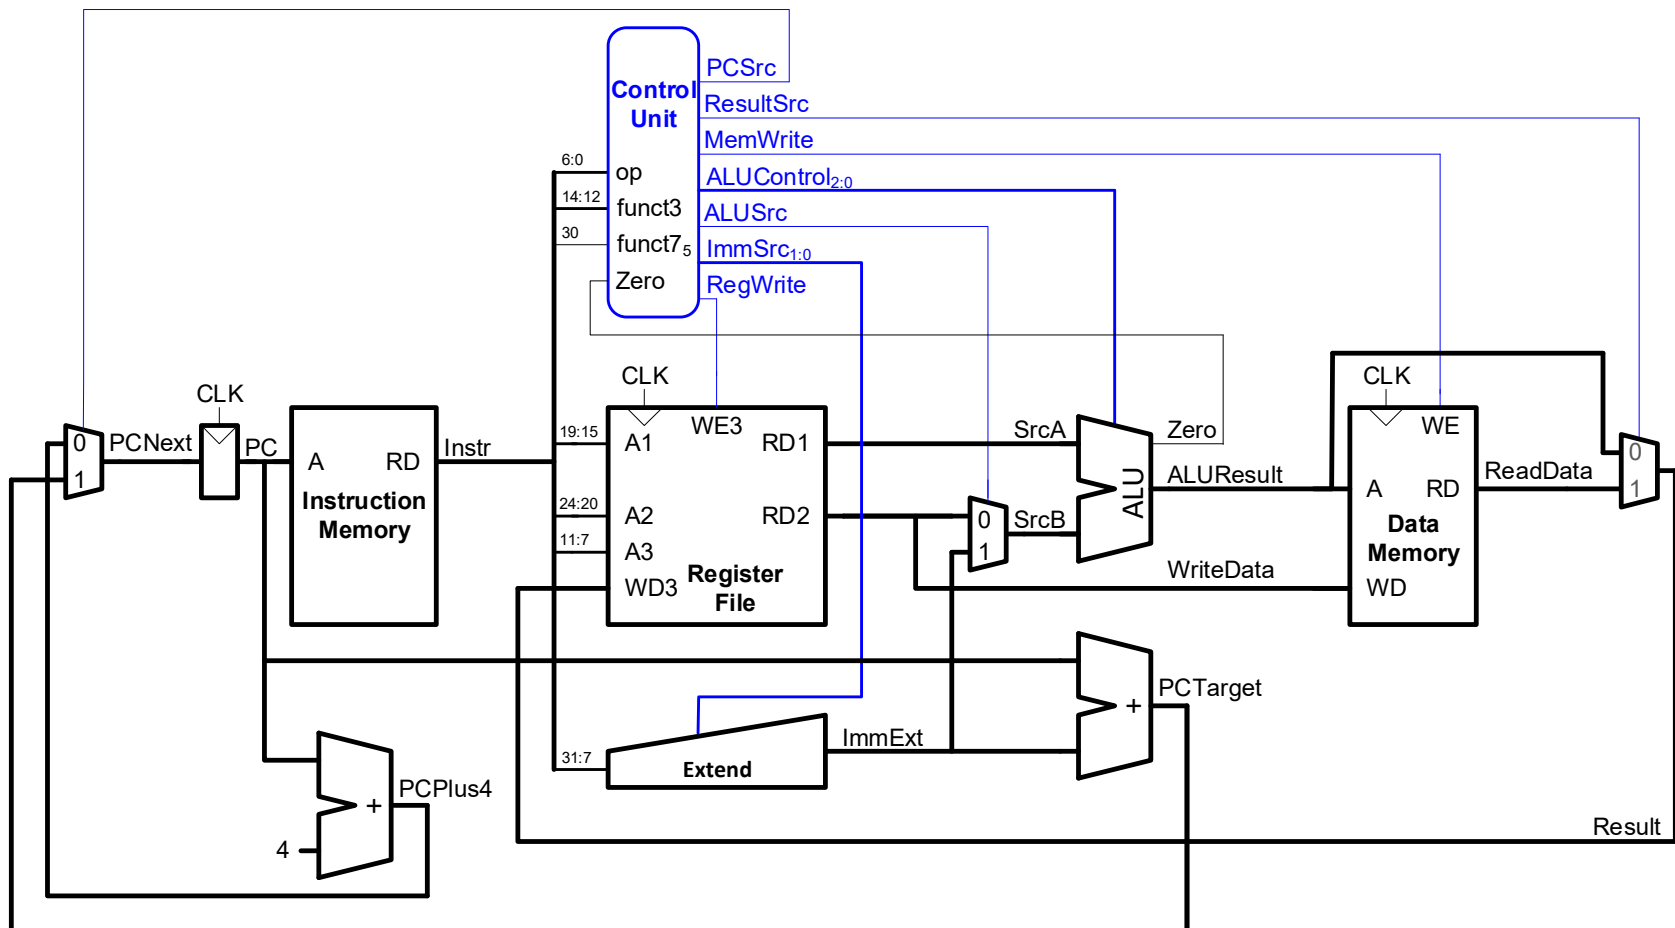


### B-Type





# Single-Cycle RISC-V Processor



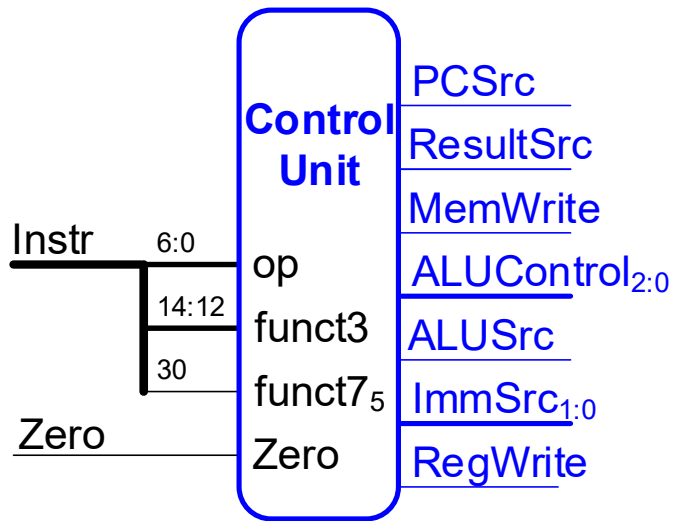
# Single-Cycle Control

---

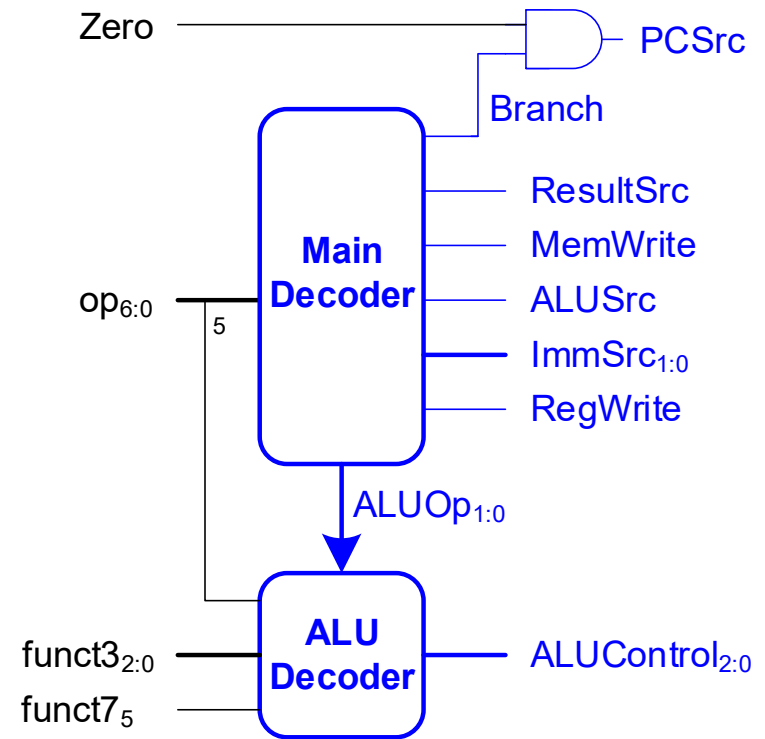
DDCA Ch7 - Part 4: RISC-V Single-Cycle Processor: Control [https://www.youtube.com/watch?v=EZb1\\_VF-yMg](https://www.youtube.com/watch?v=EZb1_VF-yMg)

# Single-Cycle Control

- High-Level View





- Low-Level View



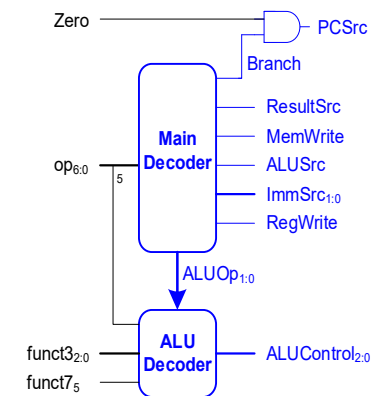
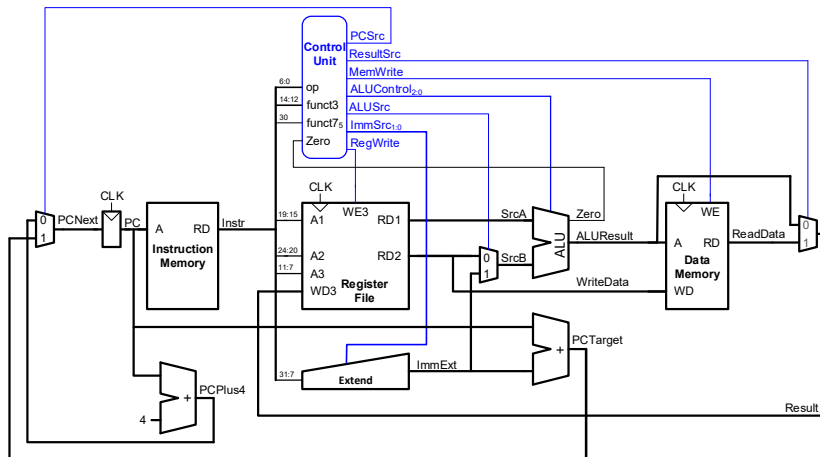
# Single-Cycle Control: Main Decoder (Animation)

op	Instr.	RegWrite	ImmSrc	ALUSrc	MemWrite	ResultSrc	Branch	ALUOp
3	<b>lw</b>							
35	<b>sw</b>							
51	<b>R-type</b>							
99	<b>beq</b>							




# Single-Cycle Control: Main Decoder

op	Instr.	RegWrite	ImmSrc	ALUSrc	MemWrite	ResultSrc	Branch	ALUOp
3	<b>lw</b>	1	00	1	0	1	0	00
35	<b>sw</b>	0	01	1	1	X	0	00
51	<b>R-type</b>	1	XX	0	0	0	0	10
99	<b>beq</b>	0	10	0	0	X	1	01




## Review: ALU

ALUControl <sub>2:0</sub>	Function
000	add
001	subtract
010	and
011	or
101	SLT



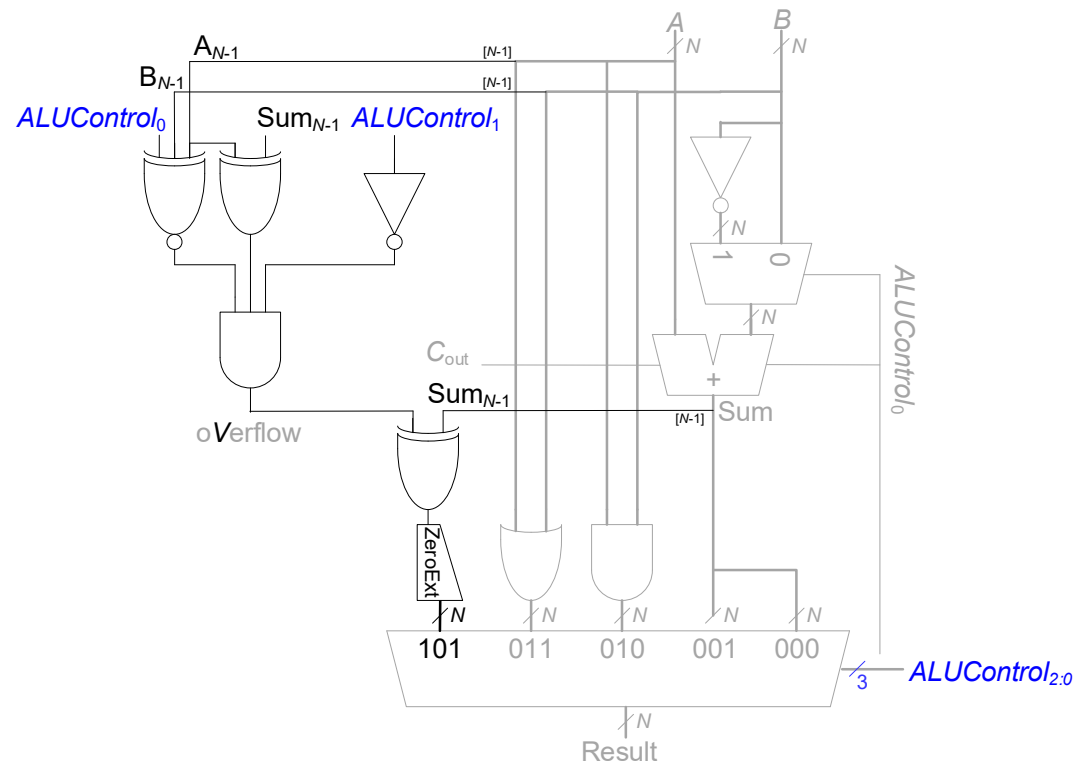
# Review: ALU

ALUControl <sub>2:0</sub>	Function
000	add
001	subtract
010	and
011	or
101	SLT



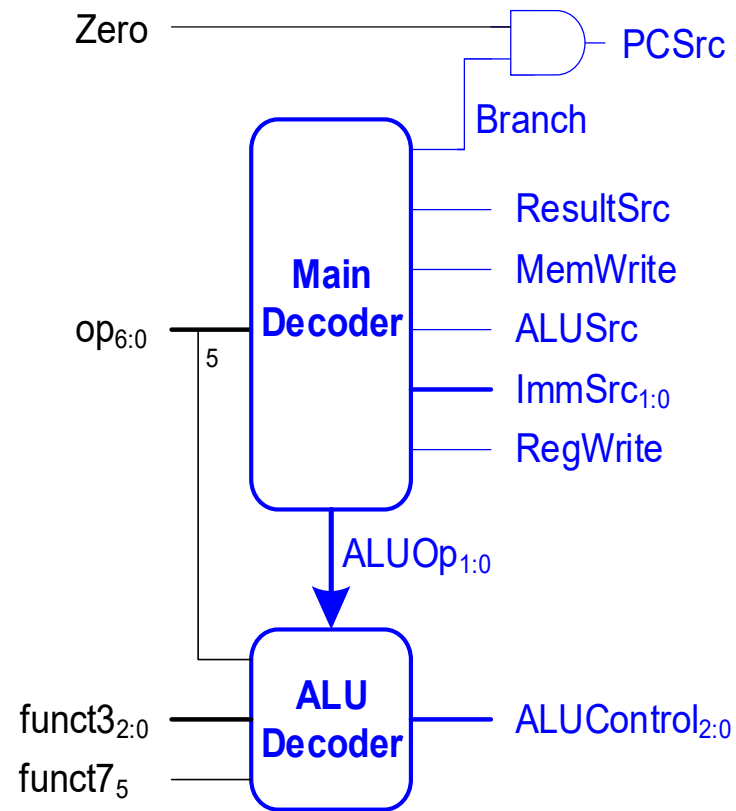
# Review: ALU

ALUControl <sub>2:0</sub>	Function
000	add
001	subtract
010	and
011	or
101	SLT



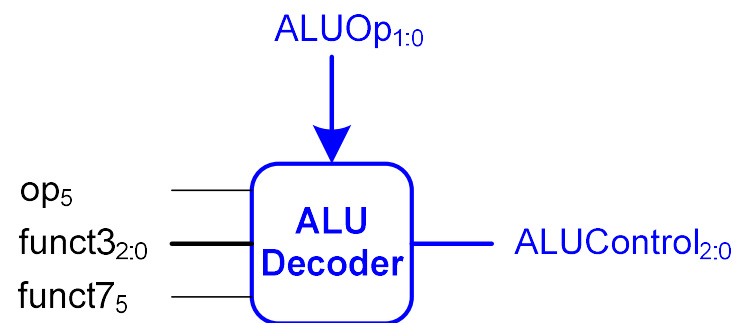


# Single-Cycle Control: ALU Decoder



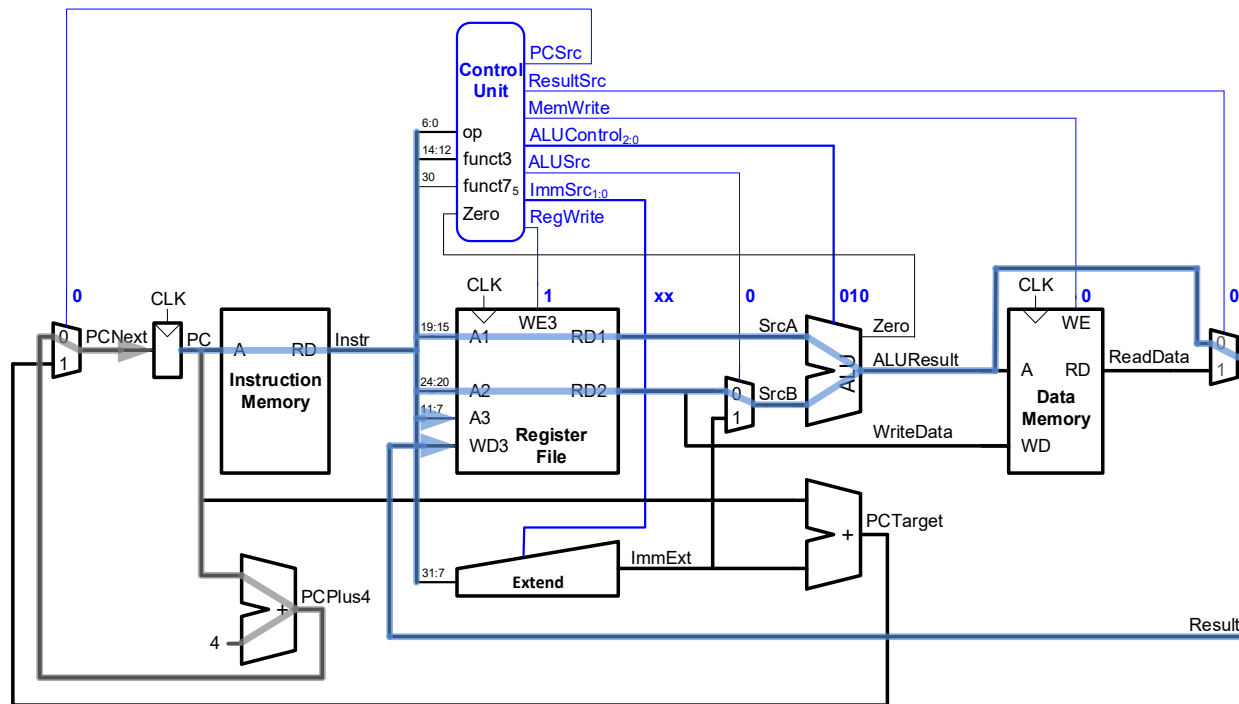
## Single-Cycle Control: ALU Decoder

$ALUOp$	funct3	$op_5, funct7_5$	Instruction	$ALUControl_{2:0}$
00	x	x	<b>lw, sw</b>	000 (add)
01	x	x	<b>beq</b>	001 (subtract)
10	000	00, 01, 10	<b>add</b>	000 (add)
	000	11	<b>sub</b>	001 (subtract)
	010	x	<b>slt</b>	101 (set less than)
	110	x	<b>or</b>	011 (or)
	111	x	<b>and</b>	010 (and)



# Example: and

op	Instruct	RegWrite	ImmSrc	ALUSrc	MemWrite	ResultSrc	Branch	ALUOp
51	R-type	1	XX	0	0	0	0	10



and x5, x6, x7  
Computer Systems

# Extending the Single-Cycle Processor

---

DDCA Ch7 - Part 5: RISC-V Single-Cycle Processor: [https://youtu.be/z6qxMFgNEM4?si=QTFcjiic\\_Hq3uRfi](https://youtu.be/z6qxMFgNEM4?si=QTFcjiic_Hq3uRfi)

## Extended Functionality: I-Type ALU

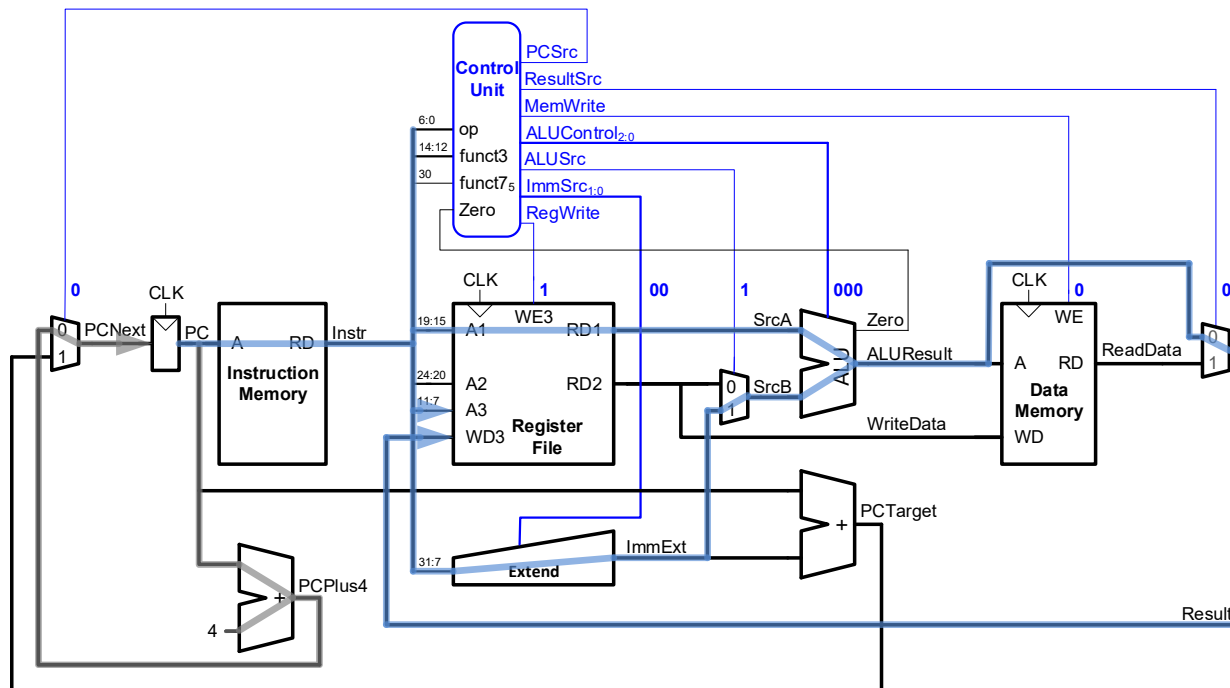
- Enhance the single-cycle processor to handle  
**I-Type ALU instructions:**  
`addi, andi, ori, and slti`
- **Similar to R-type** instructions
- But **second source** comes from **immediate**
- Change **ALUSrc** to select the immediate
- And **ImmSrc** to pick the correct immediate

## Extended Functionality: I-Type ALU

op	Instruct.	RegWrite	ImmSrc	ALUSrc	MemWrite	ResultSrc	Branch	ALUOp
3	<b>lw</b>	1	00	1	0	1	0	00
35	<b>sw</b>	0	01	1	1	X	0	00
51	R-type	1	XX	0	0	0	0	10
99	<b>beq</b>	0	10	0	0	X	1	01
19	<b>I-type</b>	1	00	1	0	0	0	10

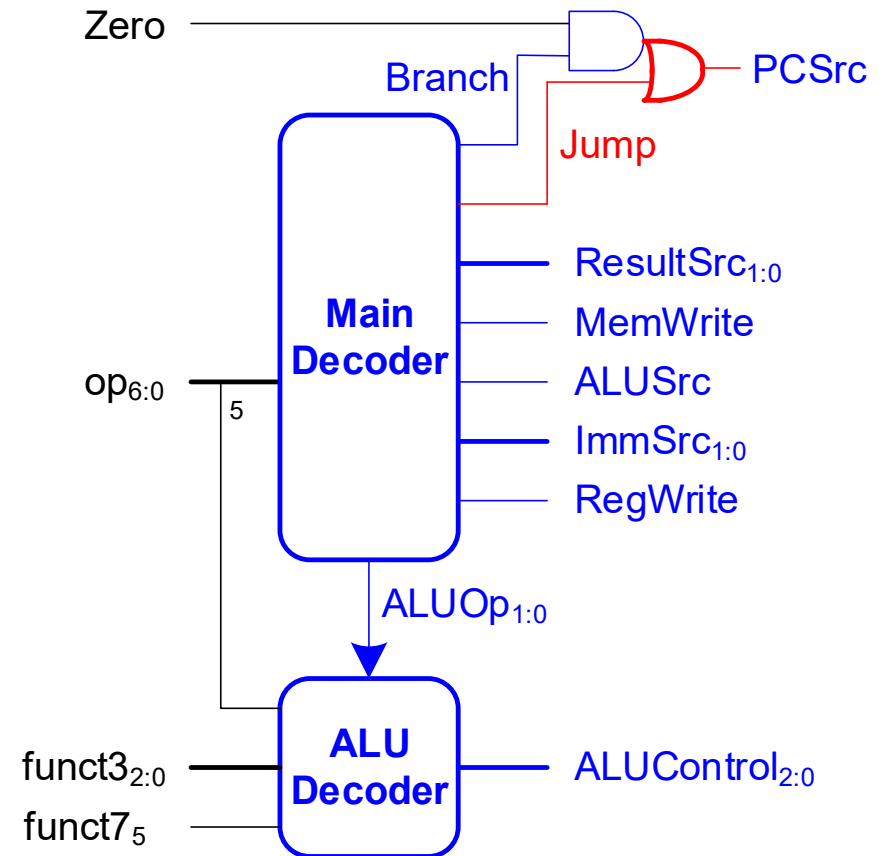
# Extended Functionality: addi

op	Instruct.	RegWrite	ImmSrc	ALUSrc	MemWrite	ResultSrc	Branch	ALUOp
19	I-type	1	00	1	0	0	0	10



## Extended Functionality: jal

- Enhance the single-cycle processor to handle jal
  - Similar to beq
- But jump is **always taken**
  - *PCSrc* should be 1
- **Immediate format** is different
  - Need a new *ImmSrc* of 11
- And jal must **compute PC+4** and store in rd
  - Take PC+4 from adder through *ResultMux*







## Extended Functionality: ImmExt

ImmSrc <sub>1:0</sub>	ImmExt	Instruction Type
00	{{20{instr[31]}}, instr[31:20]}	I-Type
01	{{20{instr[31]}}, instr[31:25], instr[11:7]}	S-Type
10	{{19{instr[31]}}, instr[31], instr[7], instr[30:25], instr[11:8], 1'b0}	B-Type
<b>11</b>	{{12{instr[31]}}, <b>instr[19:12], instr[20], instr[30:21], 1'b0</b> }	<b>J-Type</b>


### I-Type




### S-Type



### B-Type



### J-Type

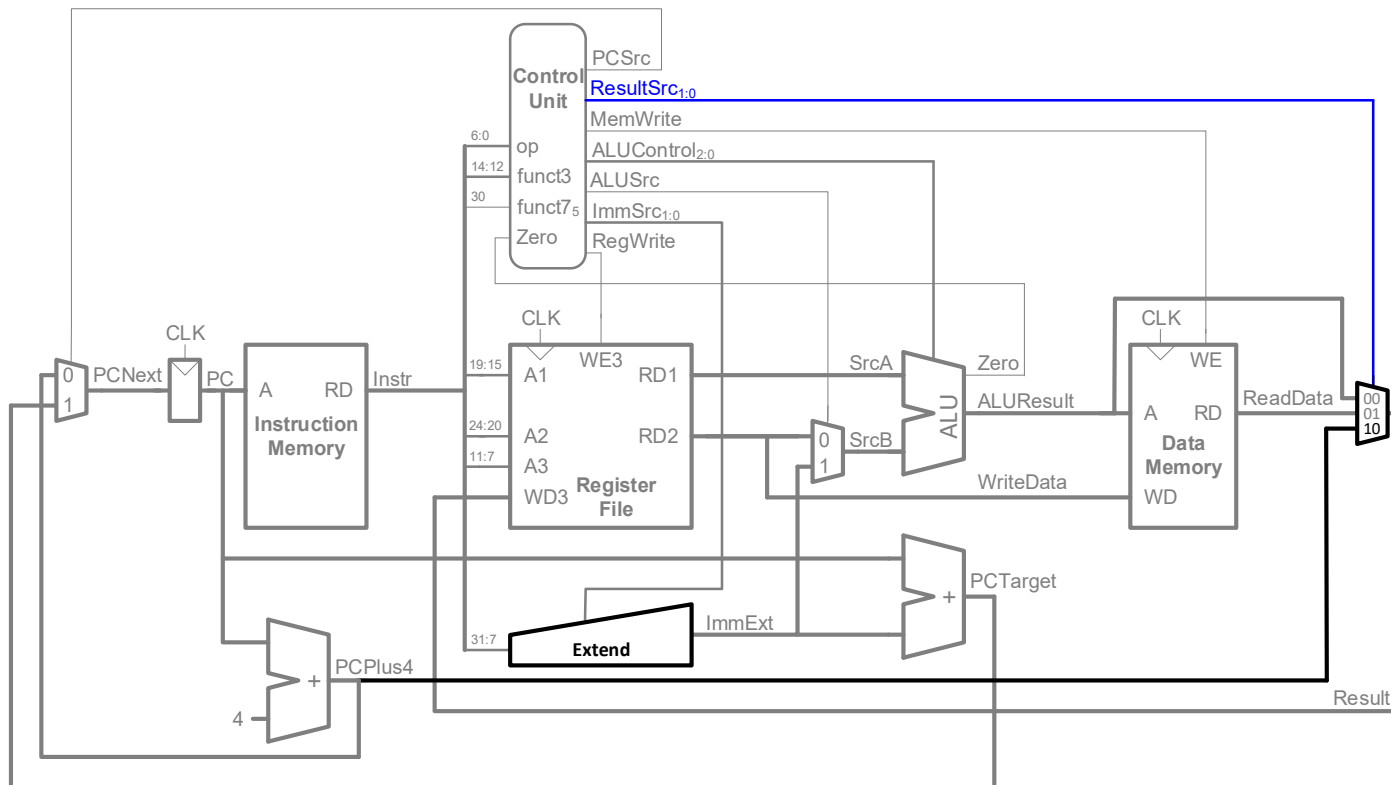


## Extended Functionality: jal

op	Instruct.	RegWrite	ImmSrc	ALUSrc	MemWrite	ResultSrc	Branch	ALUOp	Jump
3	<b>lw</b>	1	00	1	0	01	0	00	0
35	<b>sw</b>	0	01	1	1	XX	0	00	0
51	<b>R-type</b>	1	XX	0	0	00	0	10	0
99	<b>beq</b>	0	10	0	0	XX	1	01	0
19	<b>I-type</b>	1	00	1	0	00	0	10	0
<b>111</b>	<b>jal</b>	<b>1</b>	<b>11</b>	<b>X</b>	<b>0</b>	<b>10</b>	<b>0</b>	<b>XX</b>	<b>1</b>

# Extended Functionality: jal

op	Instruct.	RegWrite	ImmSrc	ALUSrc	MemWrite	ResultSrc	Branch	ALUOp	Jump
111	jal	1	11	X	0	10	0	XX	1



# Single-Cycle Performance

---

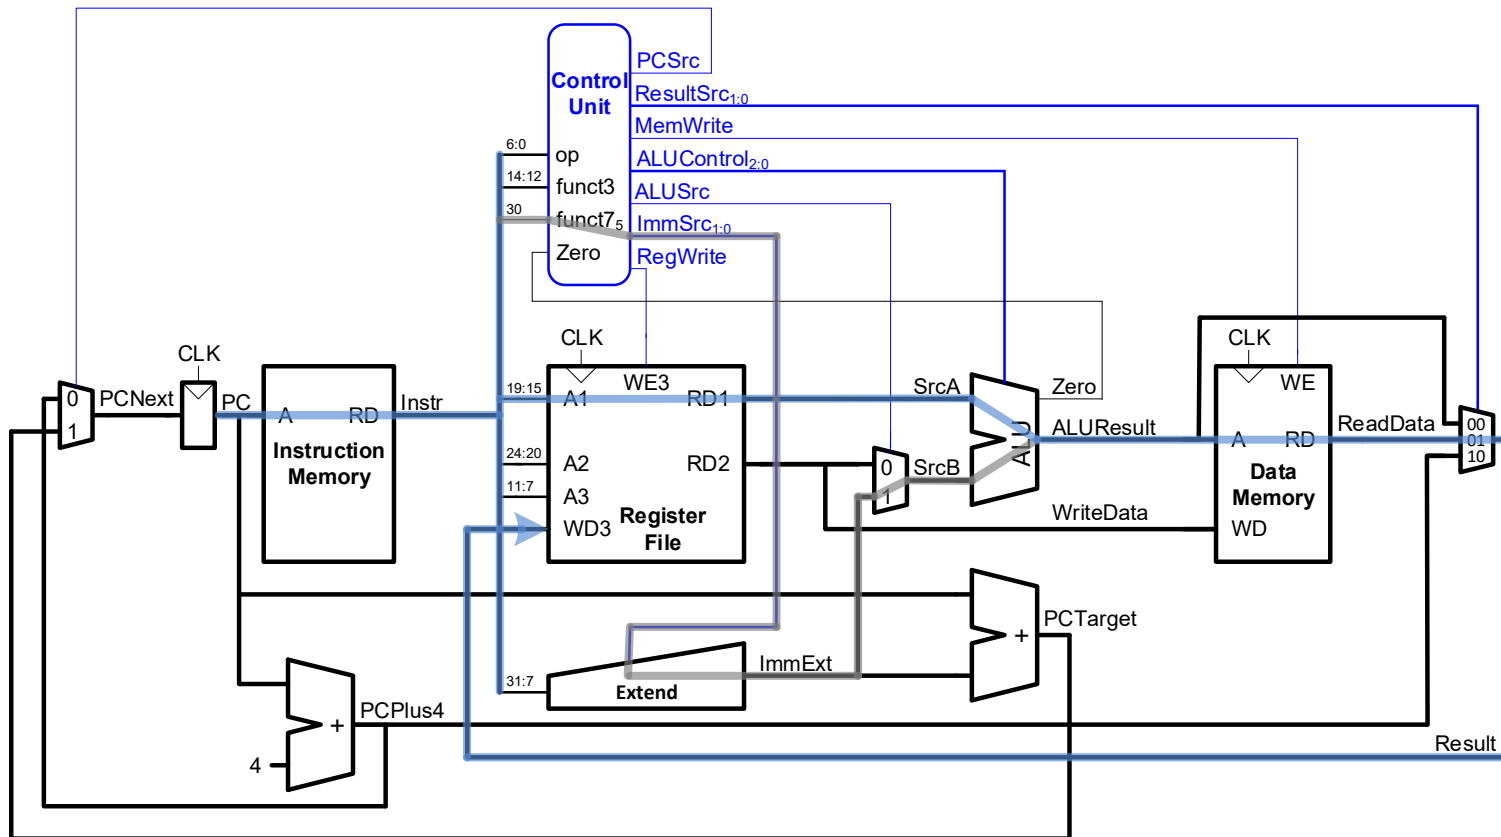
DDCA Ch7 - Part 6: RISC-V Single-Cycle Performance <https://www.youtube.com/watch?v=w82mNGranjA>

# Processor Performance

## Program Execution Time

$$\begin{aligned} &= \#instructions \times \frac{\text{cycles}}{\text{instruction}} \times \frac{\text{seconds}}{\text{cycle}} \\ &= \#instructions \times CPI \times T_C \end{aligned}$$

# Single-Cycle Processor Performance

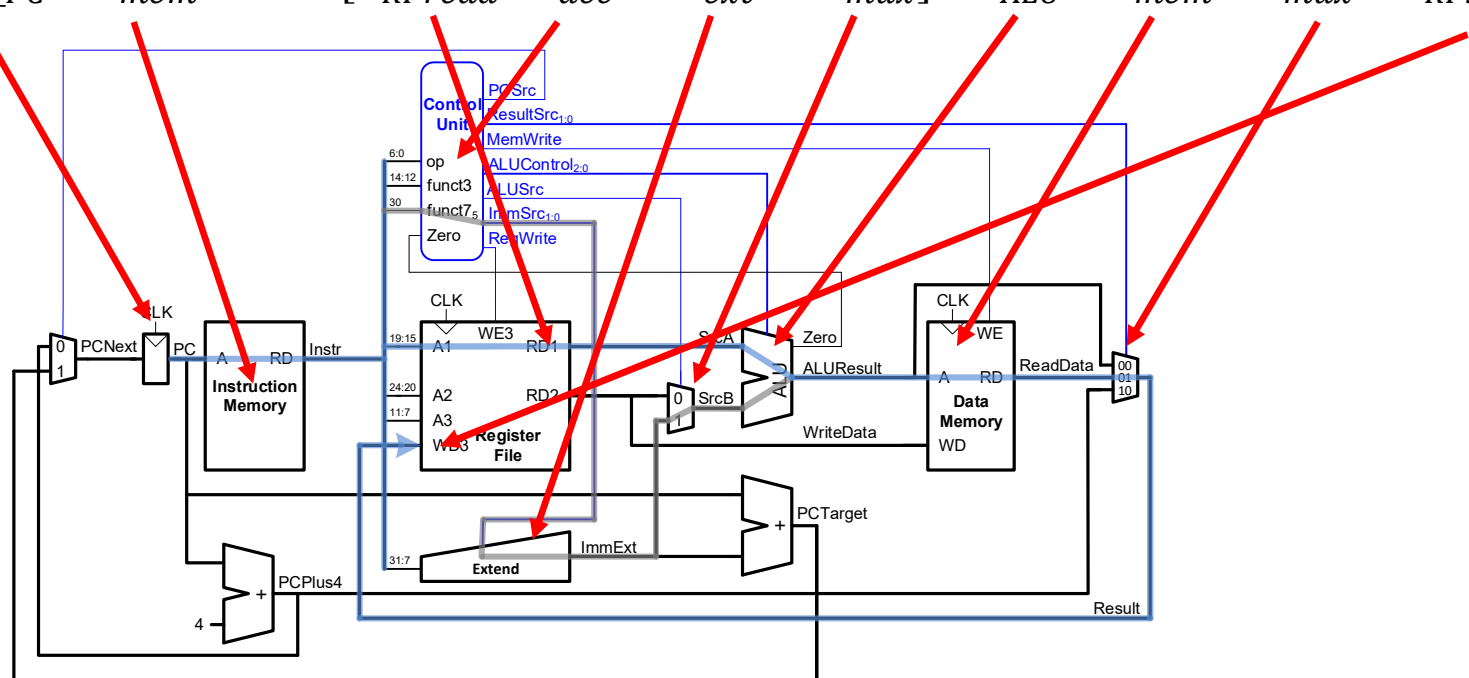


$T_c$  limited by critical path (**1w**)

# Single-Cycle Processor Performance

- Single-cycle critical path:

$$T_{c\_single} = t_{pcq\_PC} + t_{mem} + \max [t_{RFread}, t_{dec} + t_{ext} + t_{mux}] + t_{ALU} + t_{mem} + t_{mux} + t_{RFsetup}$$



## Single-Cycle Processor Performance

- **Single-cycle critical path:**

$$T_{c\_single} = t_{pcq\_PC} + t_{mem} + \max [ t_{RFread}, t_{dec} + t_{ext} + t_{mux} ] + t_{ALU} + t_{mem} + t_{mux} + t_{RFsetup}$$

- **Typically, limiting paths are:**

- memory, ALU, register file
- So,  $t_{dec} + t_{ext} + t_{mux}$  can be neglected

$$T_{c\_single} = t_{pcq\_PC} + t_{mem} + t_{RFread} + t_{ALU} + t_{mem} + t_{mux} + t_{RFsetup}$$

$$T_{c\_single} = t_{pcq\_PC} + 2t_{mem} + t_{RFread} + t_{ALU} + t_{mux} + t_{RFsetup}$$



## Single-Cycle Performance Example

Element	Parameter	Delay (ps)
Register clock-to-Q	$t_{pcq\_PC}$	40
Register setup	$t_{setup}$	50
Multiplexer	$t_{mux}$	30
AND-OR gate	$t_{AND-OR}$	20
ALU	$t_{ALU}$	120
Decoder (Control Unit)	$t_{dec}$	25
Extend unit	$t_{ext}$	35
Memory read	$t_{mem}$	200
Register file read	$t_{RFread}$	100
Register file setup	$t_{RFsetup}$	60

$$T_{c\_single} = t_{pcq\_PC} + 2t_{mem} + t_{RFread} + t_{ALU} + t_{mux} + t_{RFsetup}$$
$$= (40 + 2*200 + 100 + 120 + 30 + 60) \text{ ps} = \mathbf{750 \text{ ps}}$$

## Single-Cycle Performance Example

Program with 100 billion instructions =  $10^{11}$  instructions :

$$\begin{aligned}\text{Execution Time} &= \#instructions \times CPI \times T_c \\ &= (100 \times 10^9)(1)(750 \times 10^{-12} \text{ s}) \\ &= \mathbf{75 \text{ seconds}}\end{aligned}$$

# Multicycle RISC-V Processor

---

DDCA Ch7 - Part 7: Multicycle Processor <https://www.youtube.com/watch?v=sATaQNCC0-g>

# Single- vs. Multicycle Processor

- **Single-cycle:**
  - + simple
  - - cycle time limited by longest instruction ( $1_w$ )
  - - separate memories for instruction and data
  - - 3 adders/ALUs
- **Multicycle processor** addresses these issues by breaking instruction into shorter steps
  - shorter instructions take **fewer steps**
  - can re-use hardware
  - cycle time is faster

# Single- vs. Multicycle Processor

- **Single-cycle:**

- + simple
- - cycle time limited by longest instruction ( $1w$ )
- - separate memories for instruction and data
- - 3 adders/ALUs

- **Multicycle**

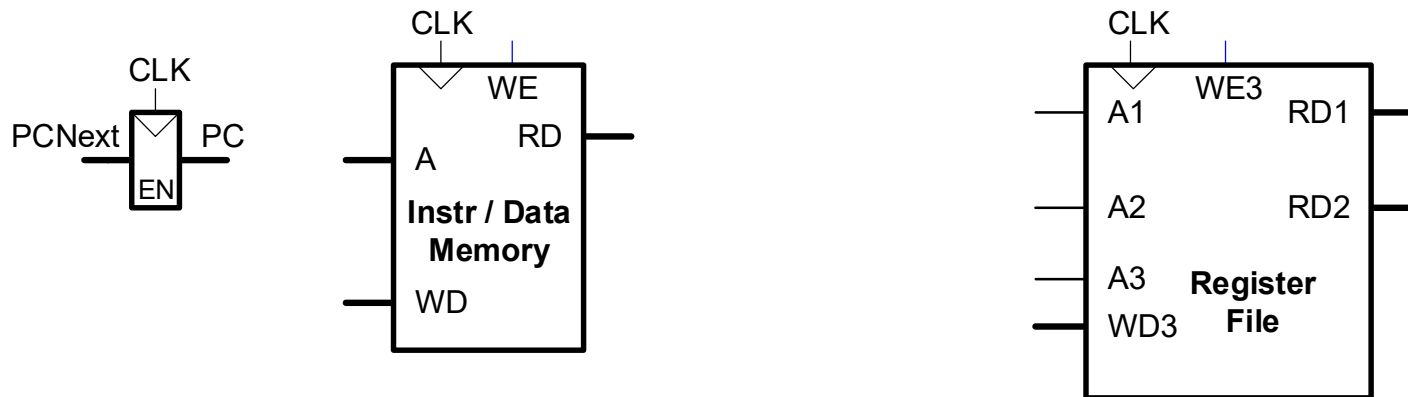
- + higher clock speed
- + simpler instructions run faster
- + reuse expensive hardware on multiple cycles
- - sequencing overhead paid many times

**Same design steps  
as single-cycle:**

- first datapath
- then control

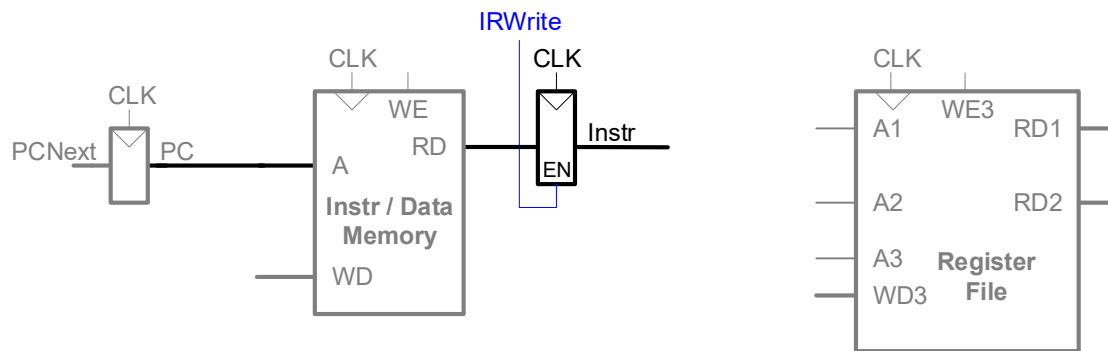
## Multicycle State Elements

- Replace separate Instruction and Data memories with a **single unified memory** – more realistic



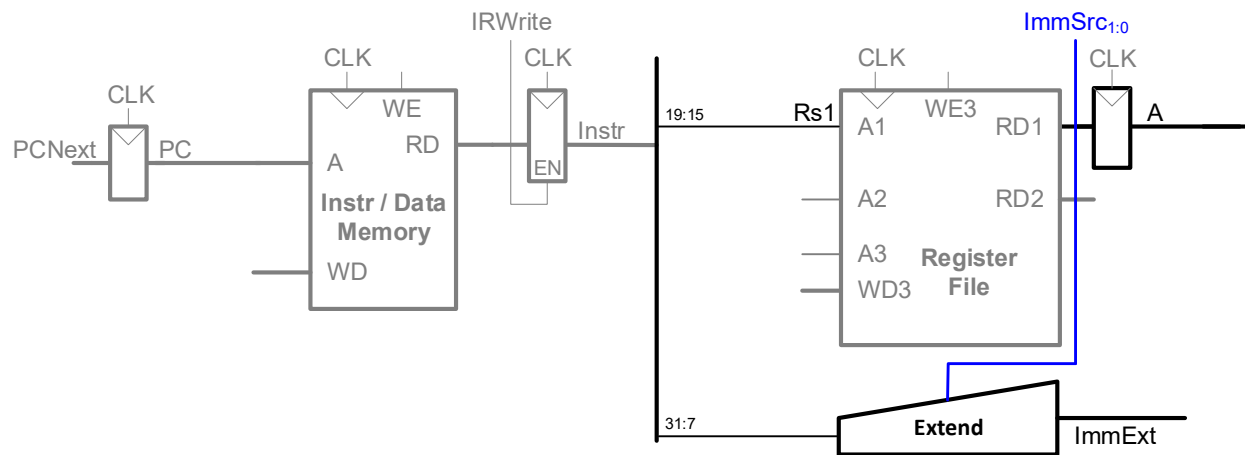
# Multicycle Datapath: Instruction Fetch

## STEP 1: Fetch instruction



## Multicycle Datapath: $l_w$ Get Sources

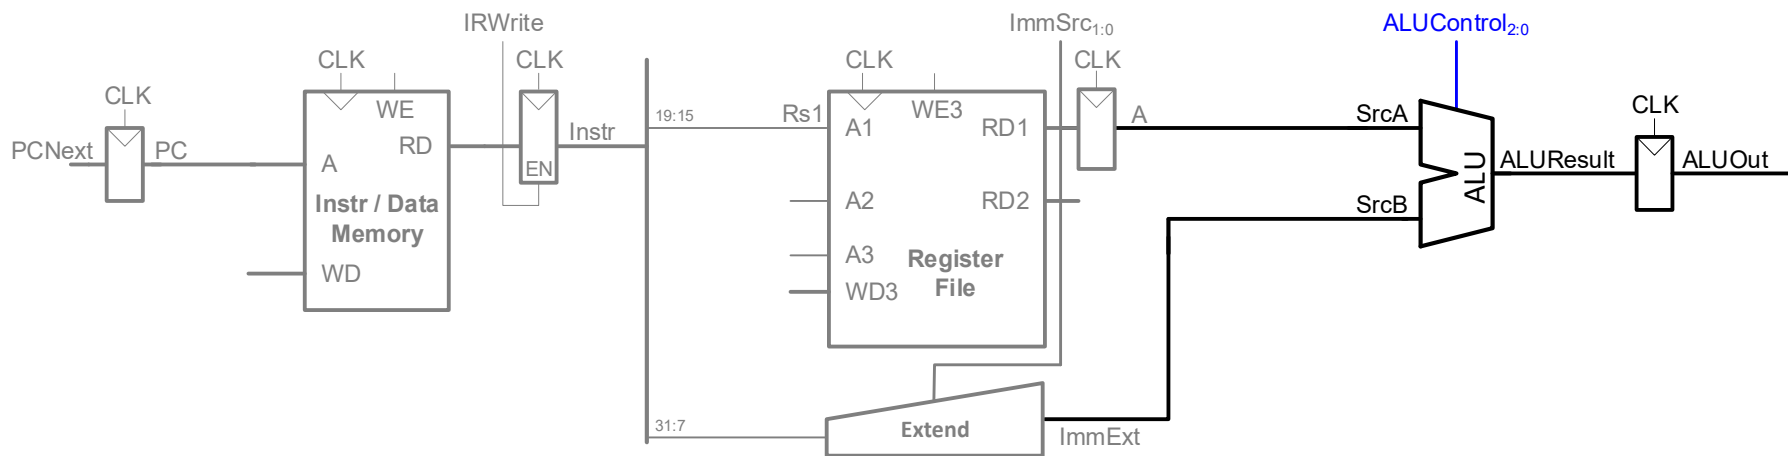
**STEP 2:** Read source operand from RF and extend immediate





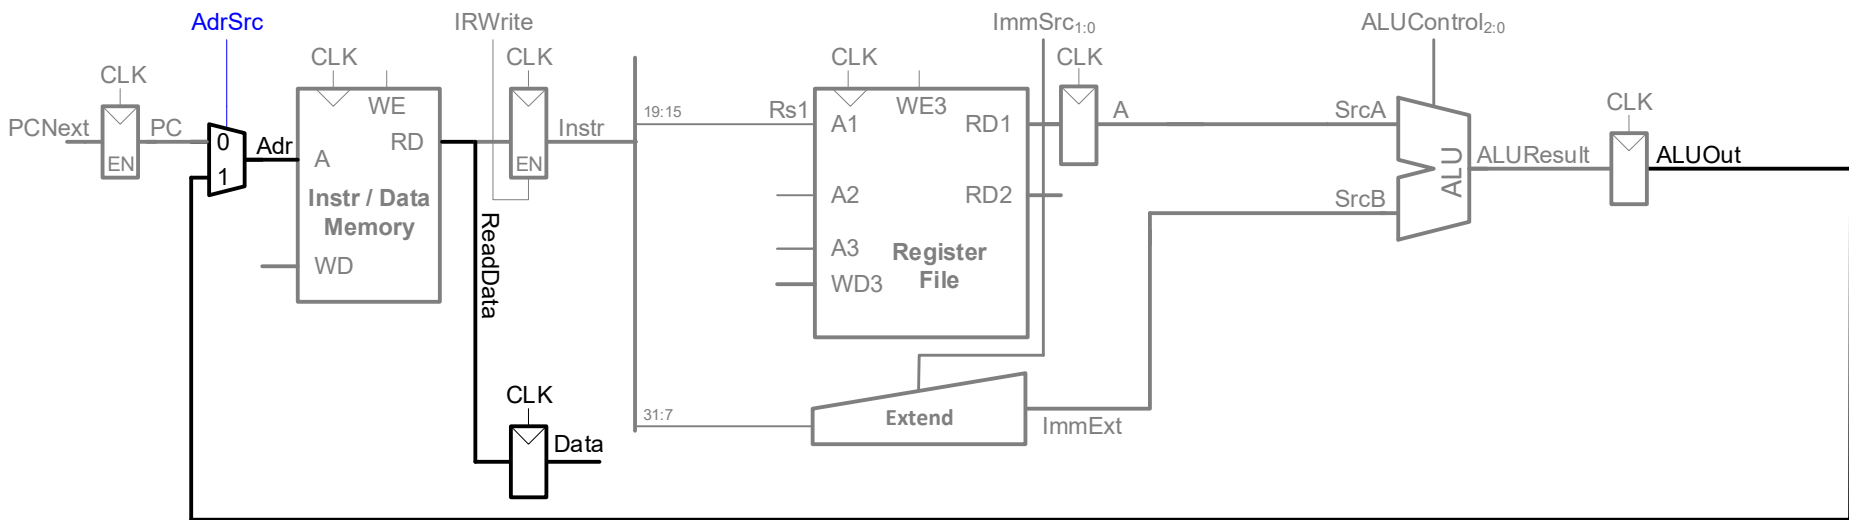
## Multicycle Datapath: $1_W$ Address

### STEP 3: Compute the memory address



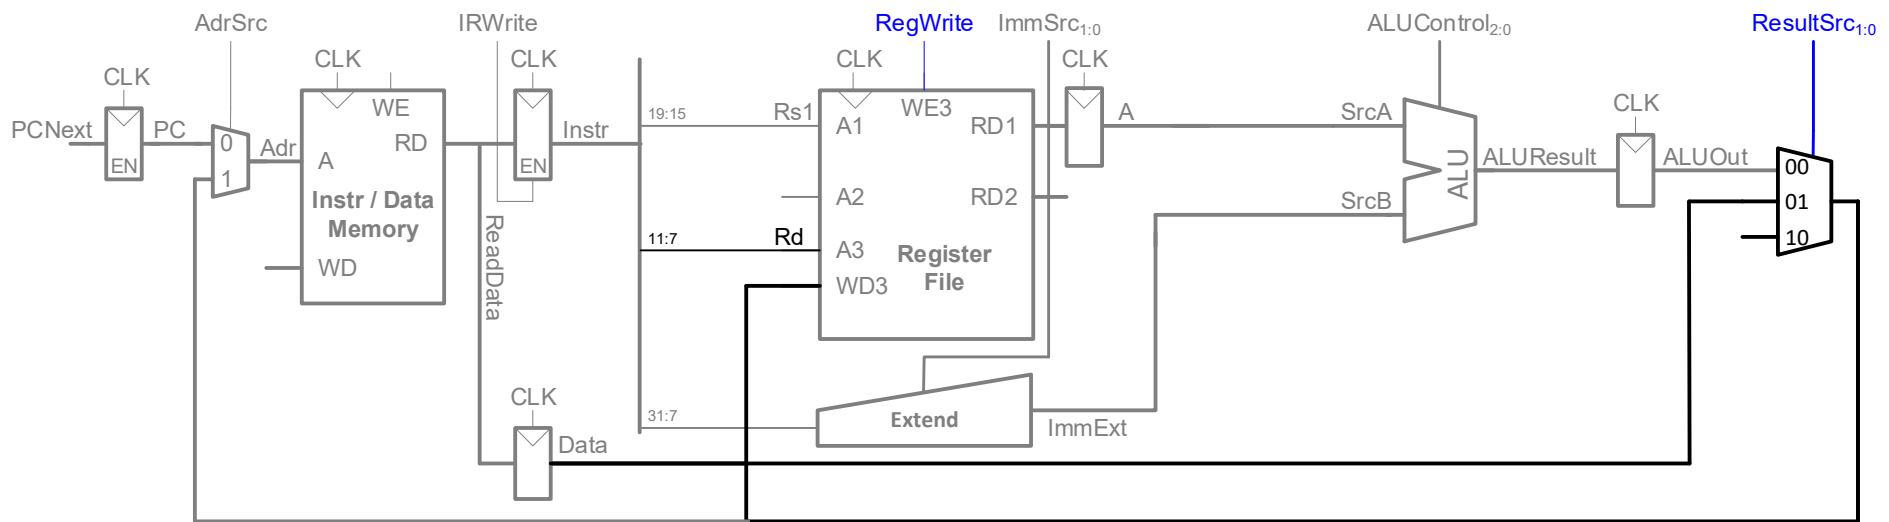
# Multicycle Datapath: $1_w$ Memory Read

## STEP 4: Read data from memory



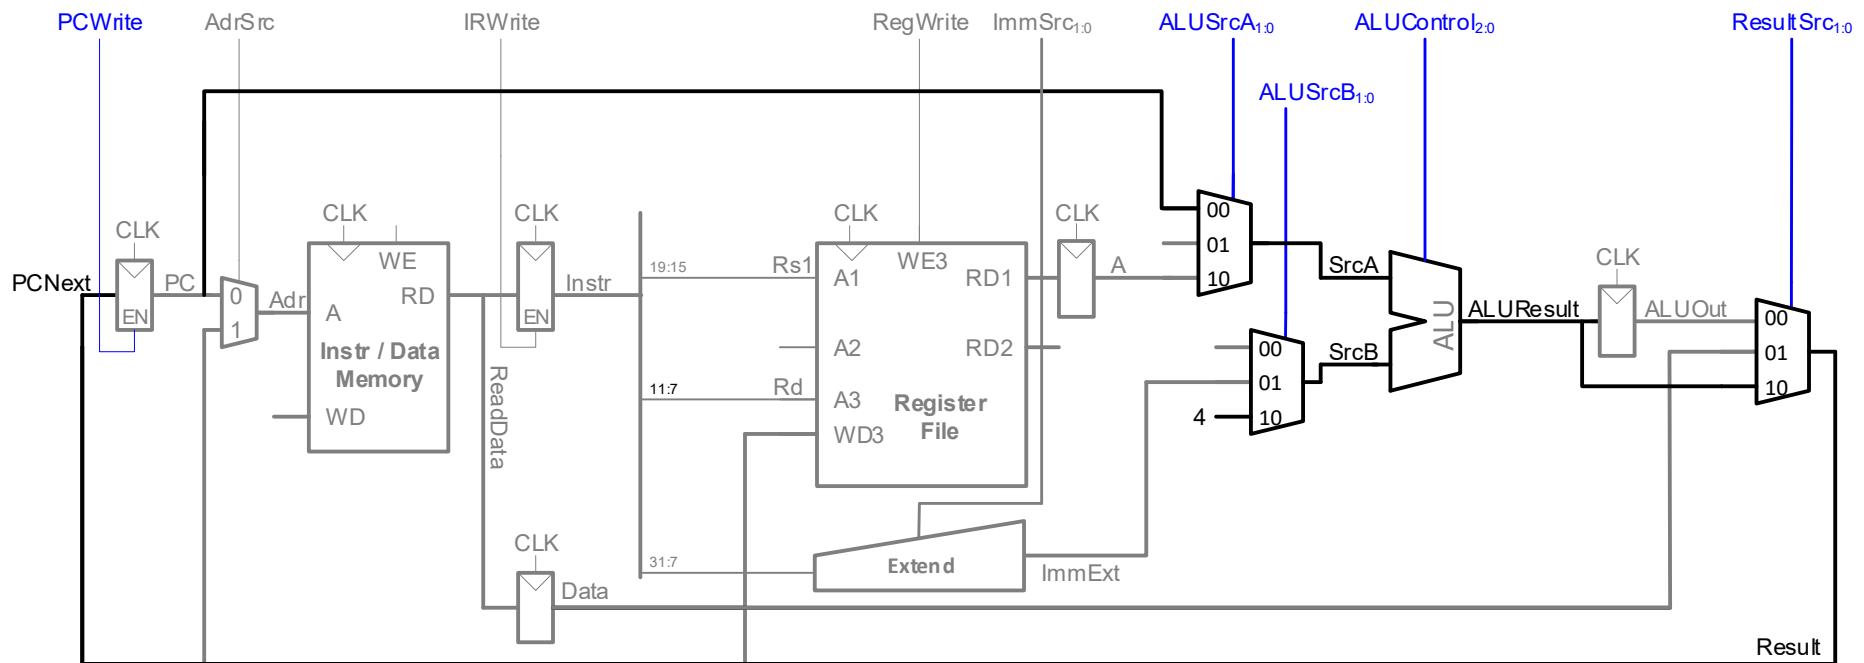
# Multicycle Datapath: $1_W$ Write Register

## STEP 5: Write data back to register file



# Multicycle Datapath: Increment PC

## STEP 6: Increment PC: $PC = PC + 4$



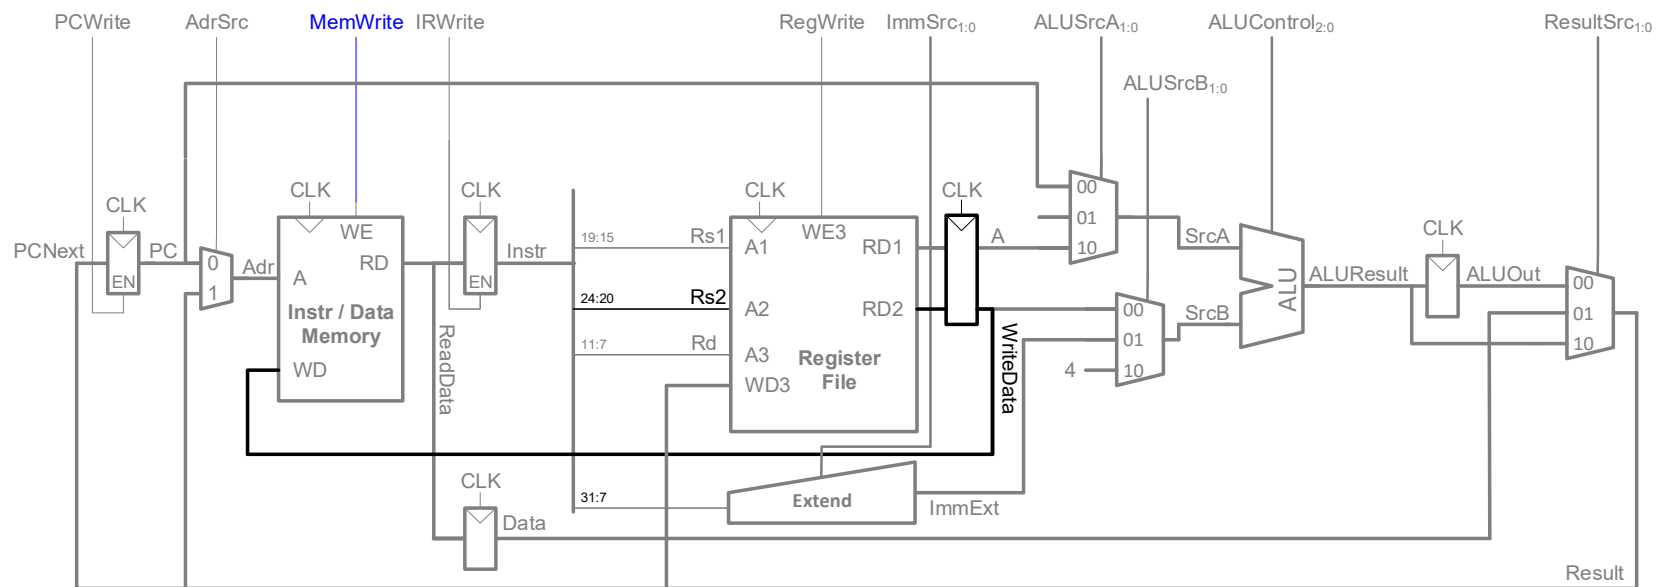
# Multicycle Datapath: Other Instructions

---

DDCA Ch7 - part 8: RISC-V Multicycle Processor - Other Instructions <https://www.youtube.com/watch?v=dnITBQQDmwU>

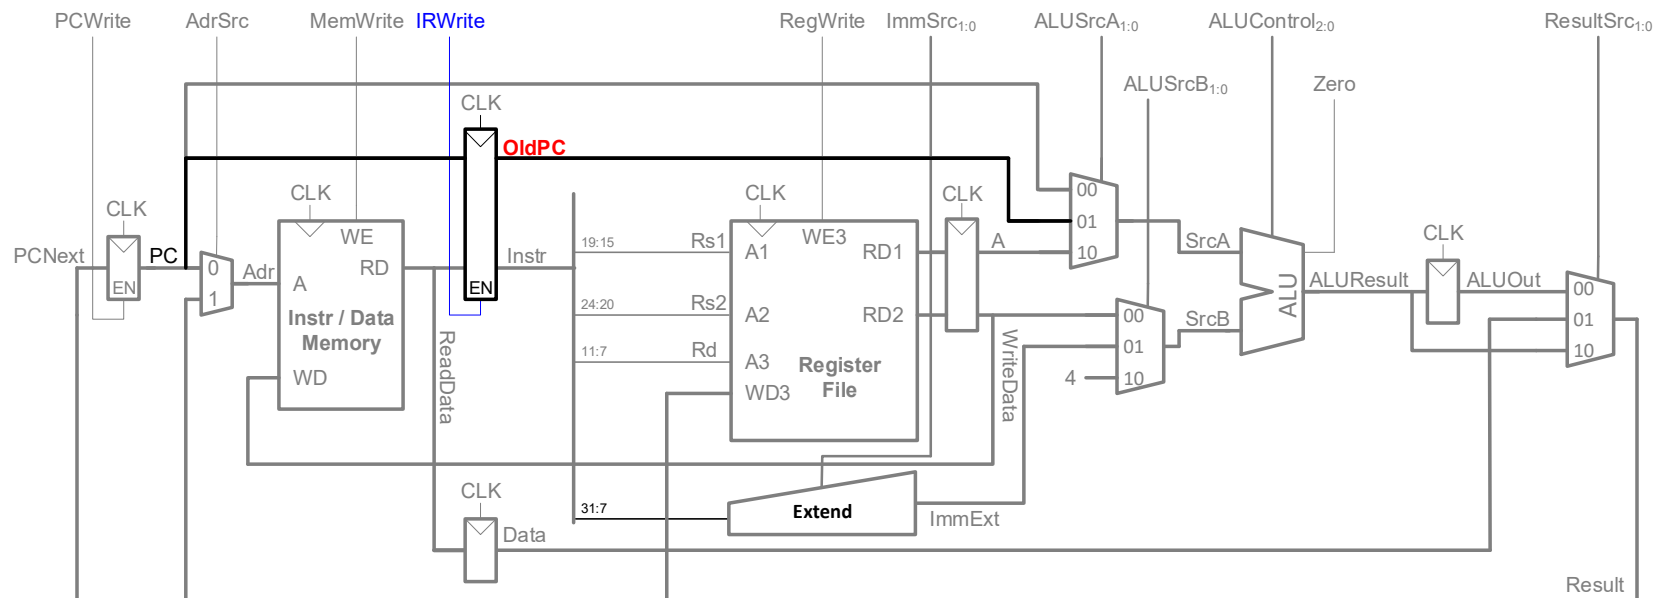
# Multicycle Datapath: $s_w$

## Write data from Register File ( $rs_2$ ) to memory



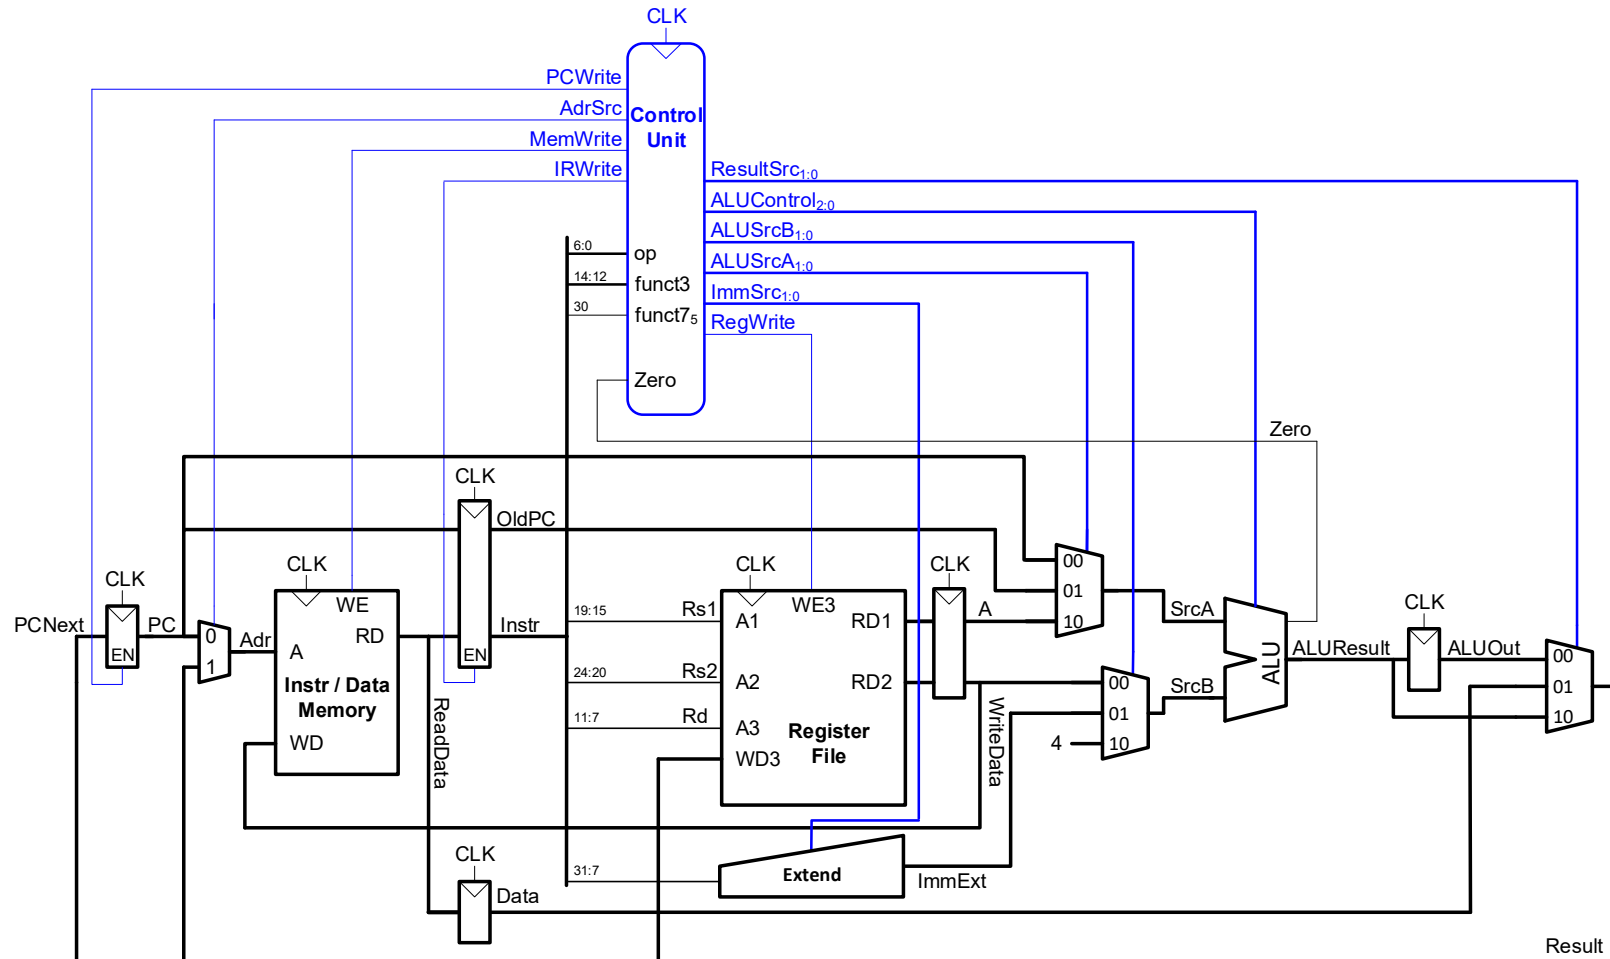
## Multicycle Datapath: beq

- Calculate branch target address:  $BTA = PC + imm$



PC is updated in Fetch stage, so need to save **old (current) PC**

# Multicycle RISC-V Processor





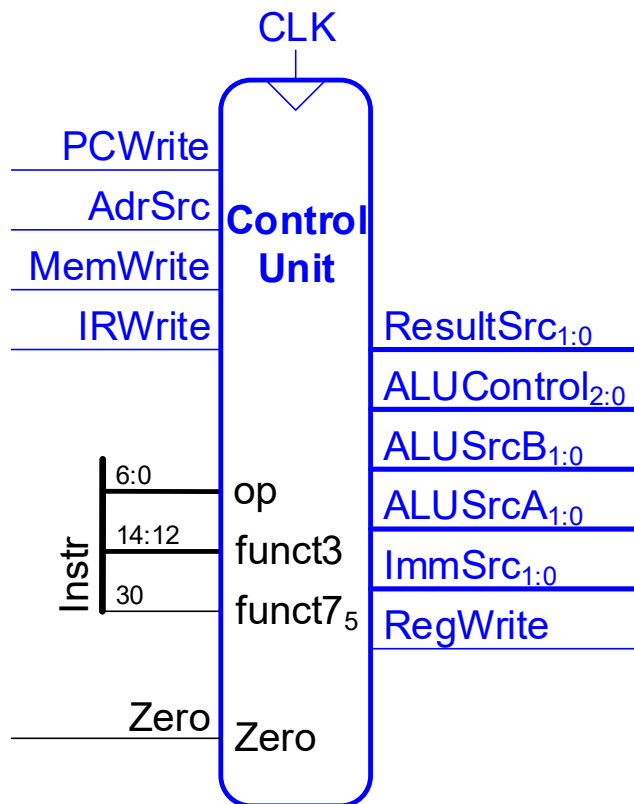
# Multicycle Control

---

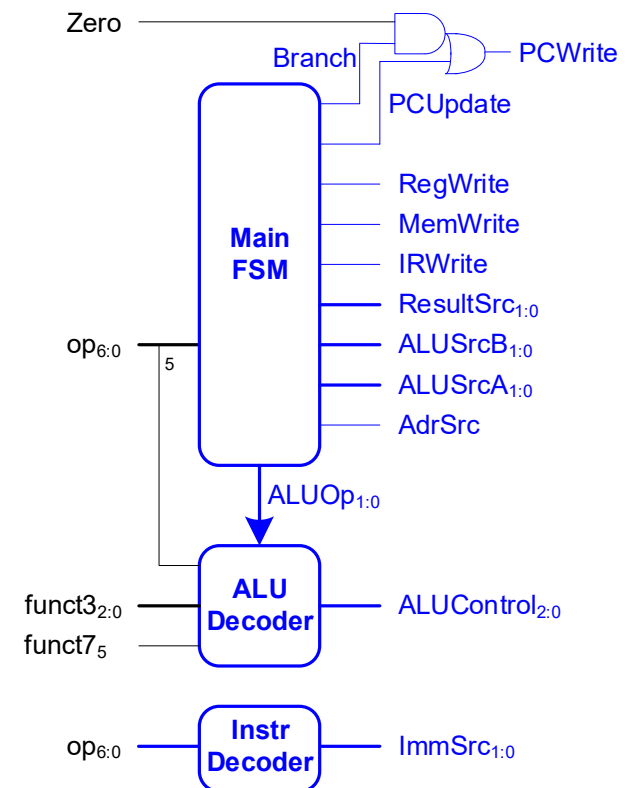
DDCA Ch7 - Part 9: RISC-V Multicycle Processor Control <https://www.youtube.com/watch?v=YUJhNTpunqI>

# Multicycle Control

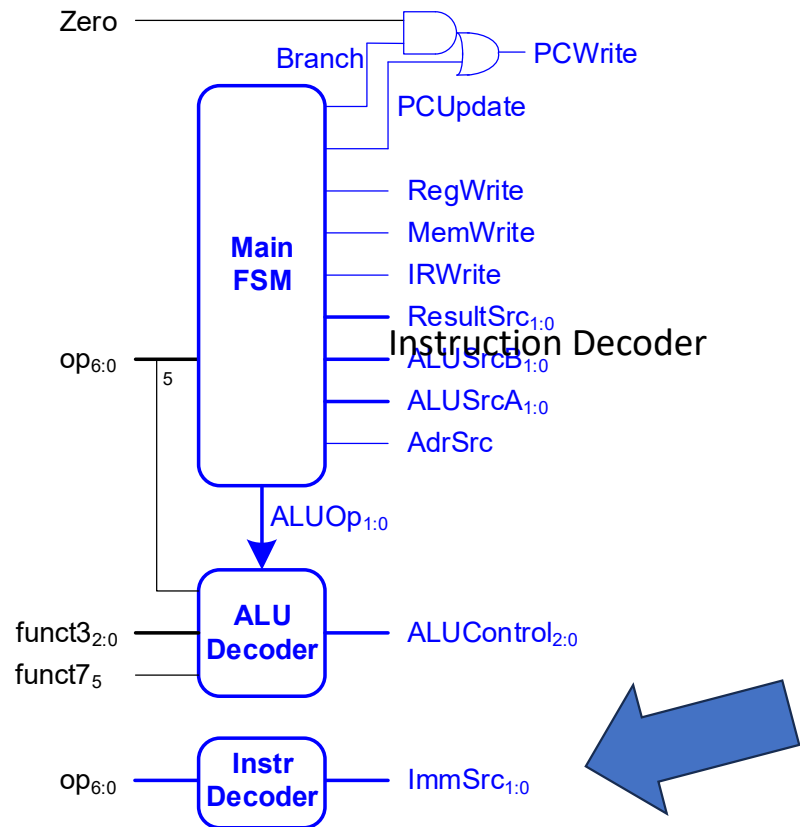
## High-Level View



## Low-Level View



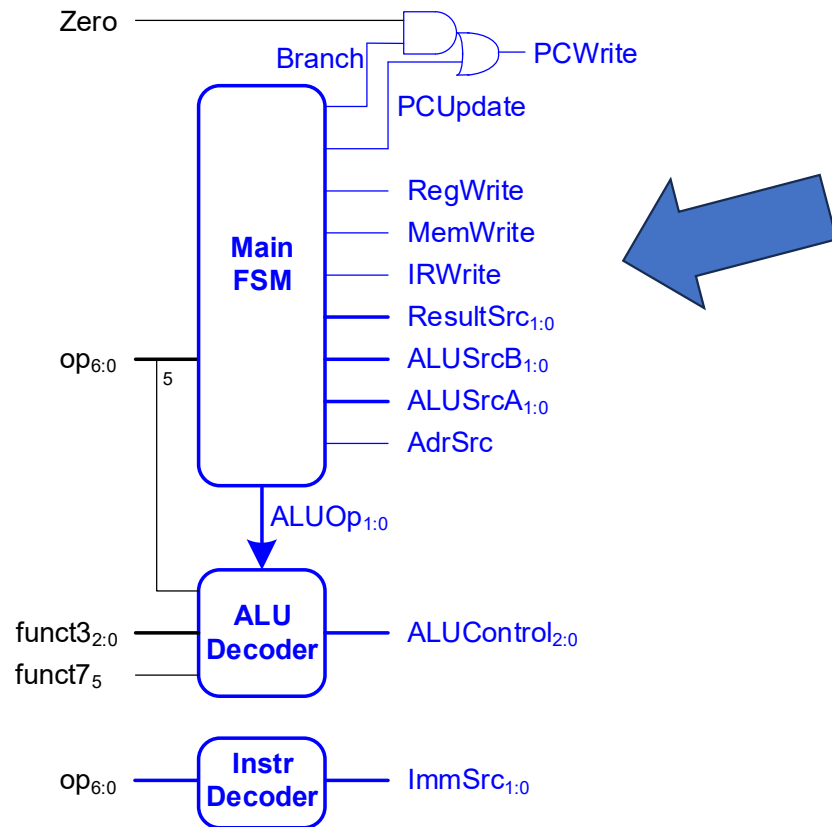
# Multicycle Control: Instruction Decoder



## Instruction Decoder

op	Instruction	ImmSrc
3	<b>lw</b>	00
35	<b>sw</b>	01
51	R-type	XX
99	<b>beq</b>	10

# Multicycle Control: Main FSM

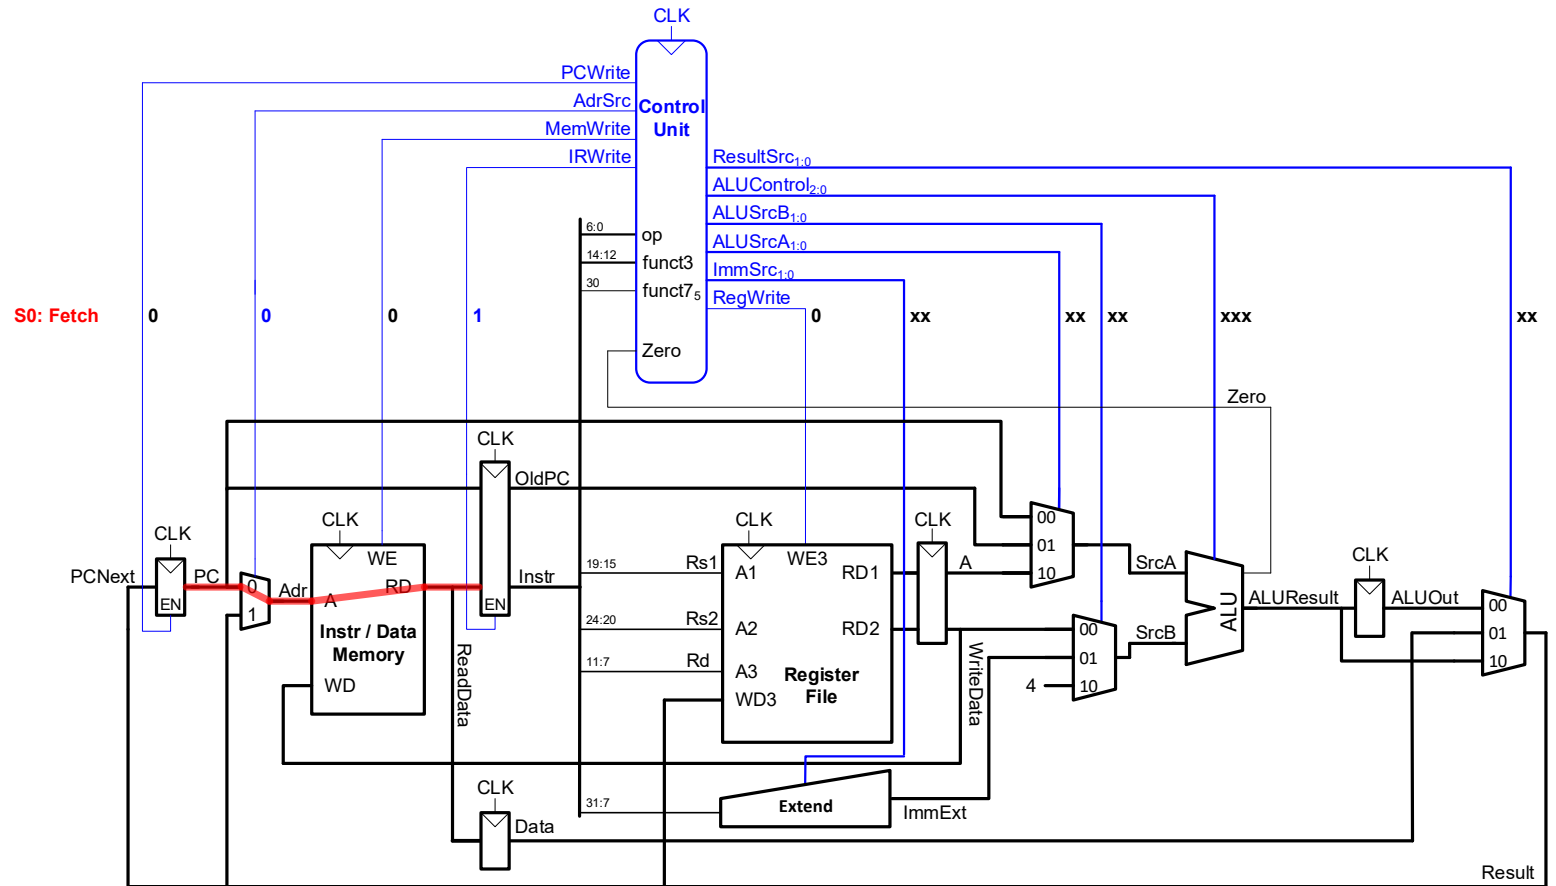



## Main FSM


To declutter FSM:

- **Write enable signals** (RegWrite, MemWrite, IRWrite, PCUpdate, and Branch) are **0** if not listed in a state.
- **Other signals are don't care** if not listed in a state

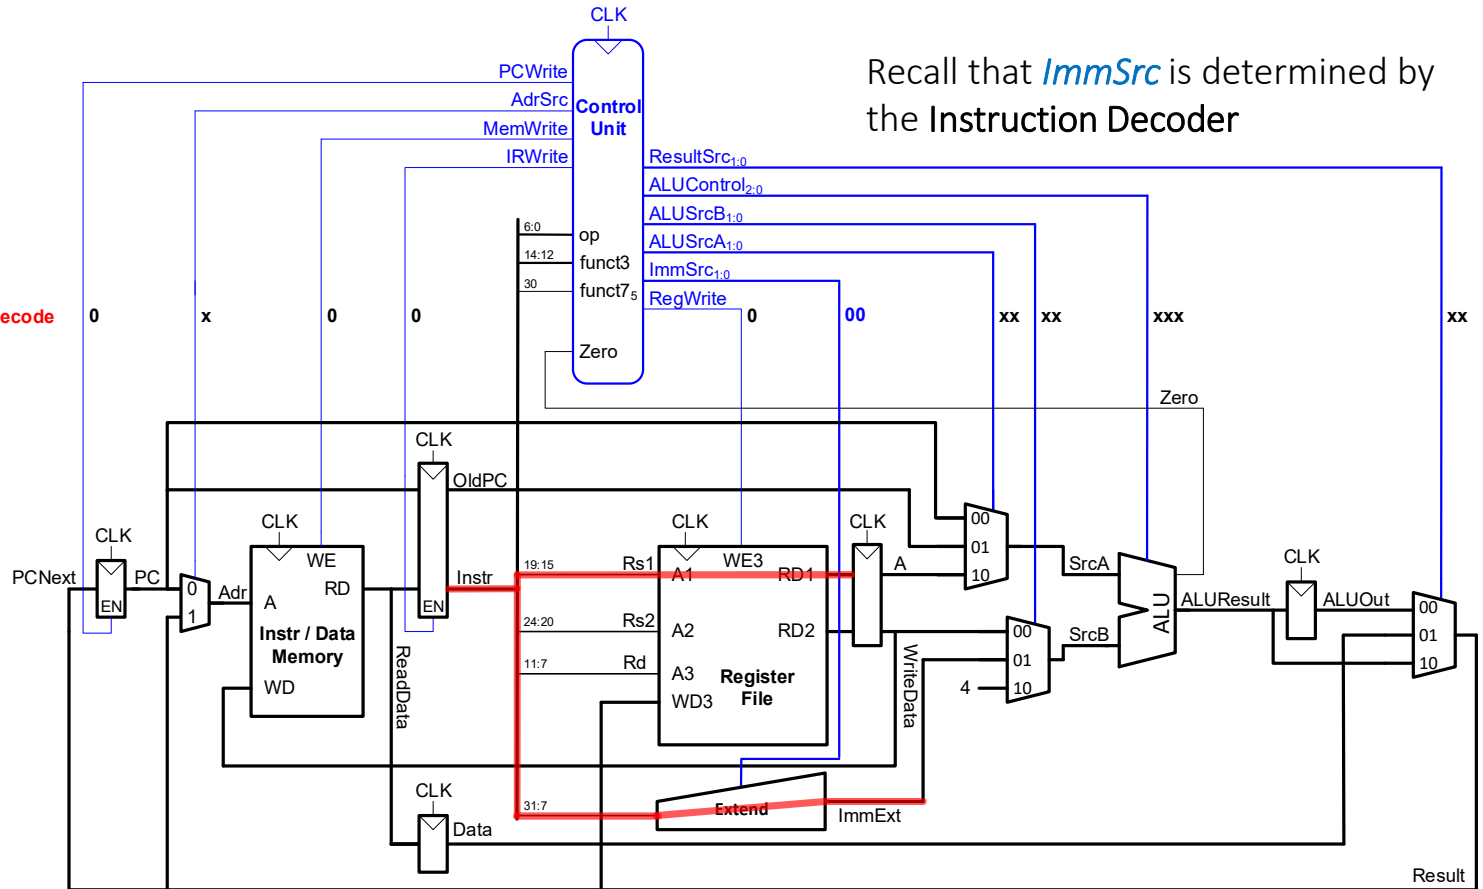
# Main FSM: Fetch



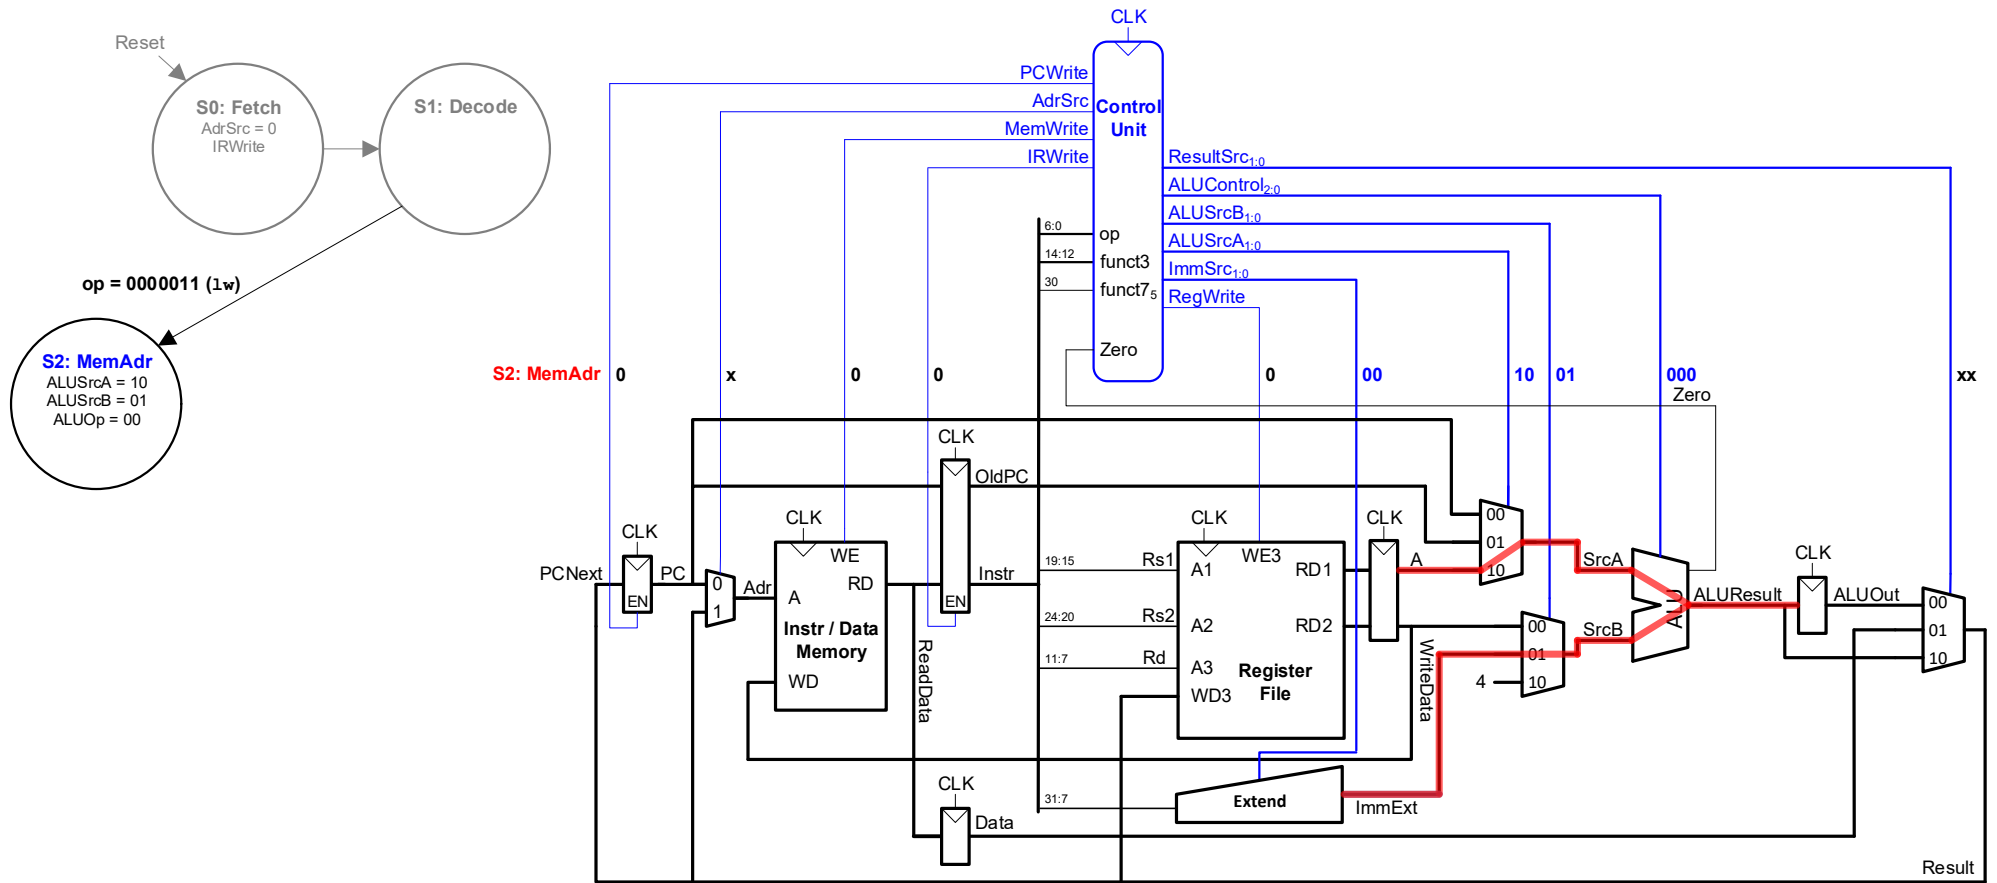
# Main FSM: Decode



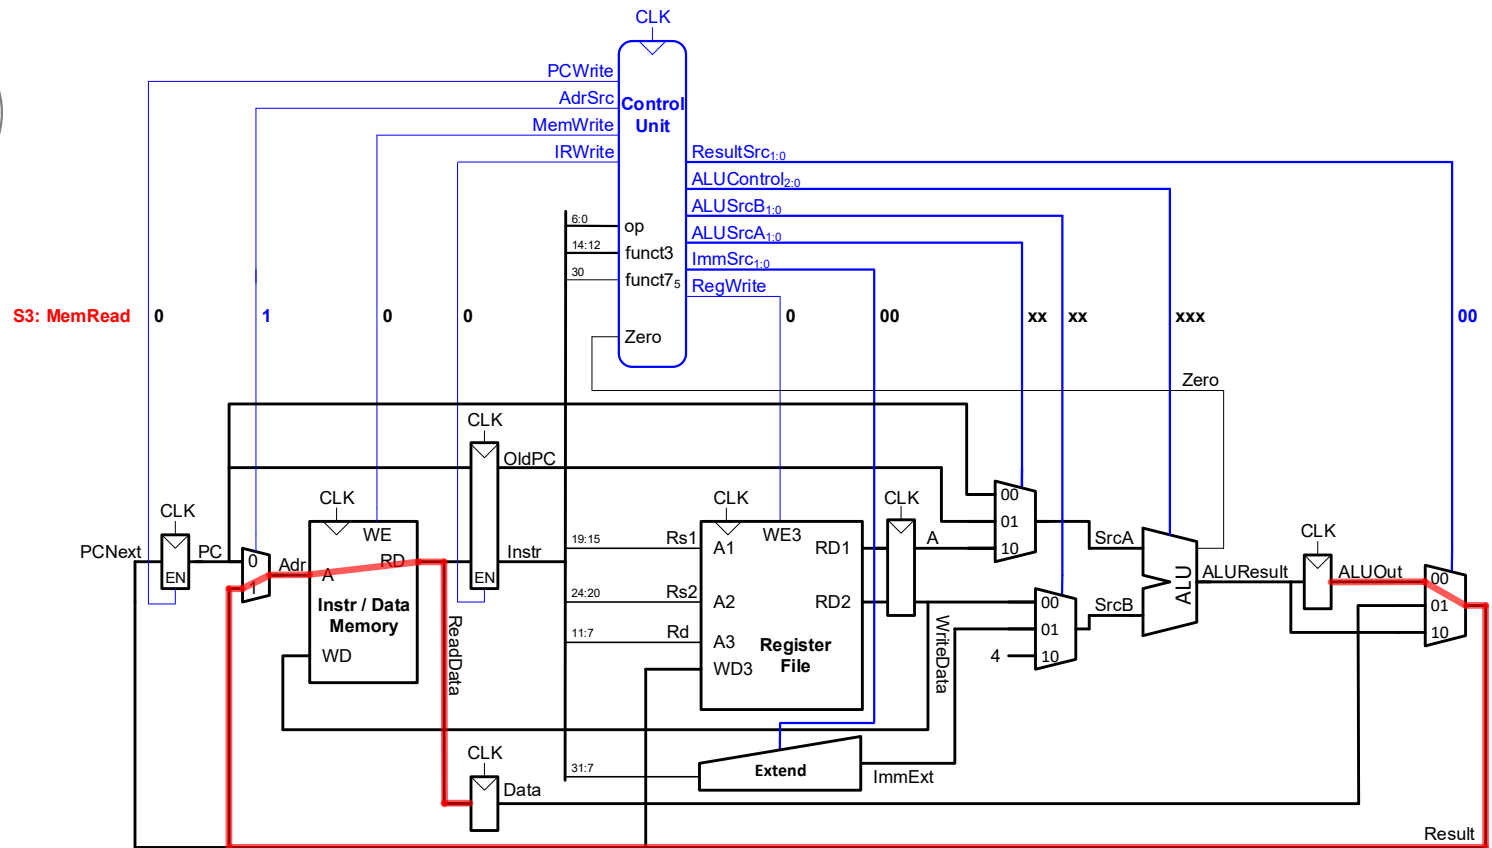
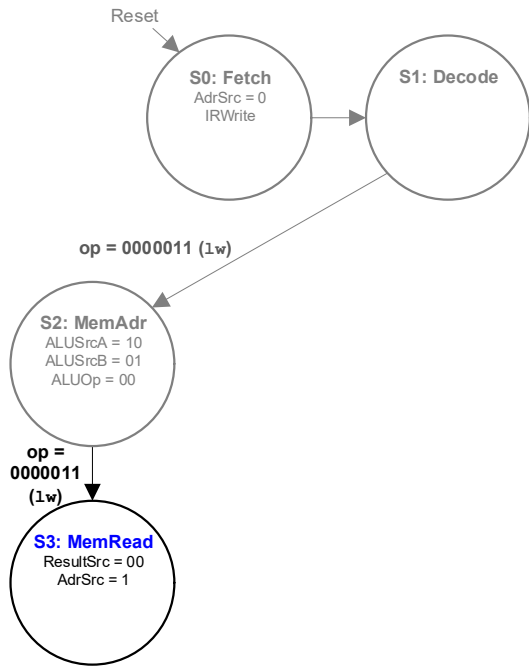
S1: Decode



# Main FSM: Address

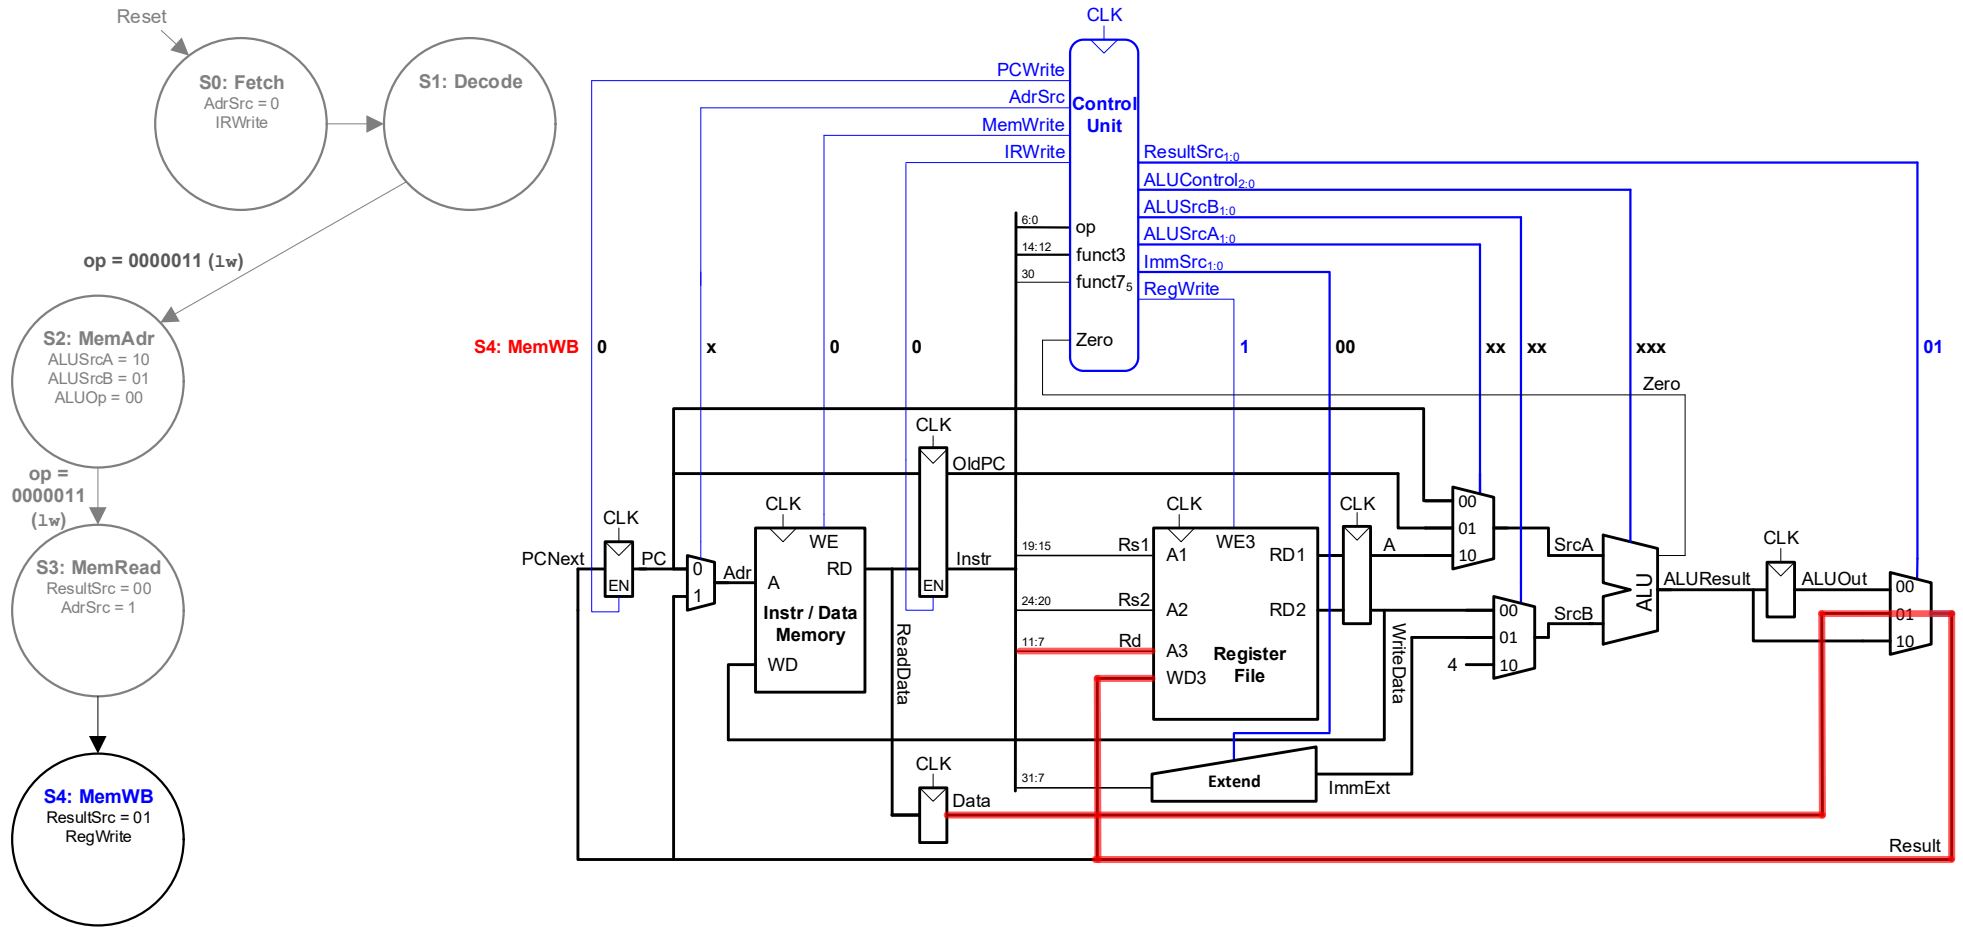


# Main FSM: Read Memory

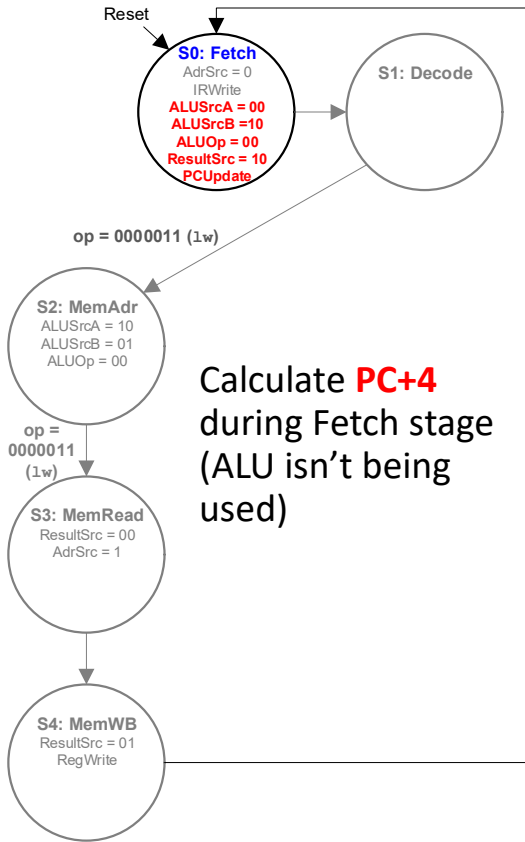




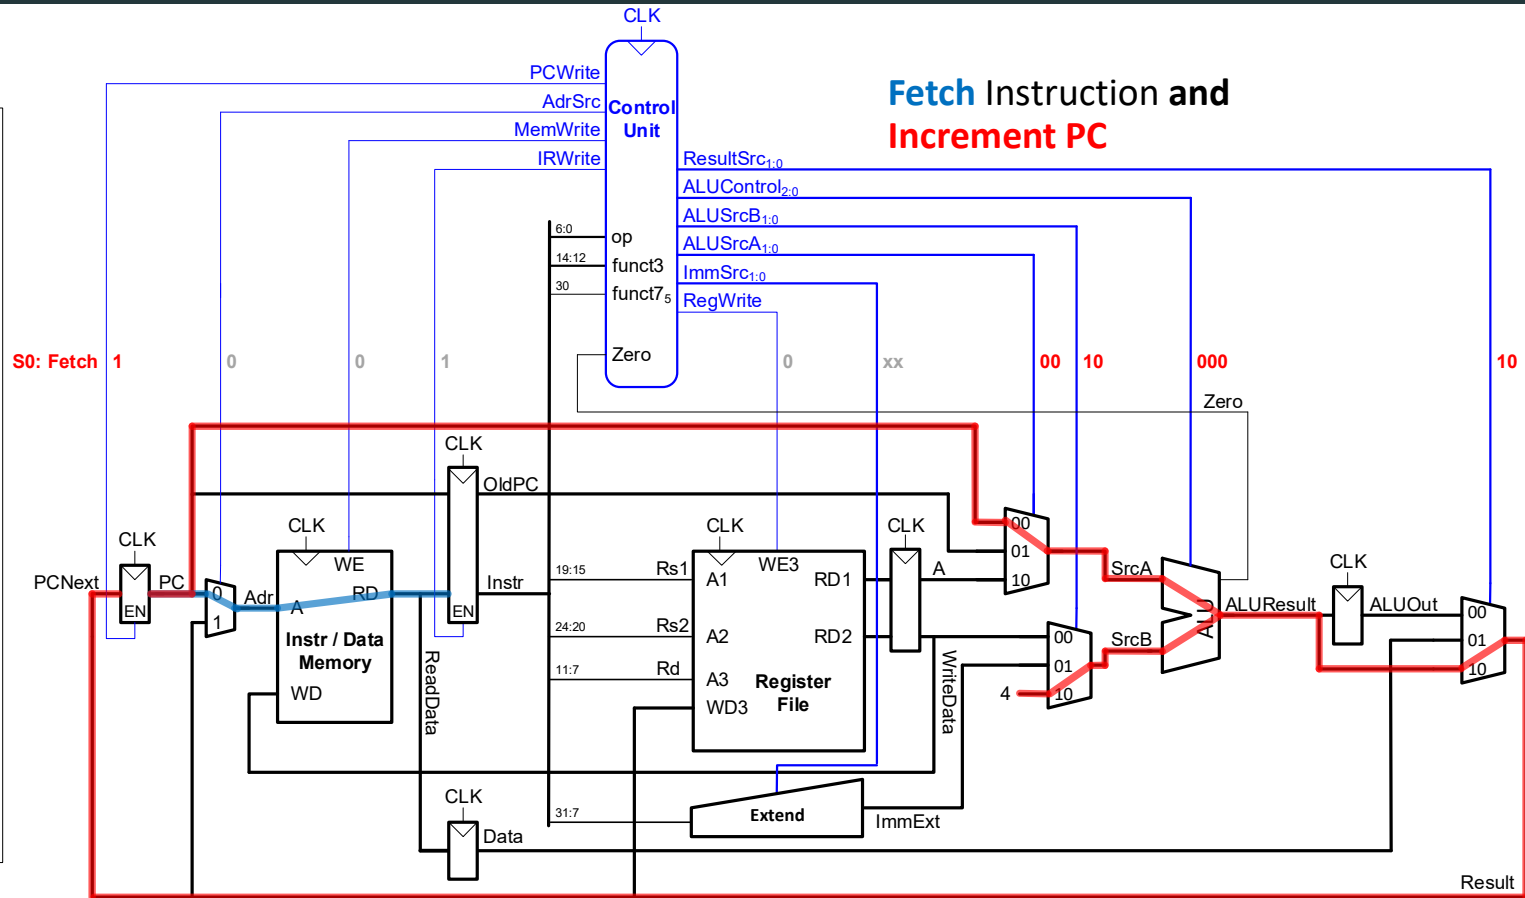
# Main FSM: Write RF



# Main FSM: Fetch Revisited



Calculate **PC+4** during Fetch stage (ALU isn't being used)



# Multicycle Control: Other Instructions

---

DDCA Ch7 - Part 10: RISC-V Multicycle Processor Control: Other Instructions

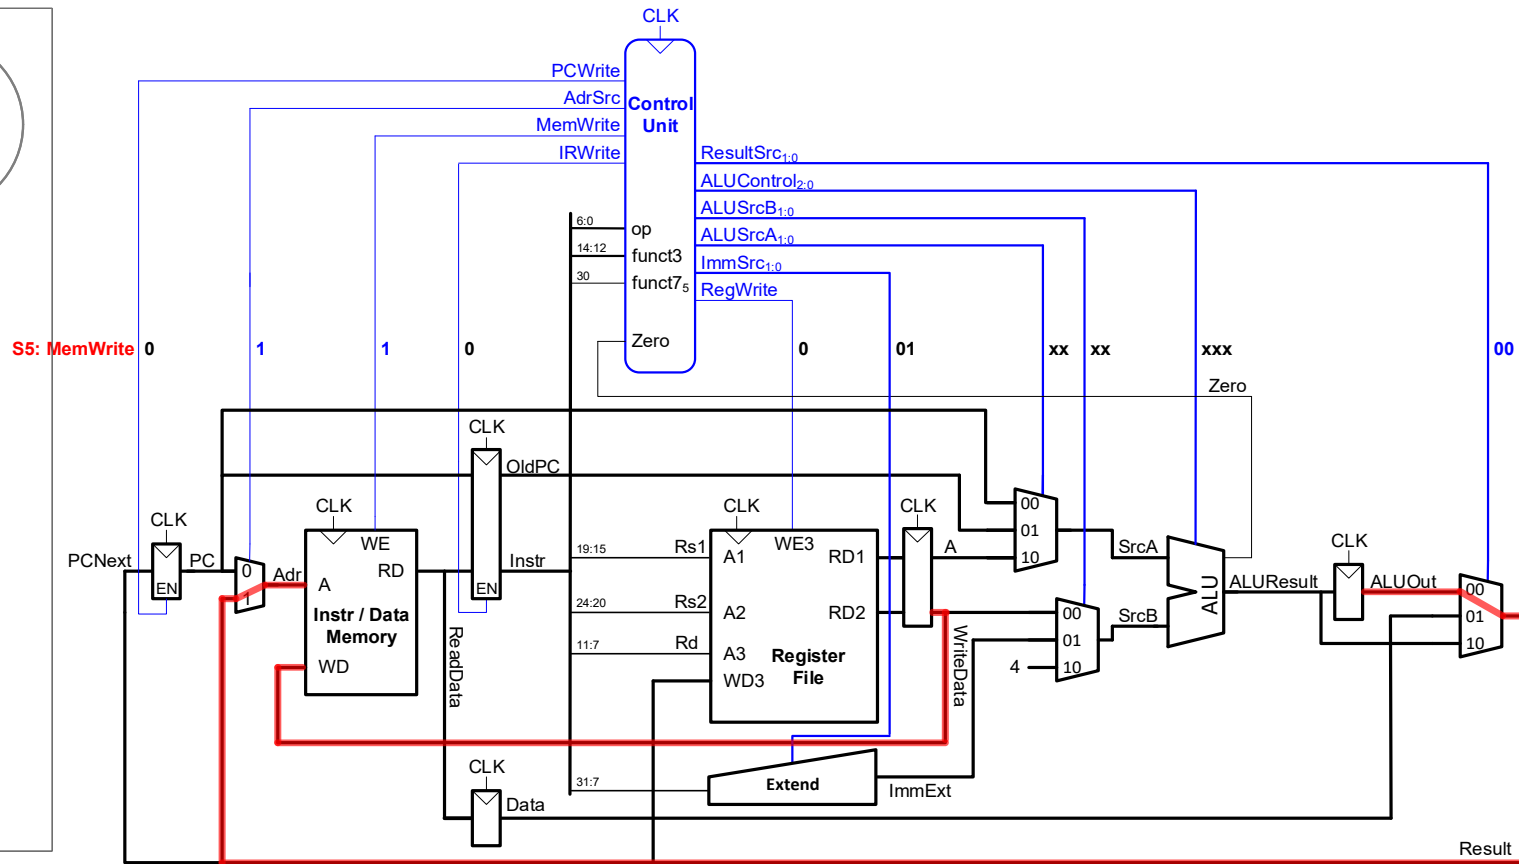
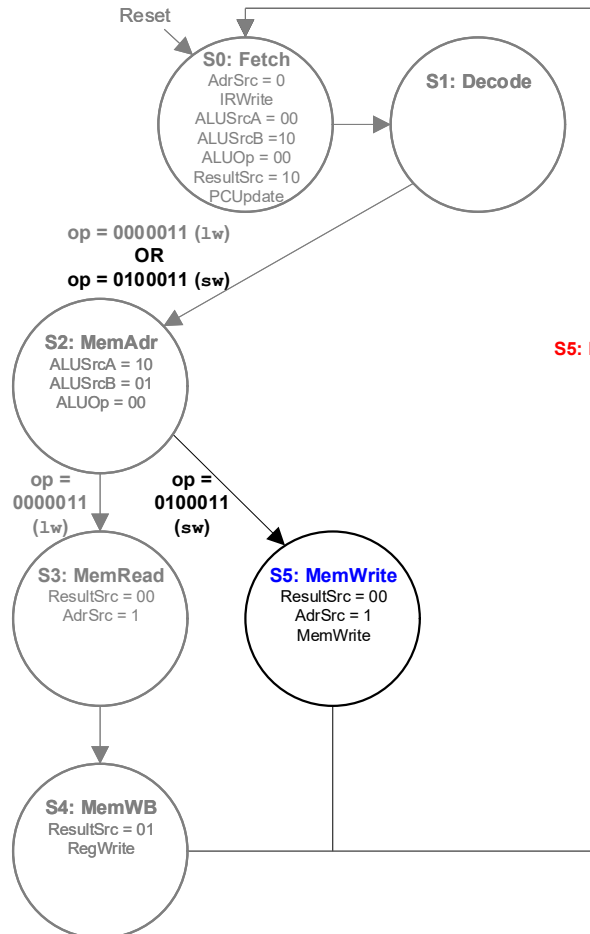
<https://www.youtube.com/watch?v=rSodrnsYXQ8>

08.04.2024

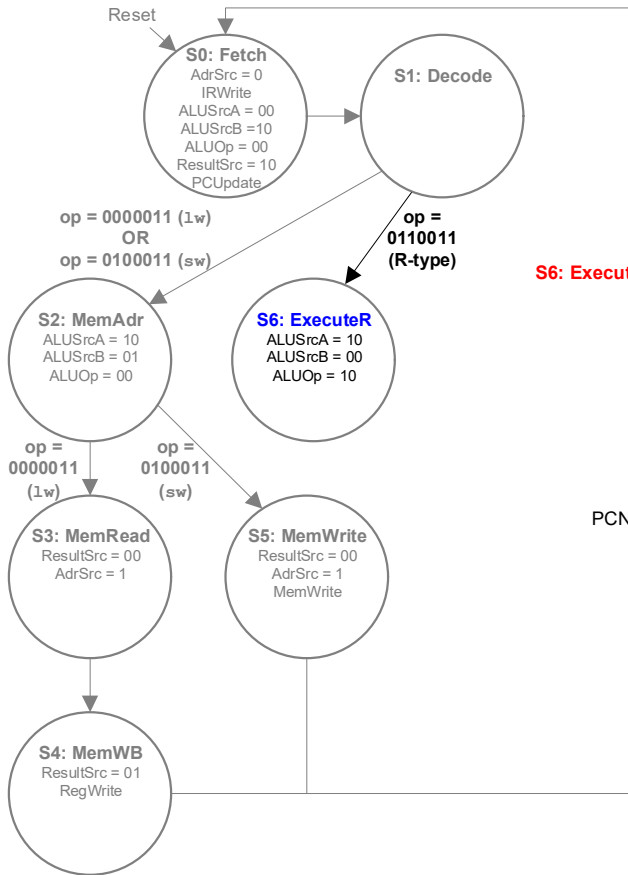
Computer Systems

75

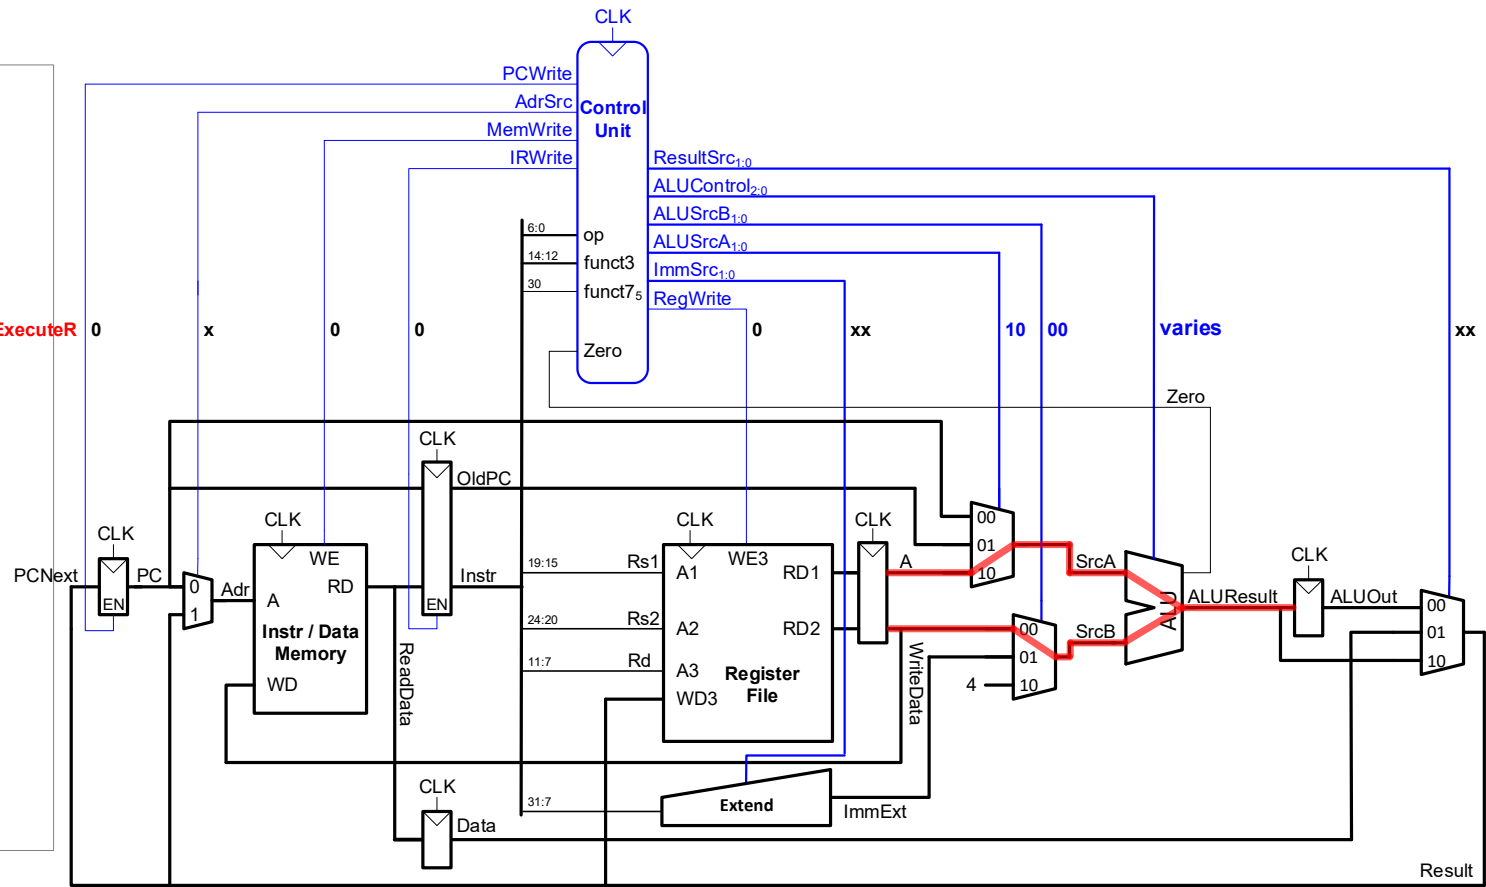
# Main FSM: *sw*



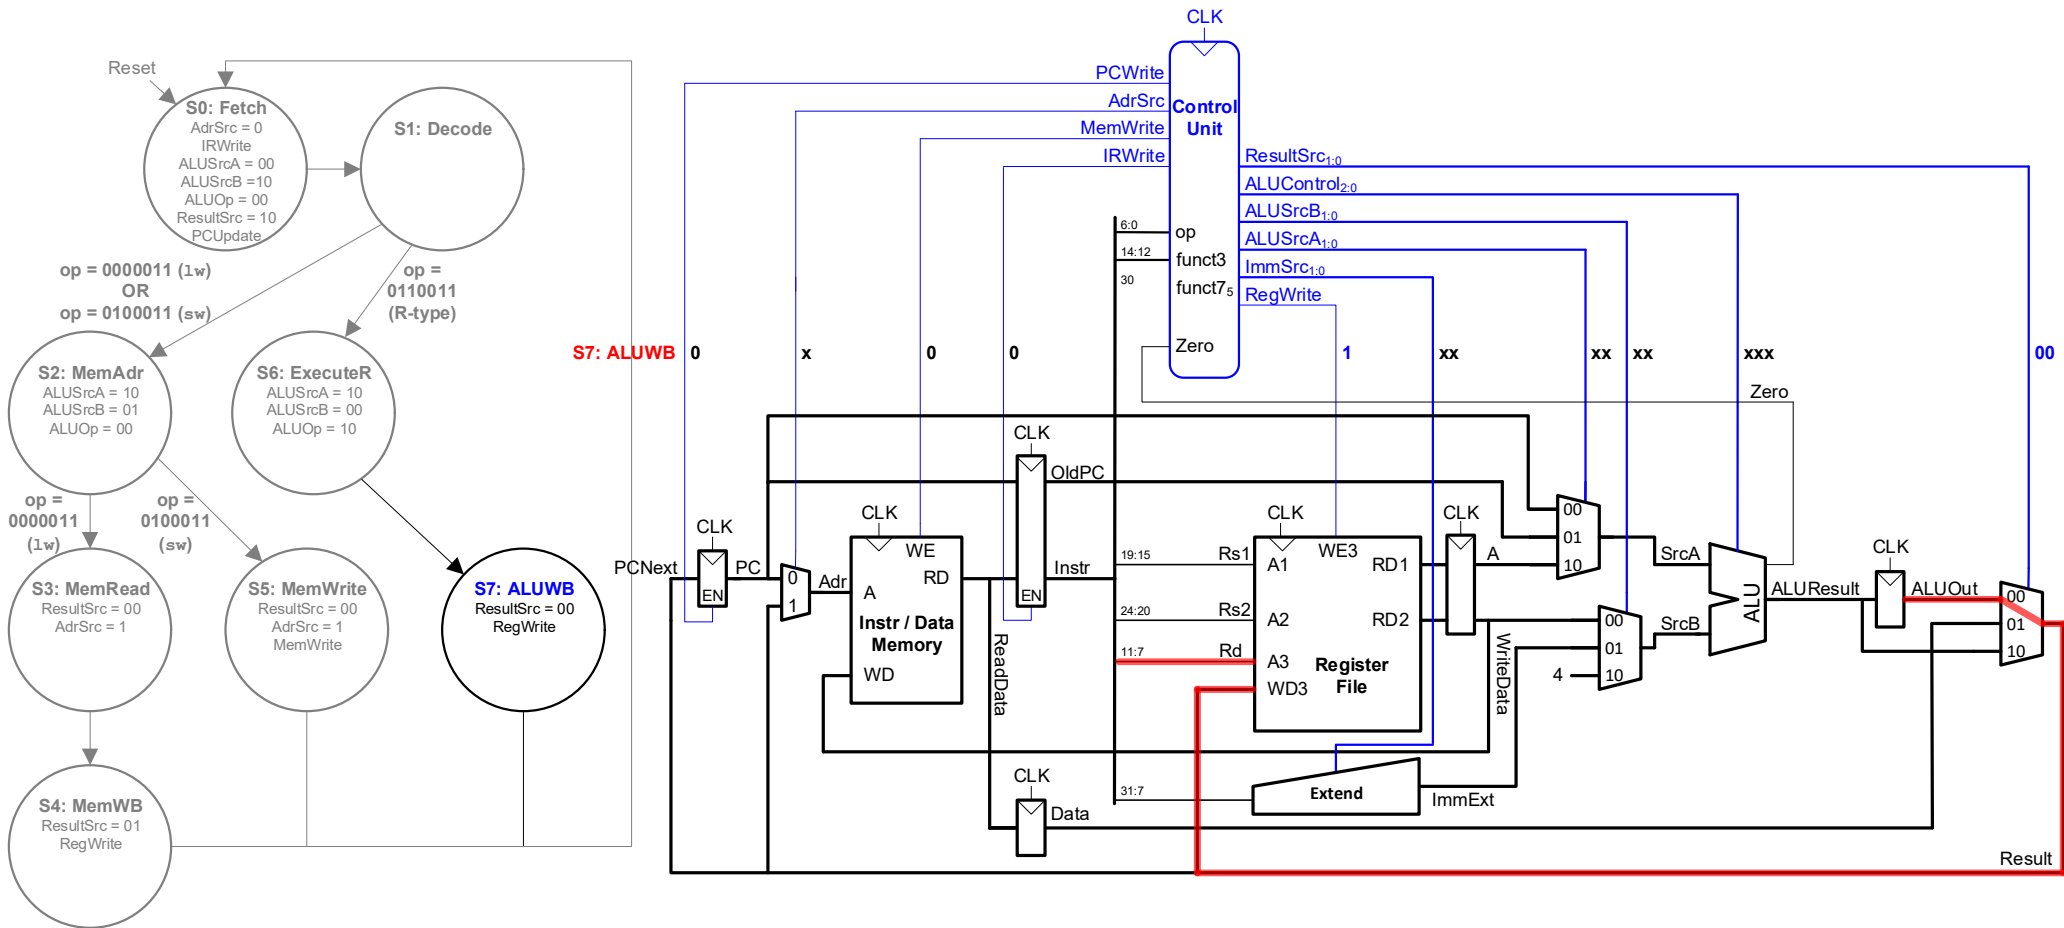
# Main FSM: R-type Execute



S6: ExecuteR

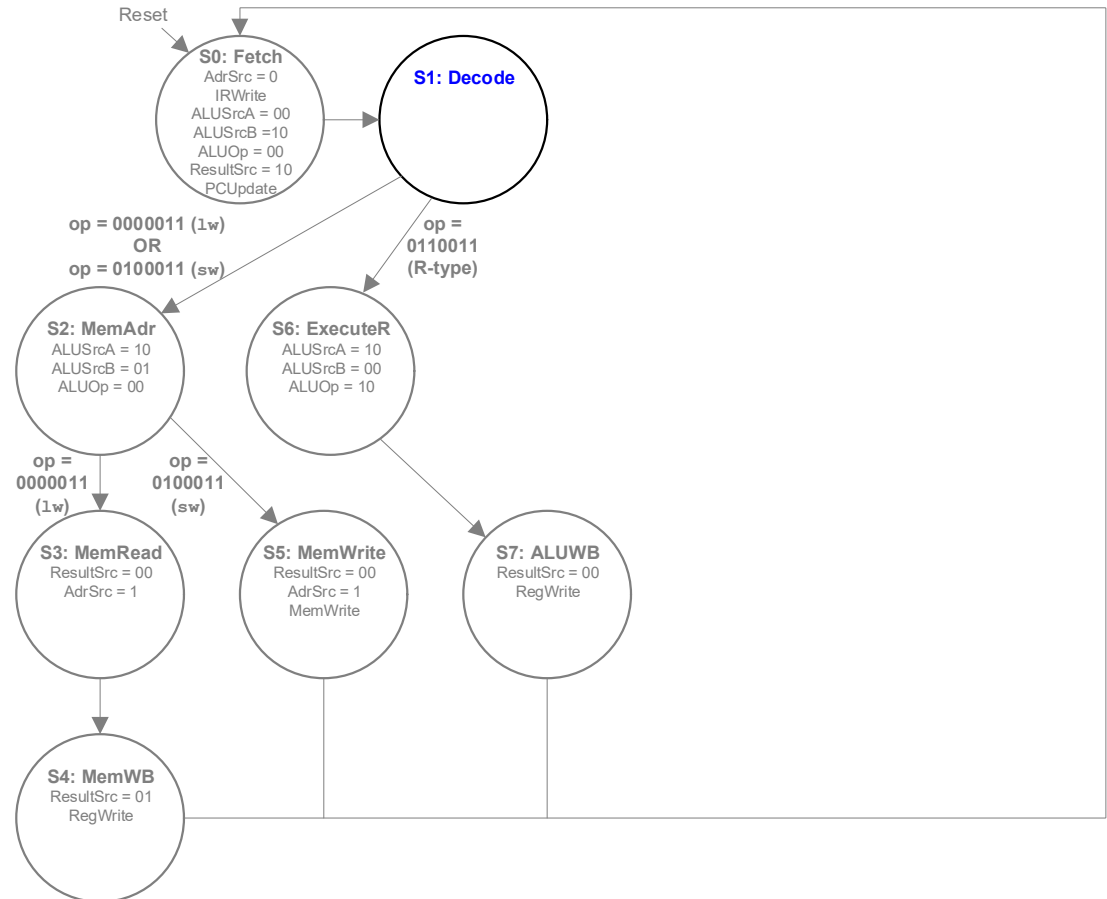


# Main FSM: R-type ALU Write Back



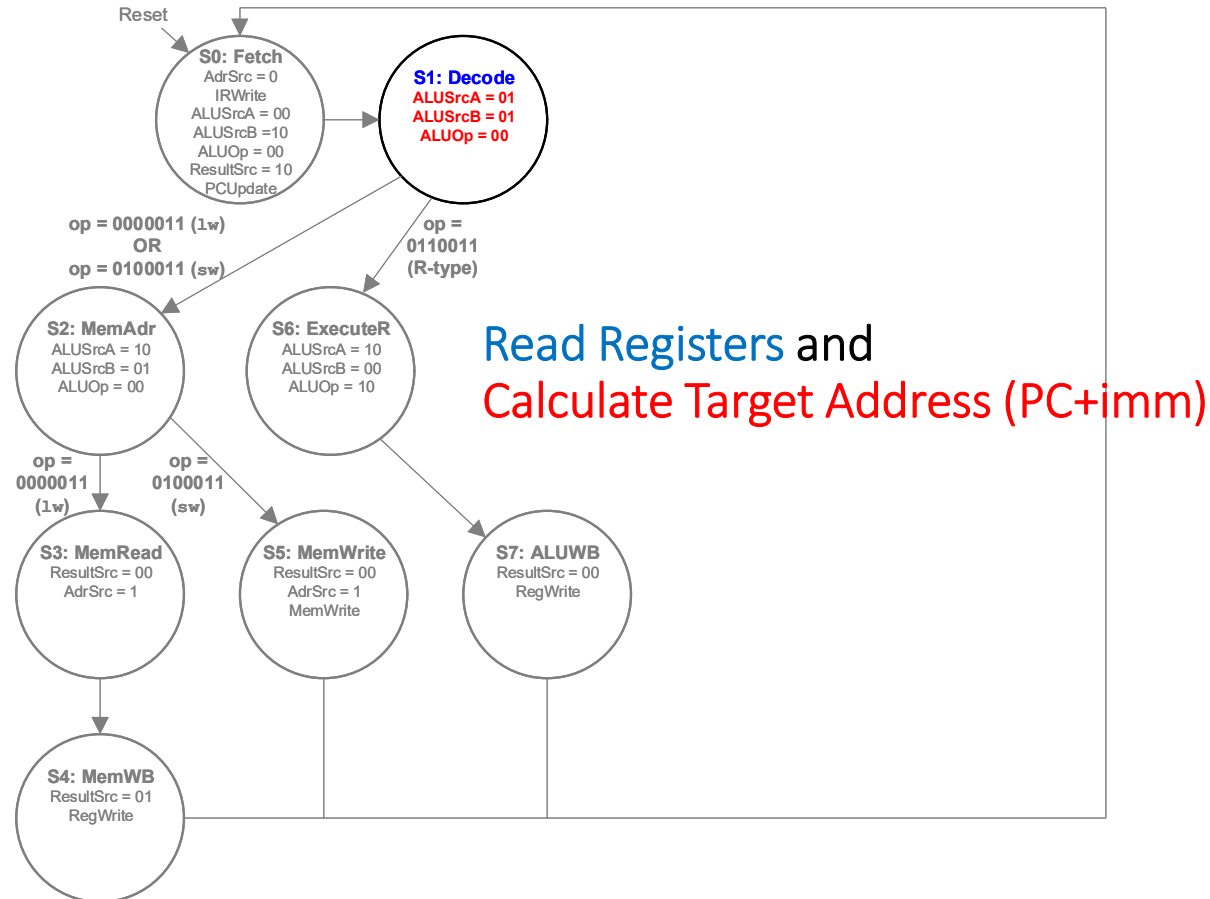
# Main FSM: R-Type ALU Write Back

- **Need to calculate:**
  - Branch Target Address
  - **rs1 - rs2** (to see if equal)
- **ALU** isn't being used in Decode stage
  - Use it to calculate Target Address (PC + imm)



# Main FSM: Decode Revisited

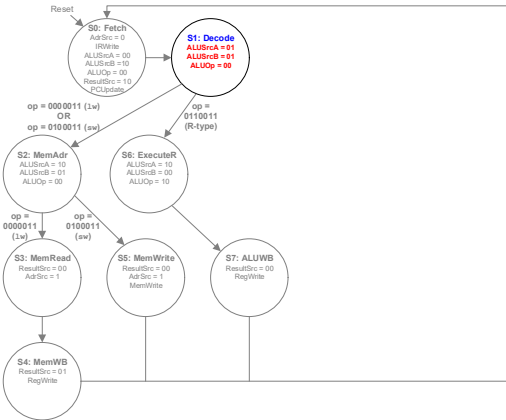
- **Need to calculate:**
  - Branch Target Address
  - **rs1 - rs2** (to see if equal)
- **ALU** isn't being used in Decode stage
  - Use it to calculate Target Address (PC + imm)



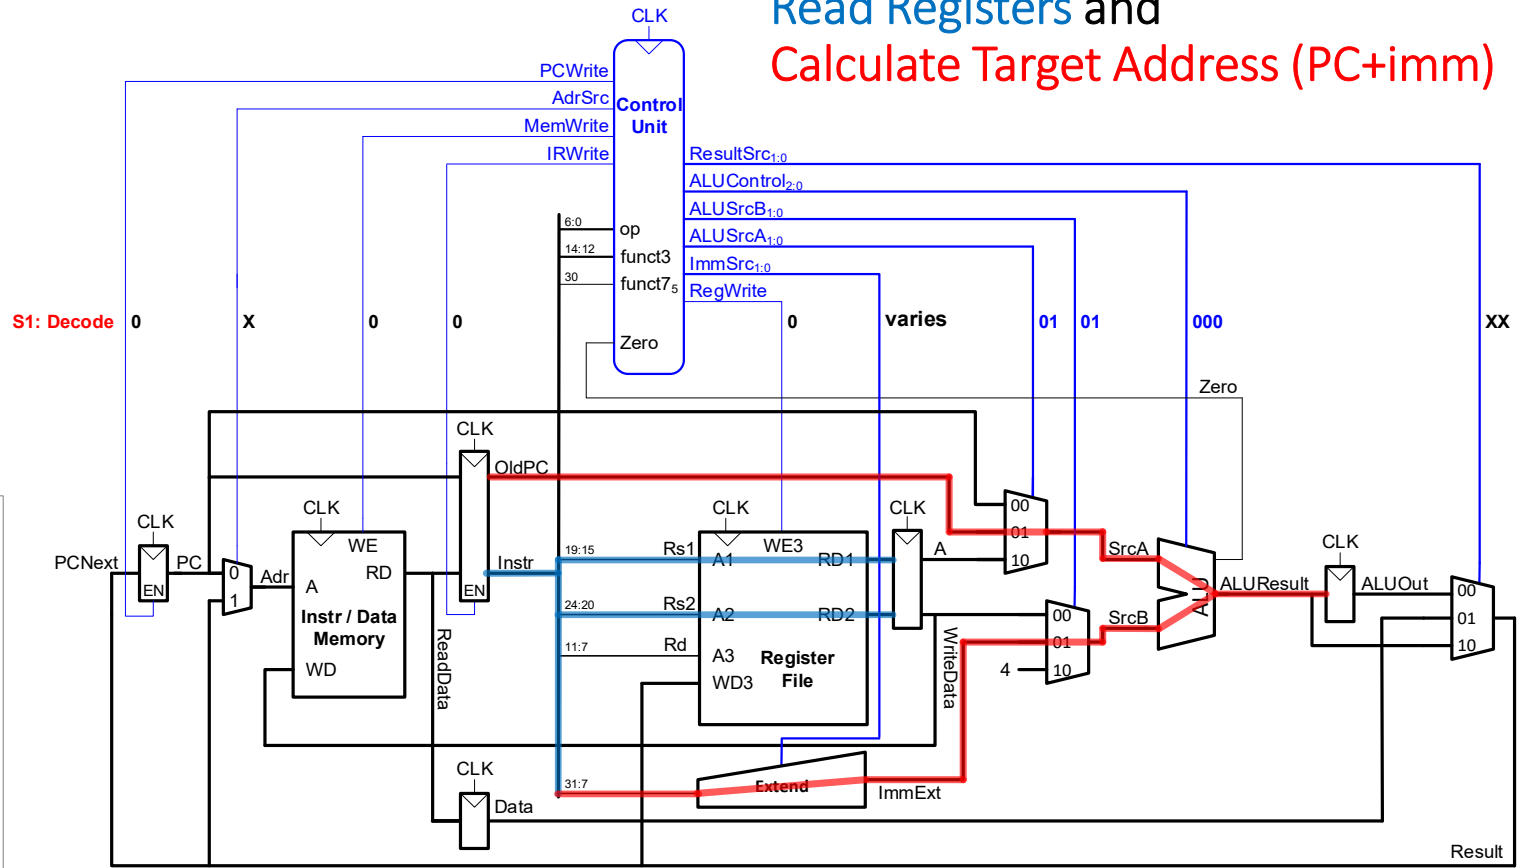


# Main FSM: Decode (Target Address)

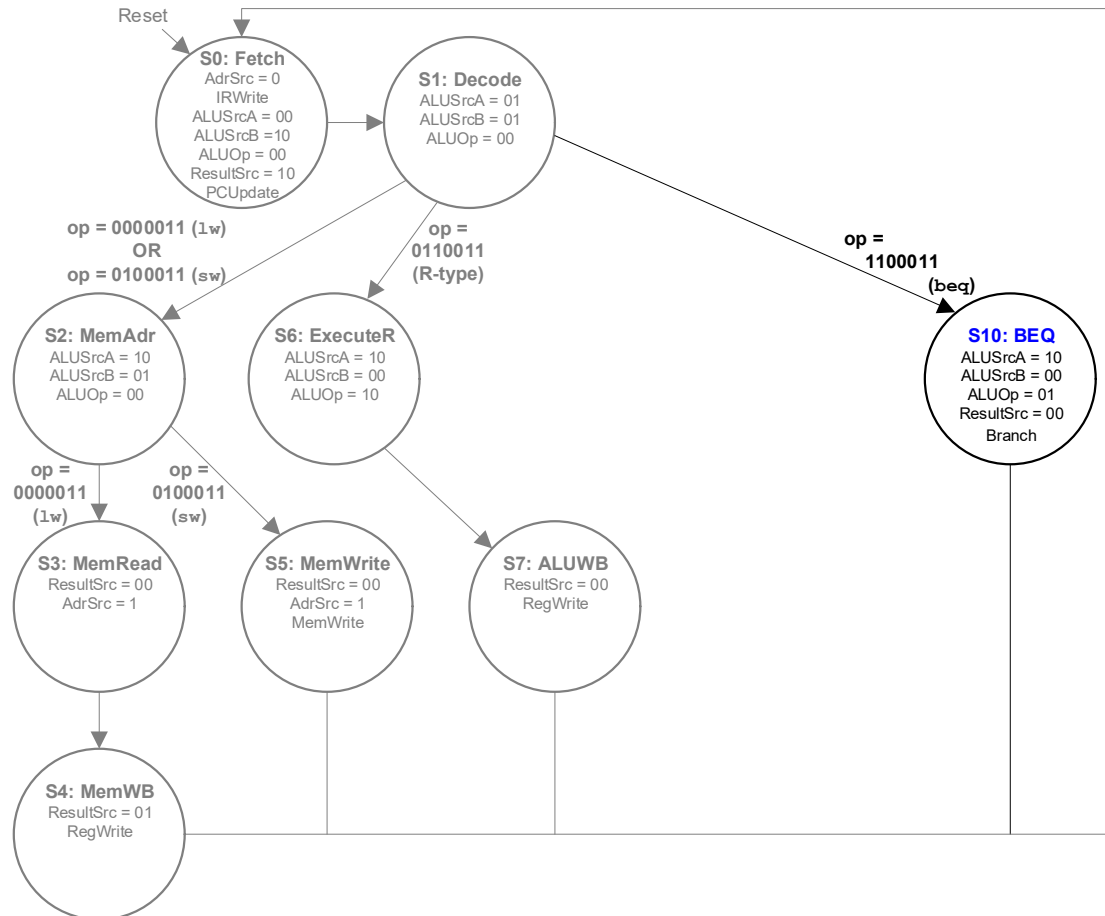
**S1: Decode**  
**ALUSrcA = 01**  
**ALUSrcB = 01**  
**ALUOp = 00**



Read Registers and  
 Calculate Target Address (PC+imm)

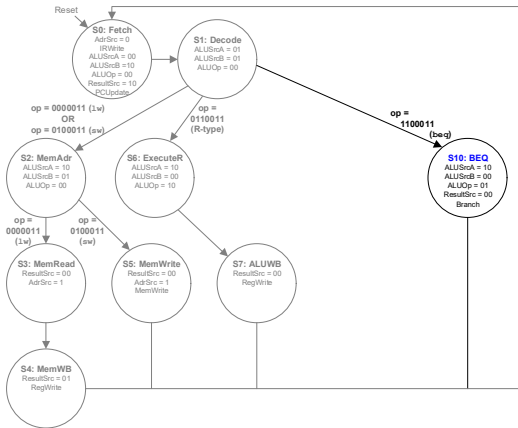


# Main FSM: beq

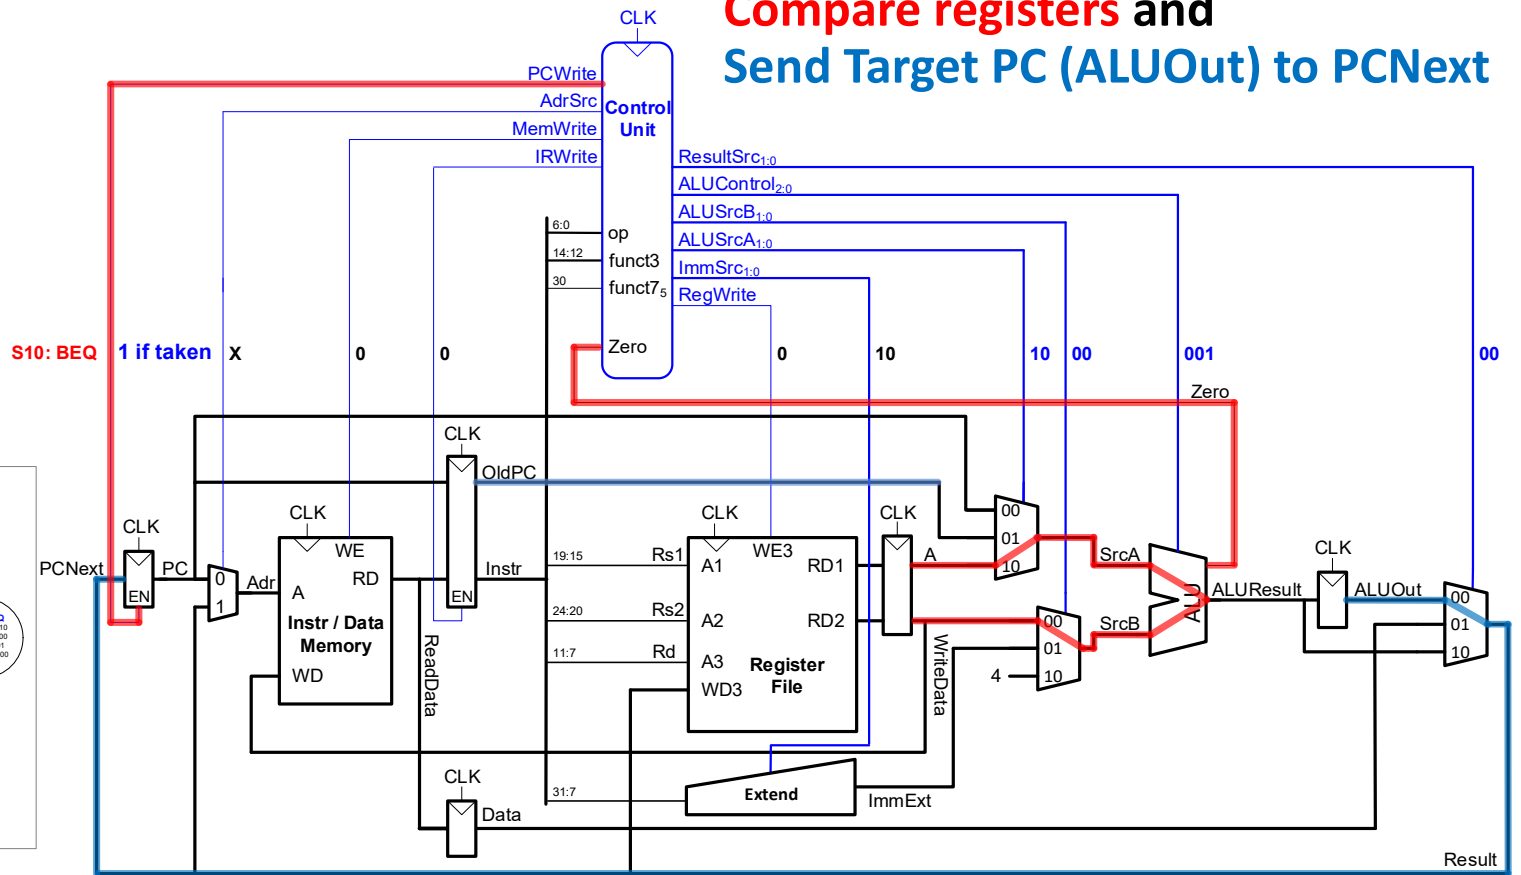


# Main FSM: beq Datapath

**S10: BEQ**  
 ALUSrcA = 10  
 ALUSrcB = 00  
 ALUOp = 01  
 ResultSrc = 00  
 Branch



**Compare registers and  
Send Target PC (ALUOut) to PCNext**

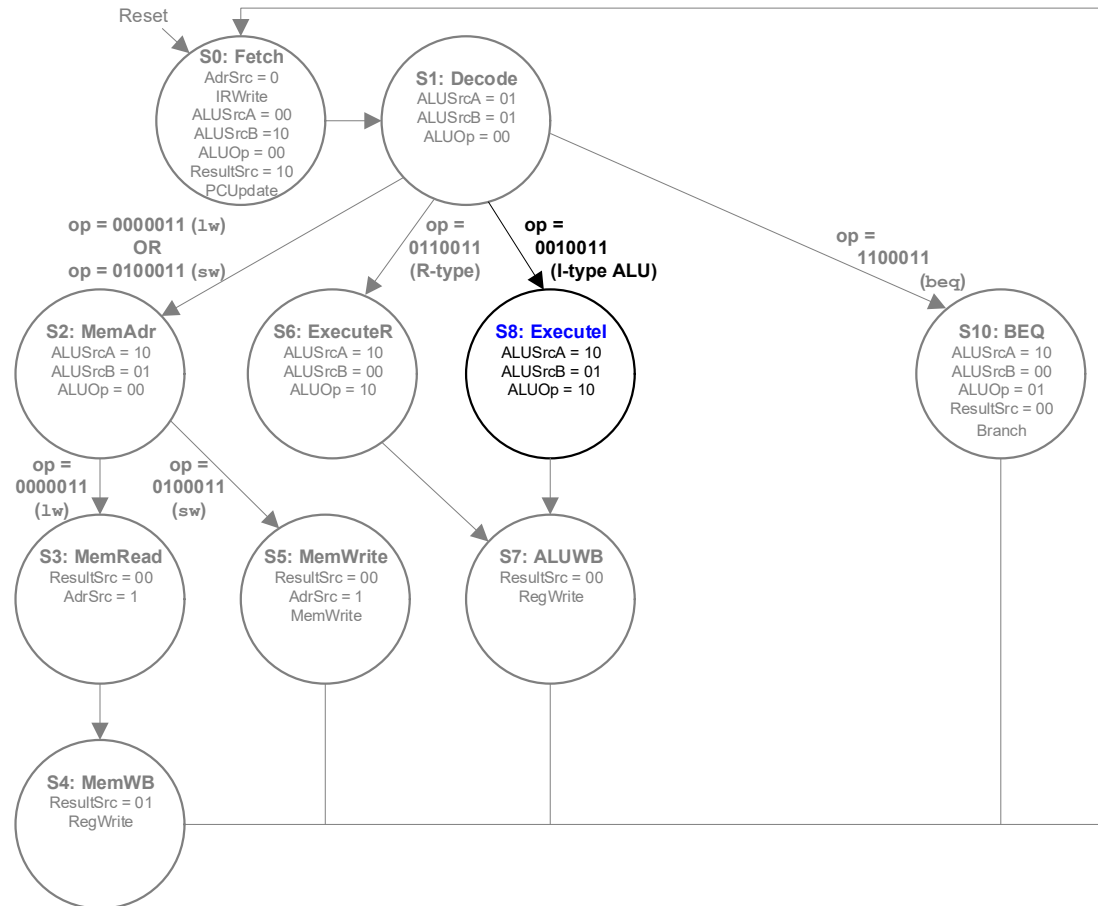


# Extending the RISC-V Multicycle Processor

---

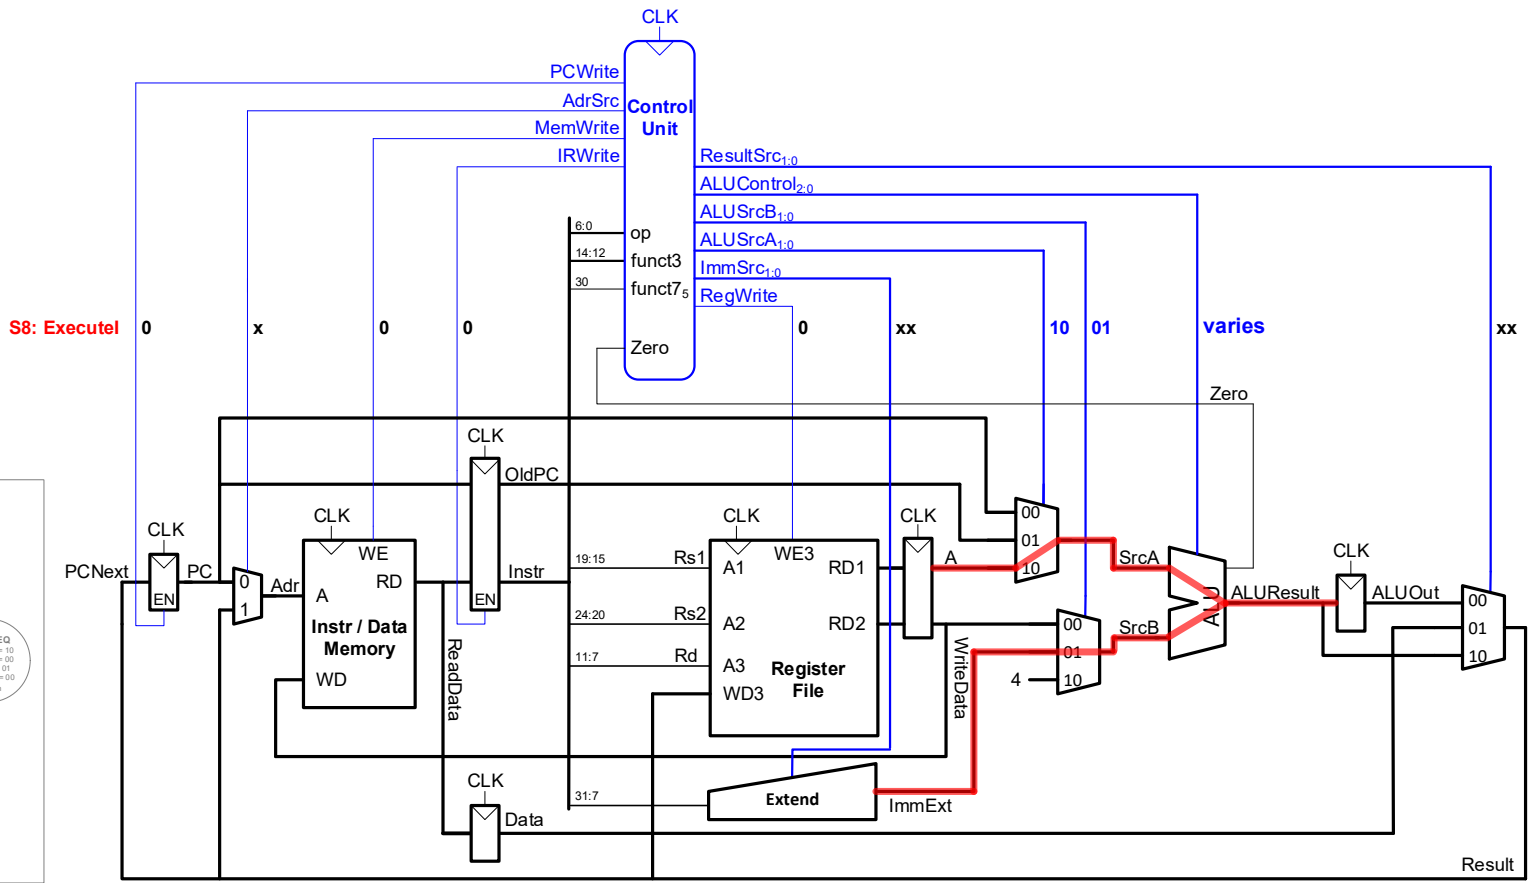
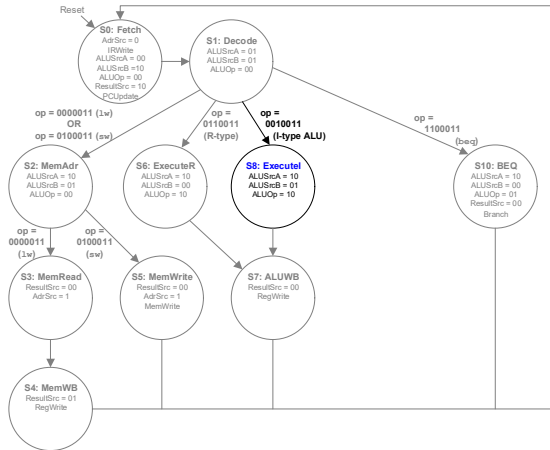
DDCA Ch7 - Part 11: Extending the RISC-V Multicycle Processor <https://www.youtube.com/watch?v=8EhVN192FRU>

# Main FSM: I-Type ALU Execute

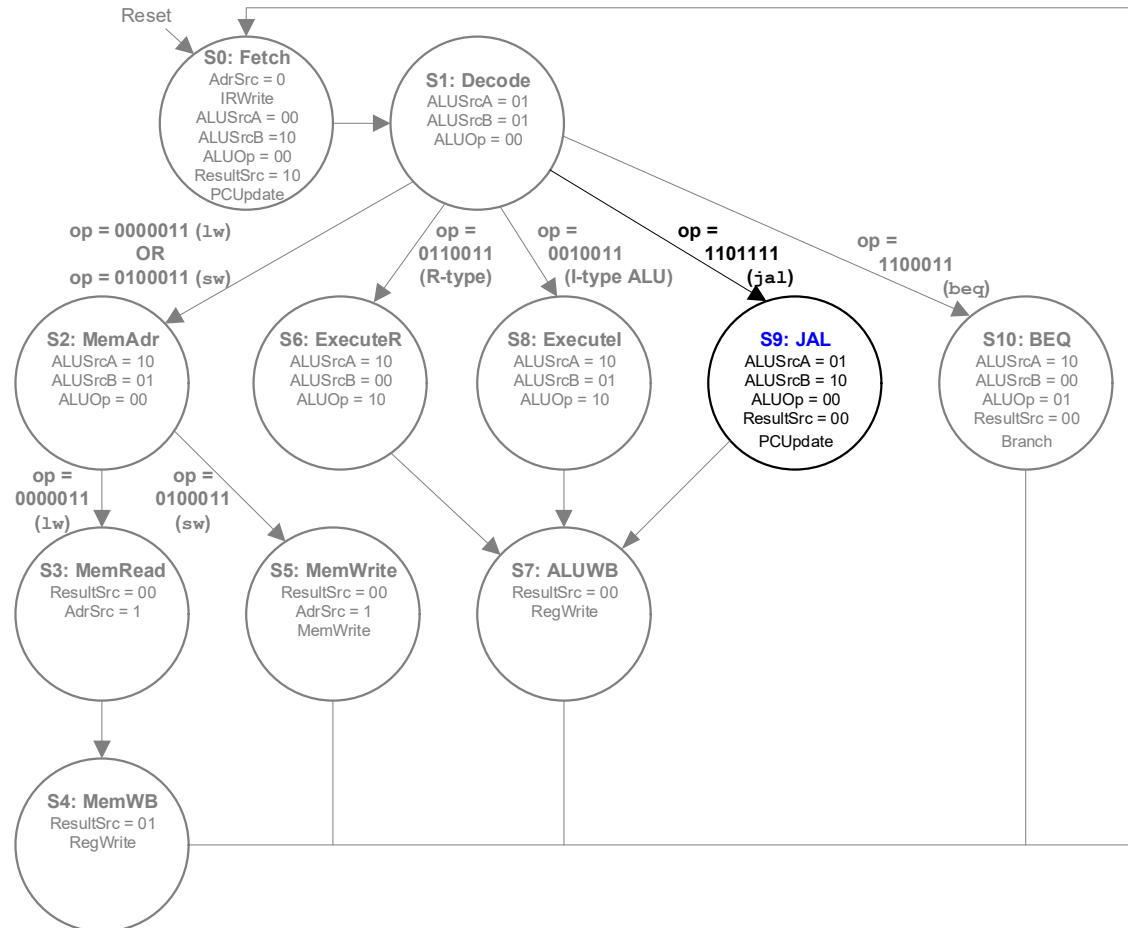


# Main FSM: I-Type ALU Exec. Datapath

**S8: Executel**  
 ALUSrcA = 10  
 ALUSrcB = 01  
 ALUOp = 10



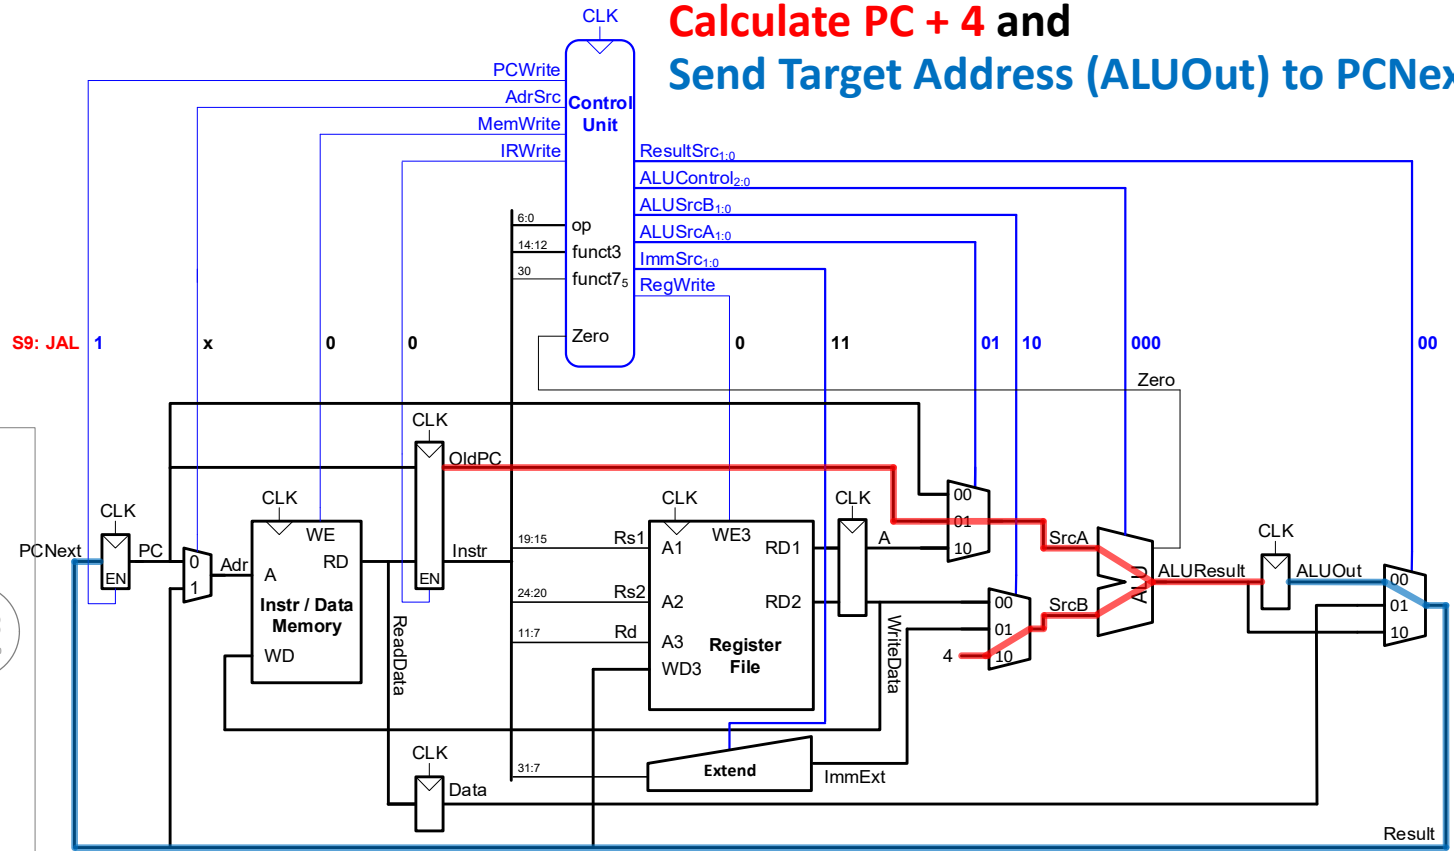
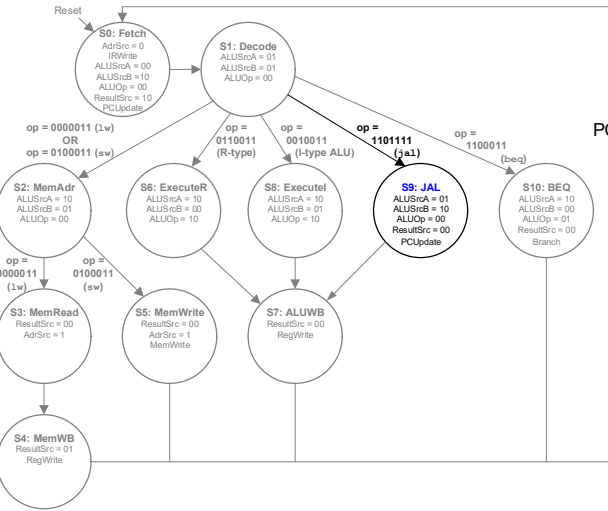
# Main FSM: jal



# Main FSM: jal Datapath

**S9: JAL**  
 ALUSrcA = 01  
 ALUSrcB = 10  
 ALUOp = 00  
 ResultSrc = 00  
 PCUpdate

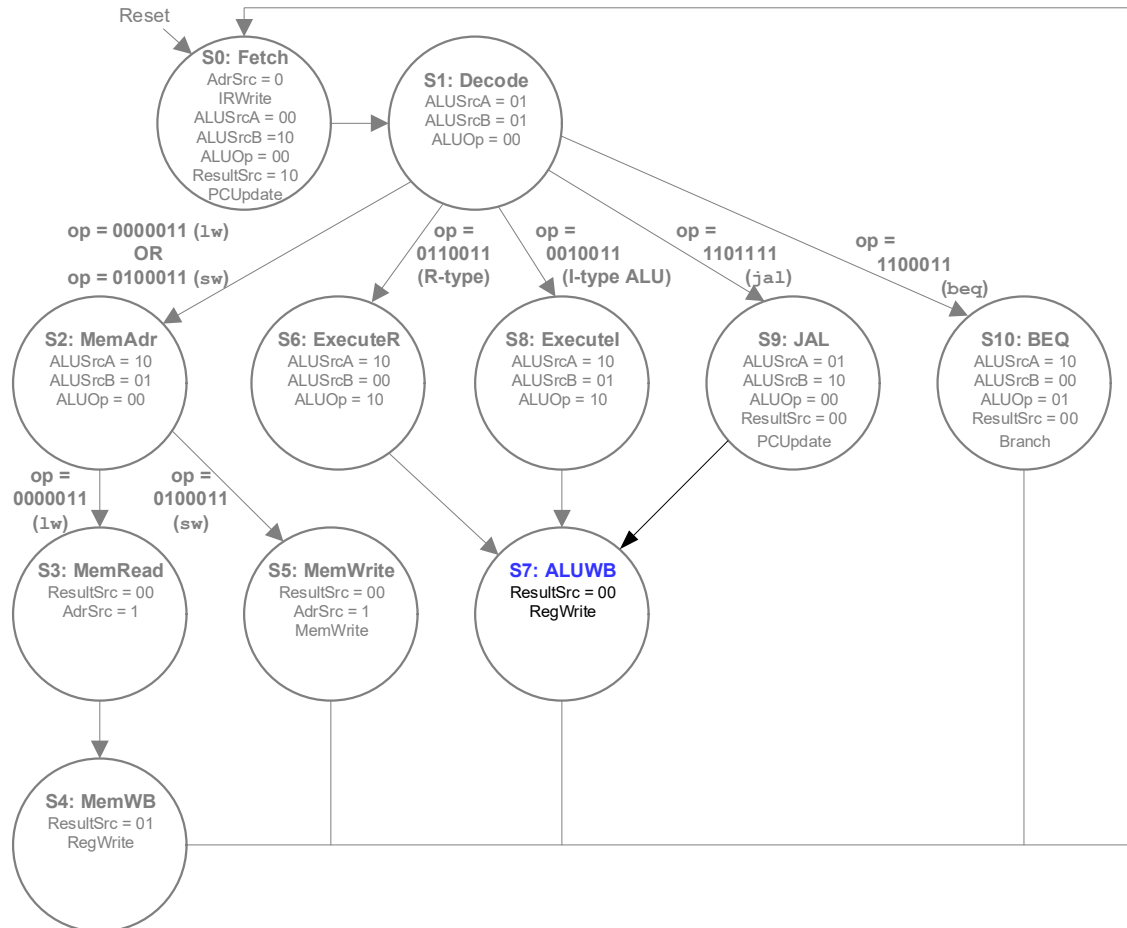
**Calculate PC + 4 and  
 Send Target Address (ALUOut) to PCNext**





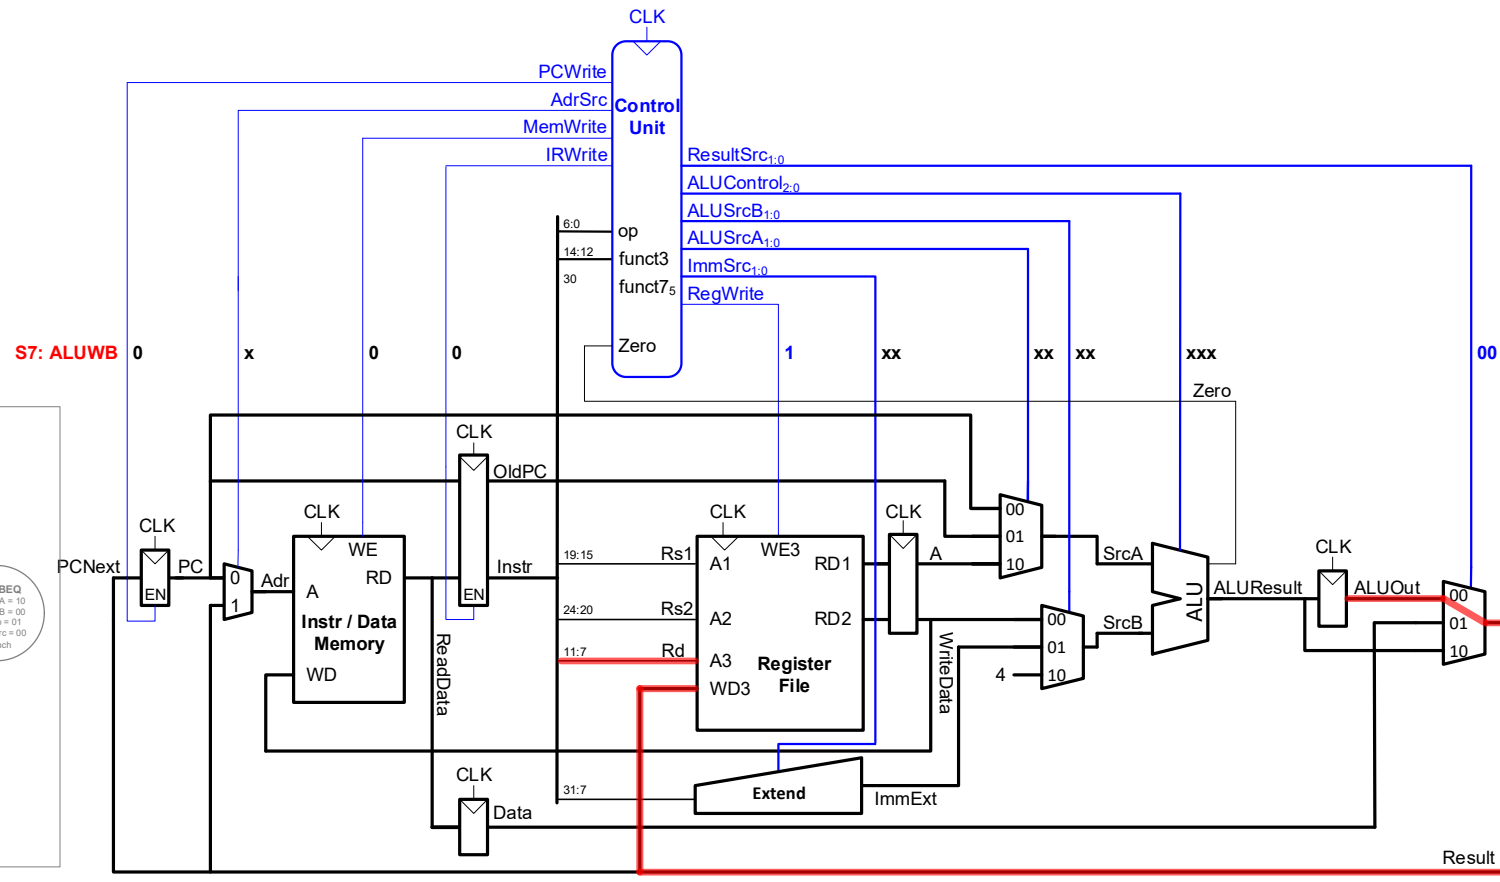
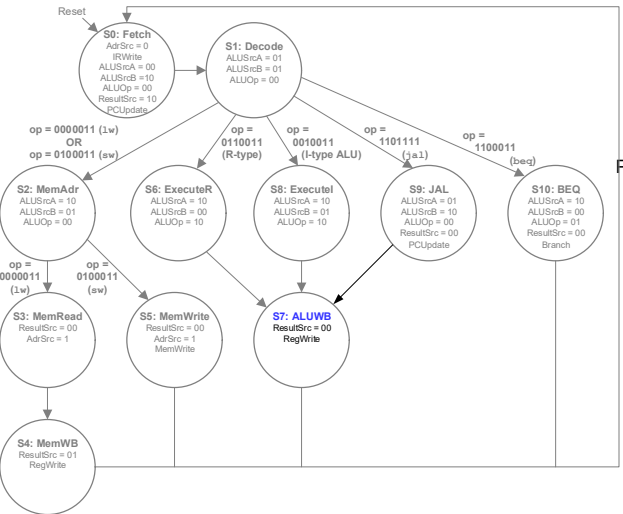
# Main FSM: jal

PC + 4 is written to rd in S7: ALUWB



# Main FSM: jal Datapath – WB same as before with PC+4

**S7: ALUWB**  
ResultSrc = 00  
RegWrite



## State

Fetch

Decode

MemAdr

MemRead

MemWB

MemWrite

ExecuteR

ExecuteI

ALUWB

BEQ

JAL

## Datapath $\mu$ Op

Instr  $\leftarrow$  Mem[PC]; PC  $\leftarrow$  PC+4

ALUOut  $\leftarrow$  PCTarget

ALUOut  $\leftarrow$  rs1 + imm

Data  $\leftarrow$  Mem[ALUOut]

rd  $\leftarrow$  Data

Mem[ALUOut]  $\leftarrow$  rd

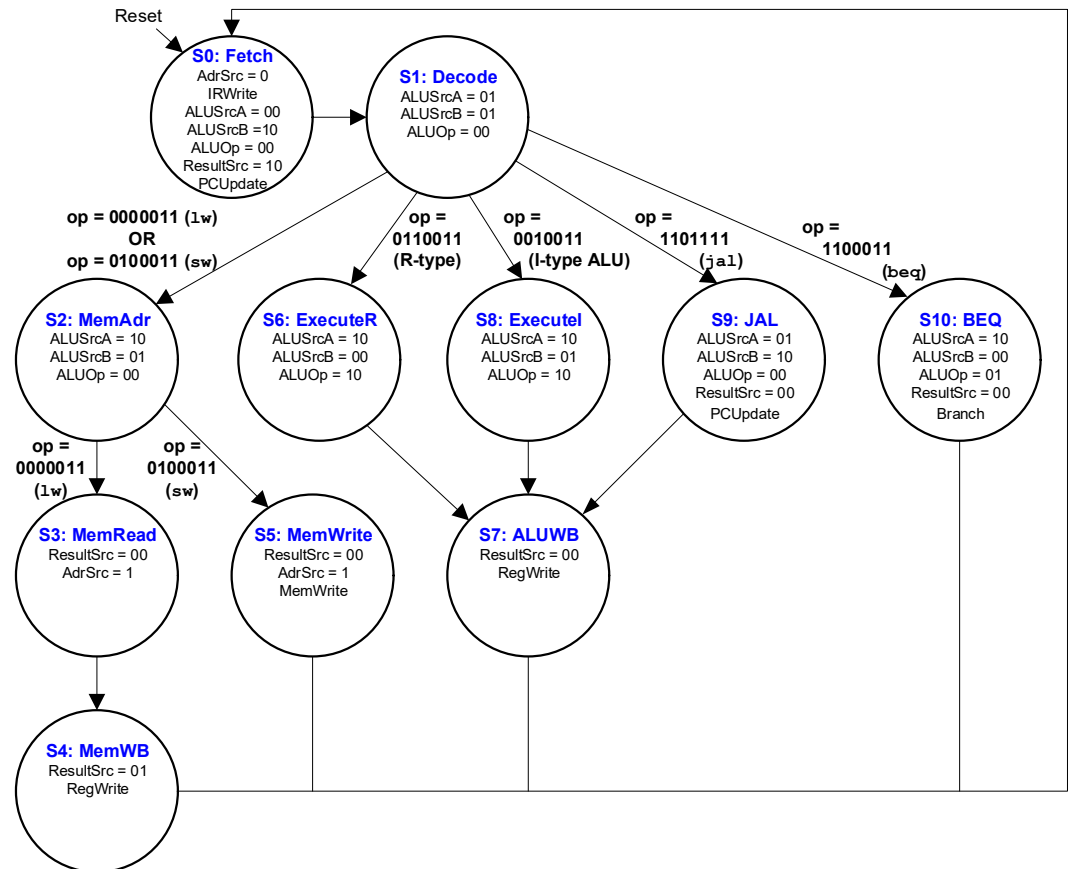
ALUOut  $\leftarrow$  rs1 op rs2

ALUOut  $\leftarrow$  rs1 op imm

rd  $\leftarrow$  ALUOut

ALUResult = rs1-rs2; if Zero, PC  $\leftarrow$  ALUOut

PC  $\leftarrow$  ALUOut; ALUOut  $\leftarrow$  PC+4



# Multicycle Performance

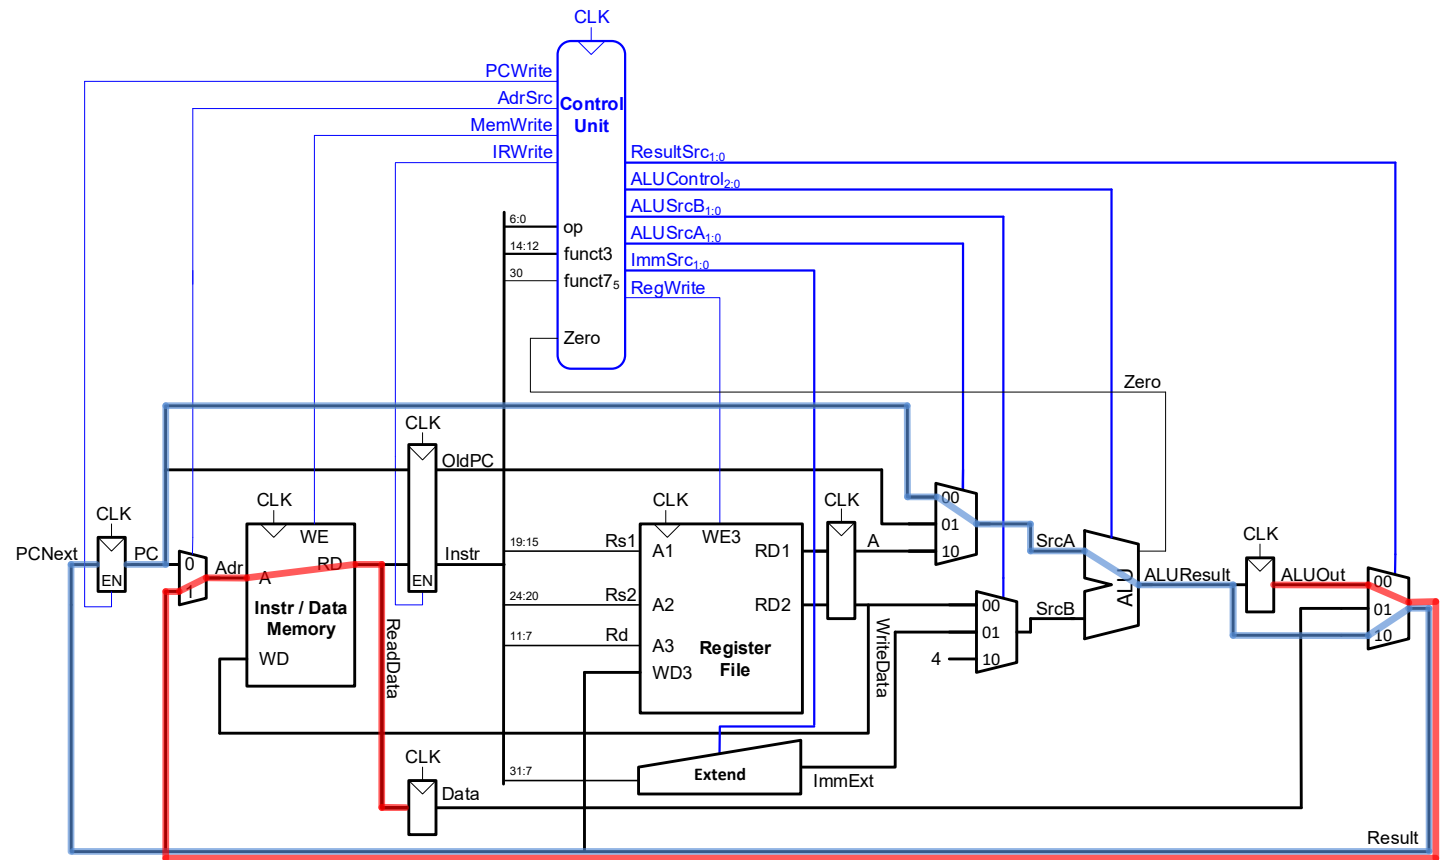
---

- Instructions take different number of cycles:
  - 3 cycles: `beq`
  - 4 cycles: R-type, `addi`, `sw`, `jal`
  - 5 cycles: `lw`
- CPI is weighted average
- SPECINT2000 benchmark:
  - 25% loads
  - 10% stores
  - 13% branches
  - 52% R-type
- **Average CPI =  $(0.13)(3) + (0.52 + 0.10)(4) + (0.25)(5) = 4.12$**

# Multicycle Critical Path

Potential Critical Paths:

- Calculate PC + 4 or
- Read Memory



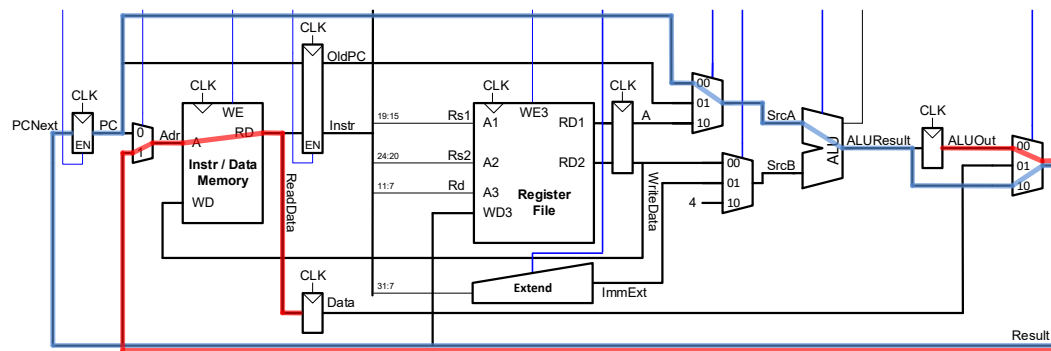
# Multicycle Processor Performance

Multicycle critical path:

- **Assumptions:**

- Registry File is faster than memory
- Writing memory is faster than reading memory

$$T_{c\_multi} = t_{pcq\_PC} + t_{dec} + 2t_{mux} + \max(t_{ALU}, t_{mem}) + t_{setup}$$



## Multicycle Performance Example

Element	Parameter	Delay (ps)
Register clock-to-Q	$t_{pcq\_PC}$	40
Register setup	$t_{setup}$	50
Multiplexer	$t_{mux}$	30
AND-OR gate	$t_{AND-OR}$	20
ALU	$t_{ALU}$	120
Decoder (Control Unit)	$t_{dec}$	25
Extend unit	$t_{dec}$	35
Memory read	$t_{mem}$	200
Register file read	$t_{RFread}$	100
Register file setup	$t_{RFsetup}$	60

$$\begin{aligned} T_{c\_multi} &= t_{pcq\_PC} + t_{dec} + 2t_{mux} + \max(t_{ALU}, t_{mem}) + t_{setup} \\ &= (40 + 25 + 2 * 30 + 200 + 50) ps = \mathbf{375 ps} \end{aligned}$$



## Multicycle Performance Example

- For a program with **100 billion** instructions executing on a **multicycle** RISC-V processor
  - **CPI** = 4.12 cycles/instruction
  - **Clock cycle time:**  $T_{c\_multi} = 375 \text{ ps}$
- **Execution Time** = (# instructions)  $\times$  CPI  $\times T_c$ 
  - =  $(100 \times 10^9)(4.12)(375 \times 10^{-12})$
  - = **155 seconds**
- This is **slower** than the single-cycle processor (75 sec.)

# Parallelism

---

DDCA Ch3 - Part 16: Parallelism <https://www.youtube.com/watch?v=xX2Crru3xCg>

# Parallelism

Two types of parallelism:

- **Spatial parallelism**
  - duplicate hardware performs multiple tasks at once
- **Temporal parallelism**
  - task is broken into multiple stages
  - also called pipelining
  - for example, an assembly line

## Parallelism

- **Token:** Group of inputs processed to produce group of outputs
- **Latency:** Time for one token to pass from start to end
- **Throughput:** Number of tokens produced per unit time

Parallelism increases throughput

## Parallelism Example 1/3

- Ben Bitdiddle bakes cookies to celebrate traffic light controller installation
  - **5 minutes** to roll cookies
  - **15 minutes** to bake
- What is the latency and throughput without parallelism?
  - Latency (when is the first cookie finished?)
  - Throughput (how many cookies can Ben finish in an hour?)

$$\text{Latency} = 5 + 15 = 20 \text{ minutes} = 1/3 \text{ hour}$$
$$\text{Throughput} = 1 \text{ tray} / 1/3 \text{ hour} = 3 \text{ trays/hour}$$

## Parallelism Example 2/3

What is the latency and throughput if Ben uses parallelism?

- **Spatial parallelism:**

Ben asks Allysa P. Hacker to help, using her own oven

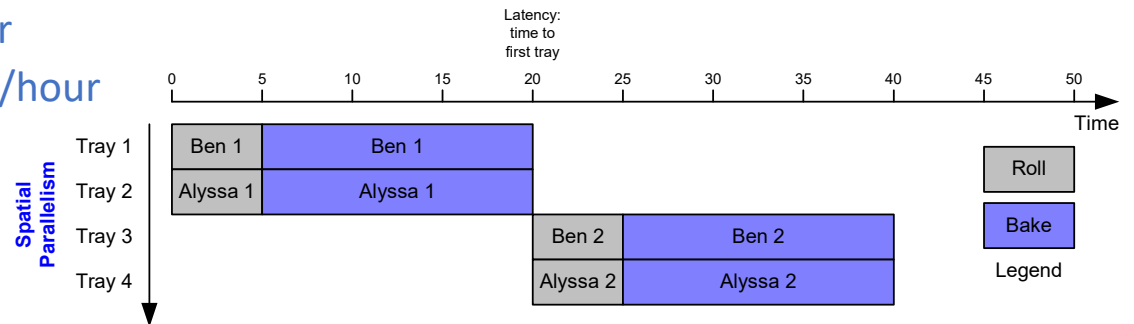
- **Temporal parallelism:**

- two stages: rolling and baking
- He uses two trays
- While first batch is baking, he rolls the second batch, etc.

# Spatial vs. Temporal Parallelism

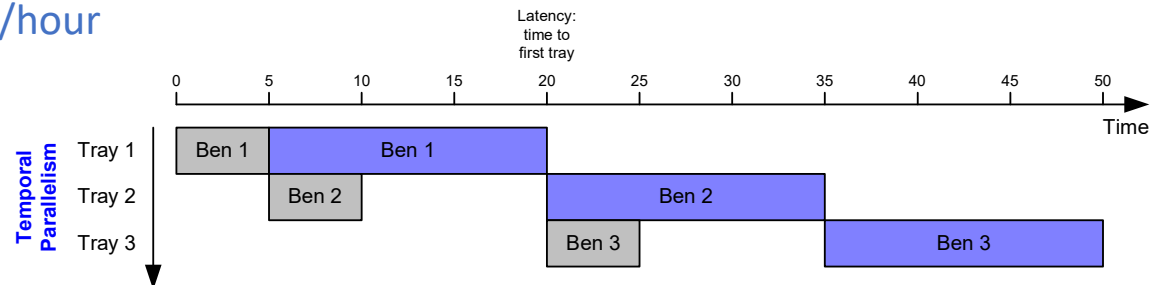
- Spatial Parallelism

- Latency = 5 + 15 = 20 minutes = 1/3 hour
- Throughput = 2 trays/ 1/3 hour = 6 trays/hour



- Temporal Parallelism

- Latency = 5 + 15 = 20 minutes = 1/3 hour
- Throughput = 1 trays/ 1/4 hour = 4 trays/hour



# Parallelism in Circuits



# Pipelined RISC-V Processor

---

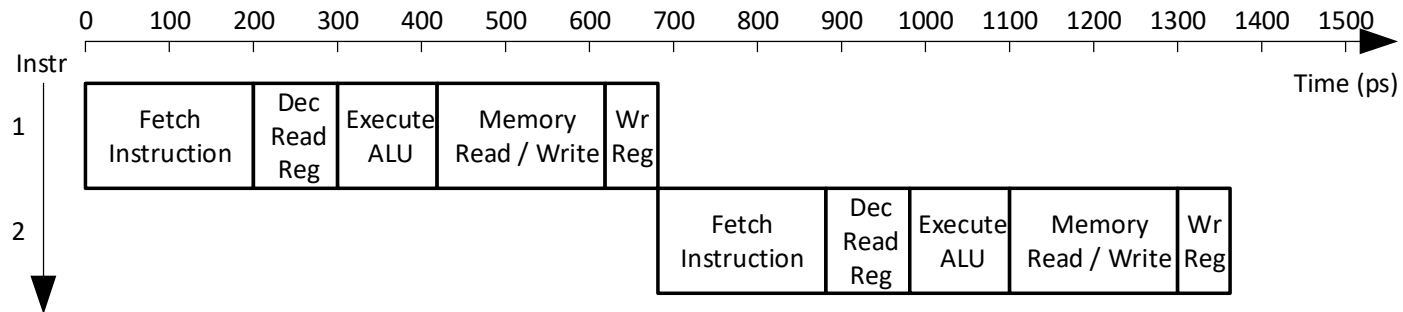
DDCA Ch7 - Part 13: Pipelined Processor <https://www.youtube.com/watch?v=UZdURUwQMmk>

# Pipelined RISC-V Processor

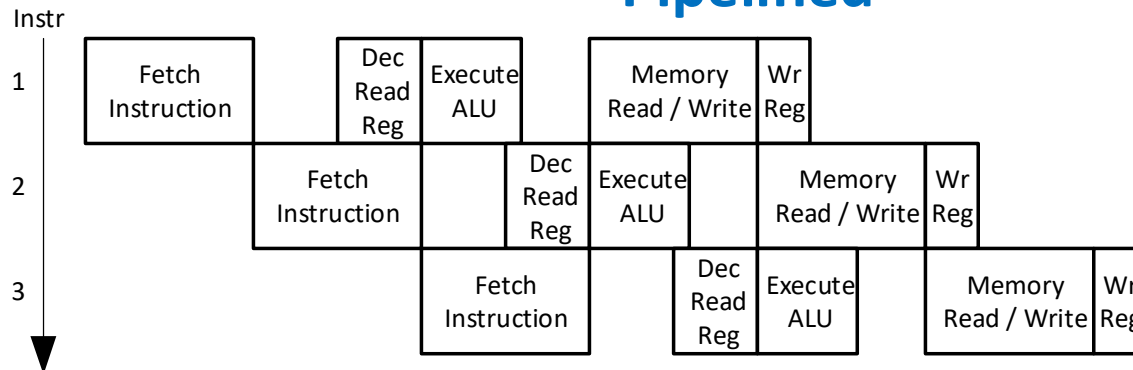
- **Temporal parallelism**
- Divide single-cycle processor into **5 stages**:
  - Fetch
  - Decode
  - Execute
  - Memory
  - Writeback
- Add **pipeline registers** between stages

# Single-Cycle vs. Pipelined Processor

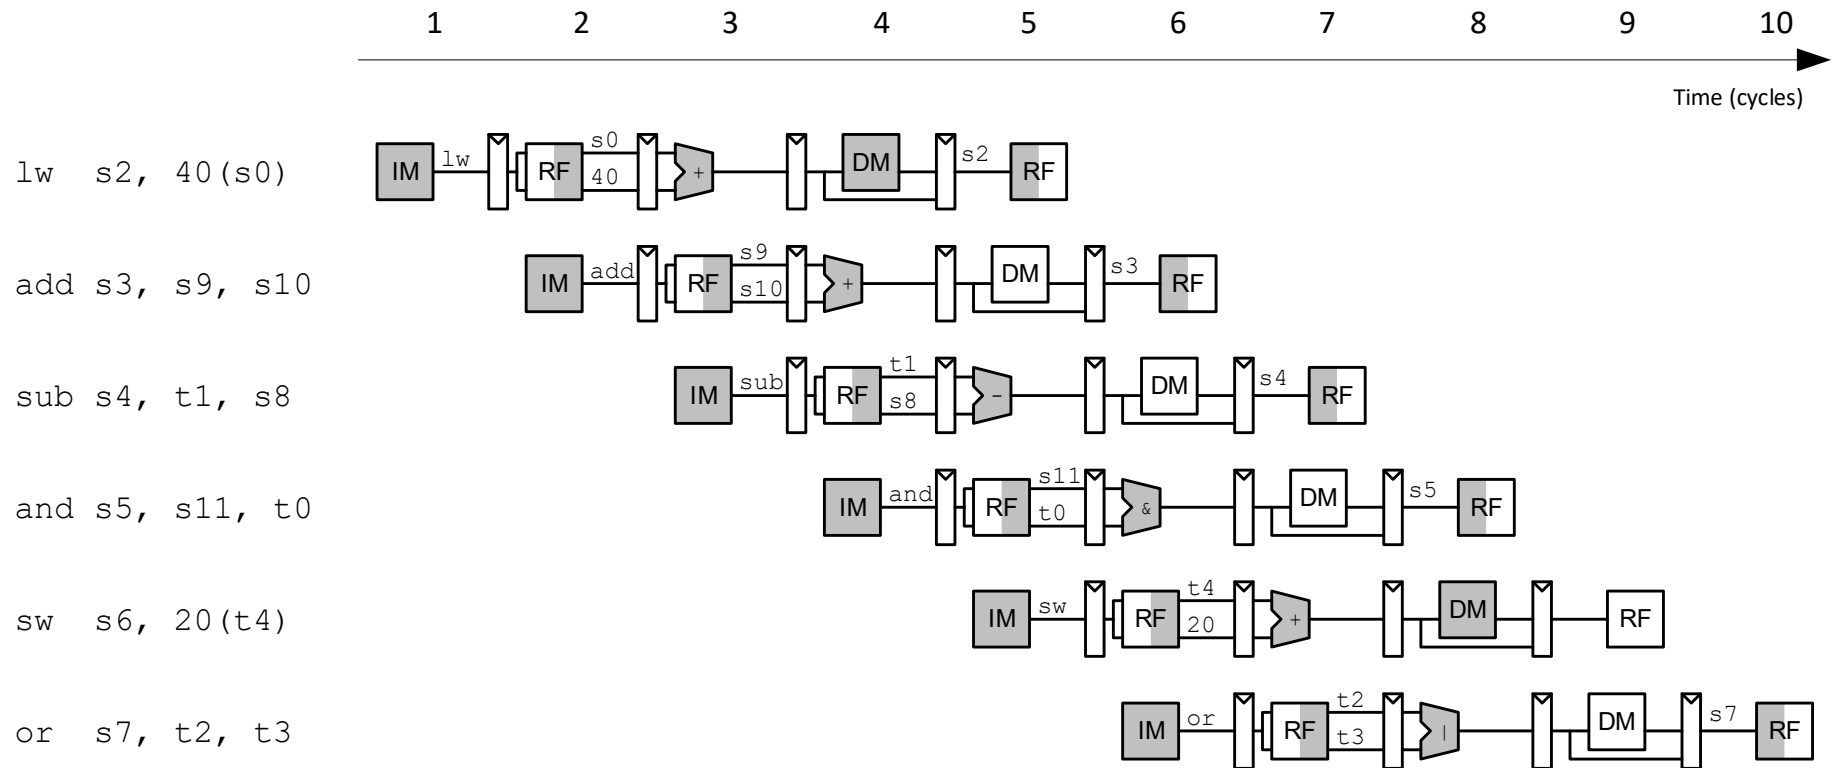
## Single-Cycle



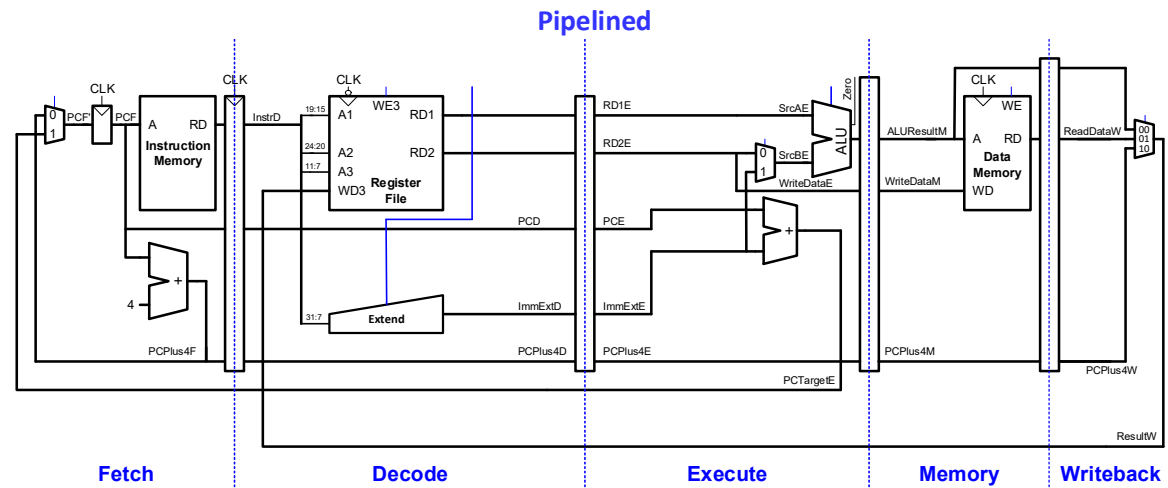
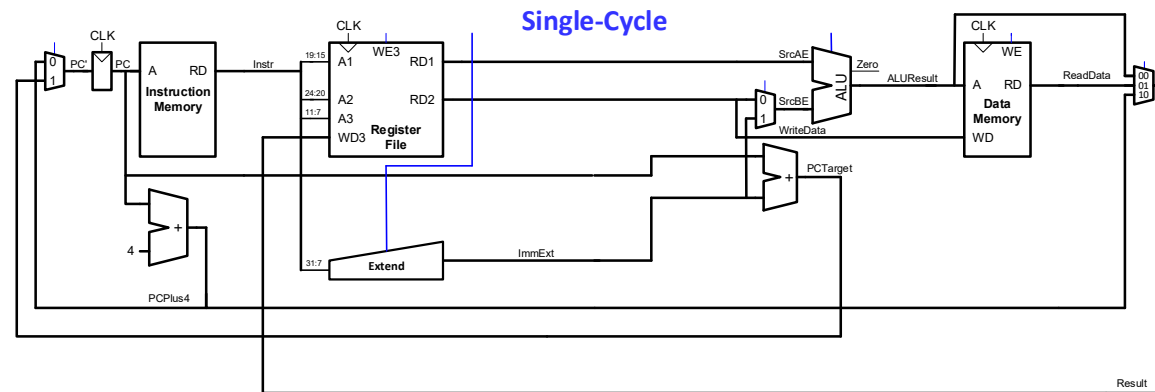
## Pipelined



# Pipelined Processor Abstraction

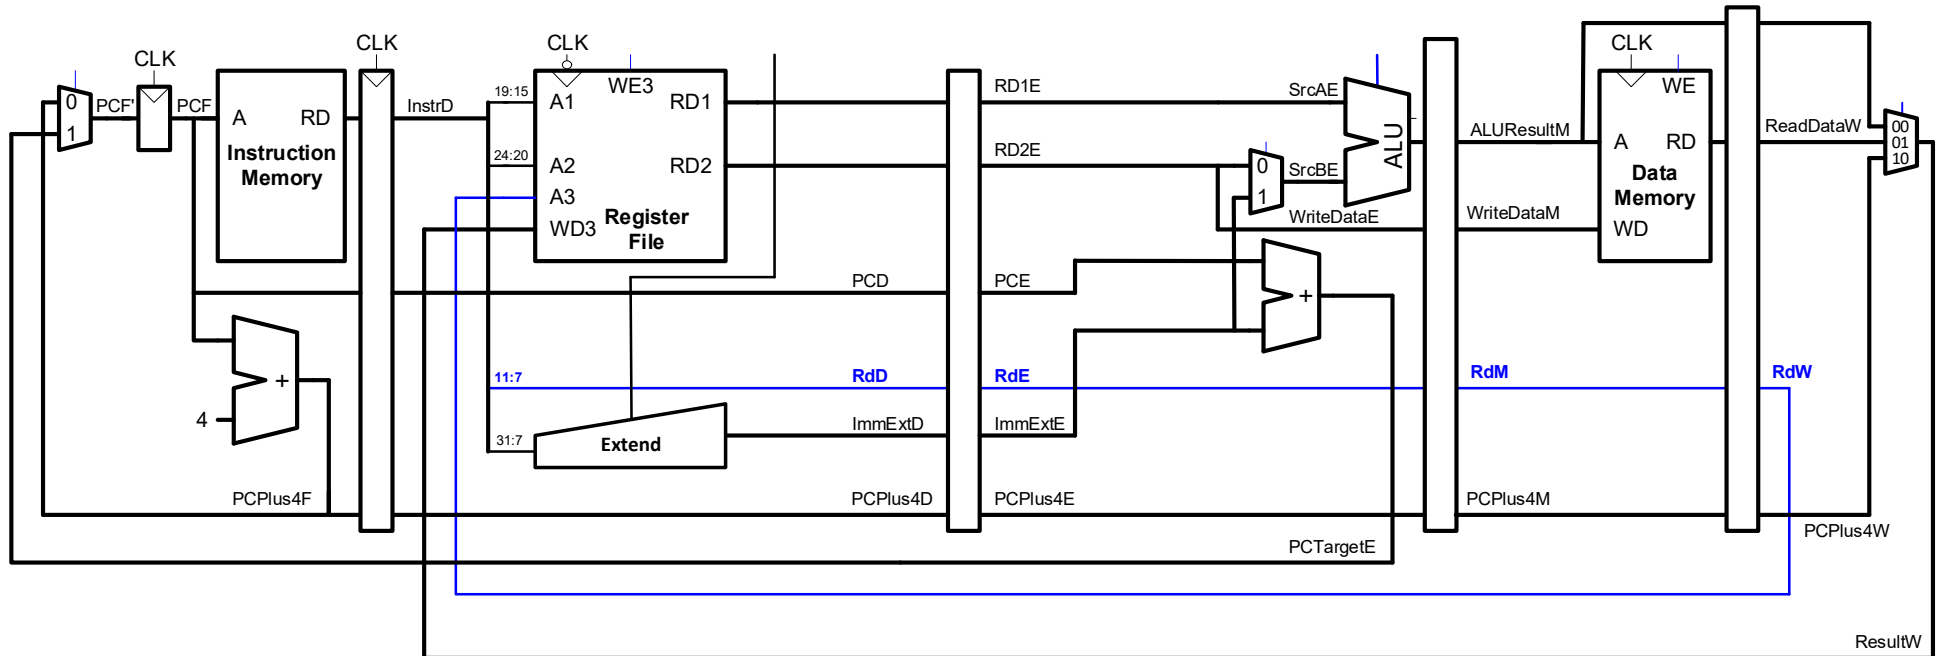


# Single-Cycle & Pipelined Datapaths



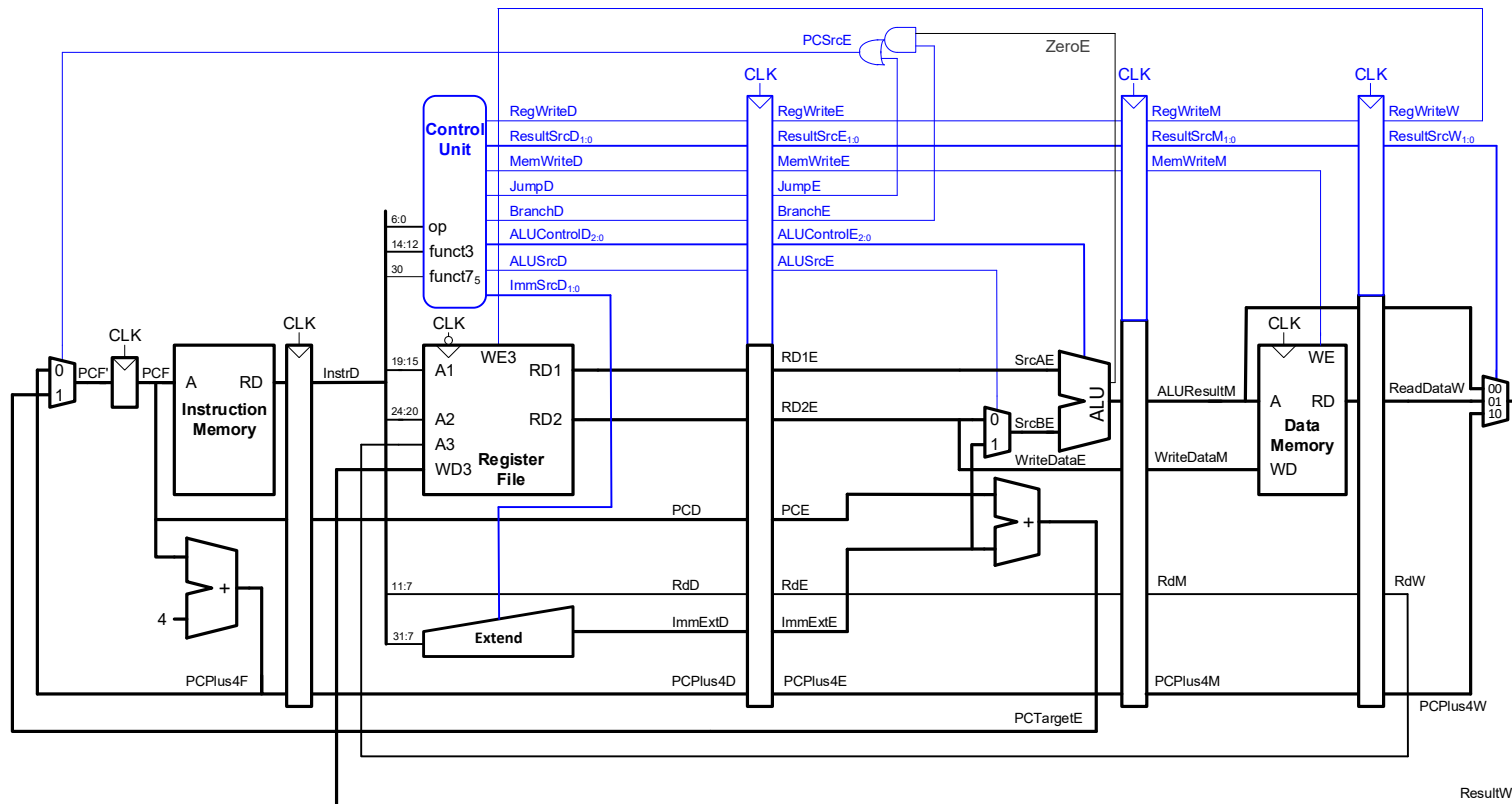
Signals in Pipelined Processor are appended with first letter of stage (i.e., PC**F**, PC**D**, PC**E**, PC**M**, PC**W**).

## Corrected Pipelined Datapath



- **Rd** must arrive at same time as **Result**
- Register file written on **falling edge** of **CLK**

# Pipelined Processor with Control



- **Same control unit** as single-cycle processor
- **Control signals travel with** the instruction (drop off when used)

# Pipelined Processor Hazards

---

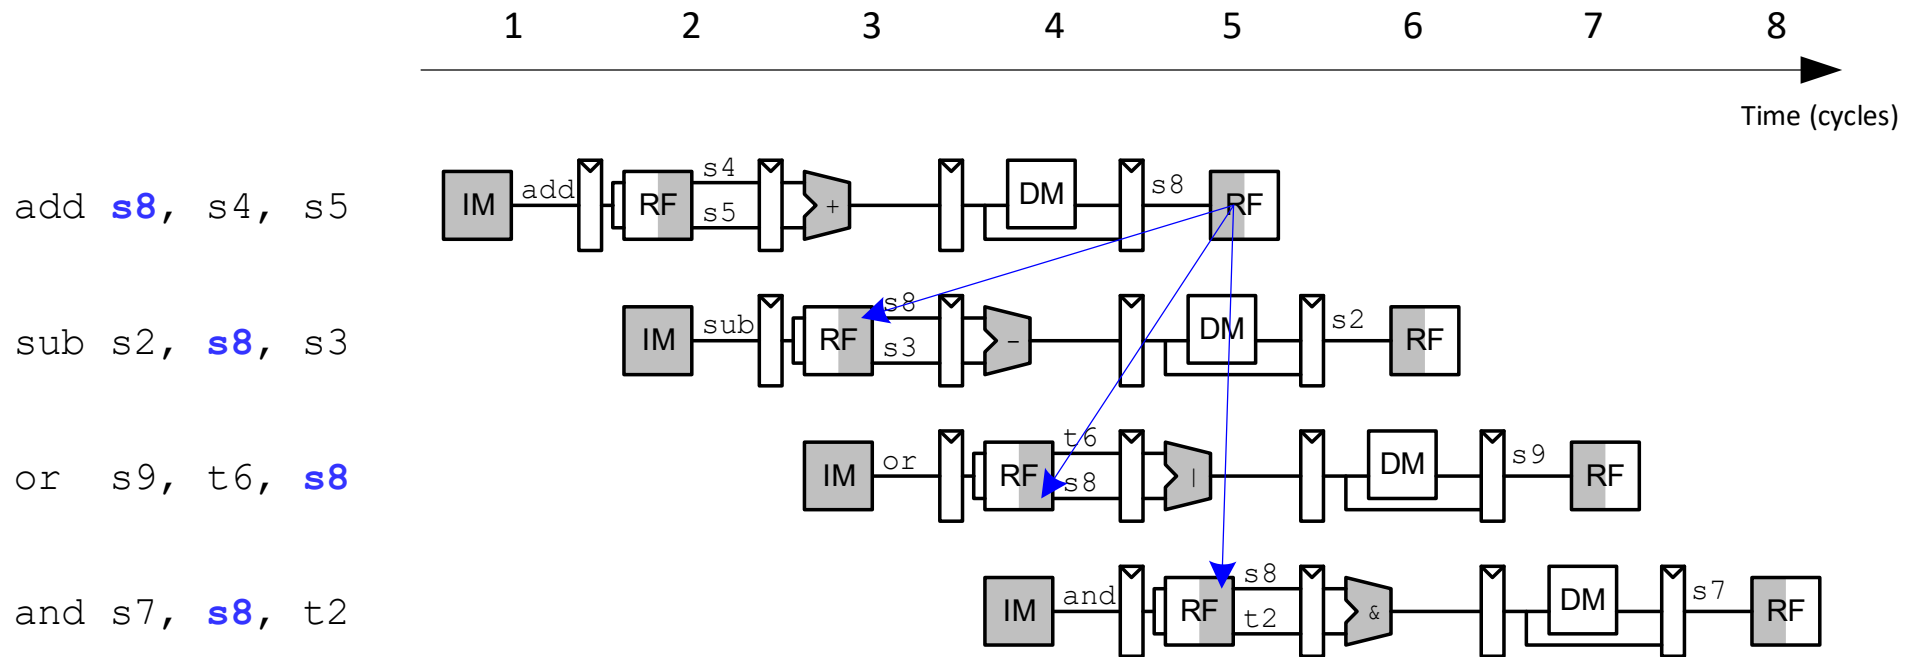
DDCA Ch7 - Part 14: Pipelined Processor Data Hazards <https://www.youtube.com/watch?v=zuegcg6ZSFQ>



## Pipelined Hazards

- When an instruction depends on result from instruction that hasn't completed
- Types:
  - **Data hazard:**  
register value not yet written back to register file
  - **Control hazard:**  
next instruction not decided yet (caused by branch)

# Data Hazard

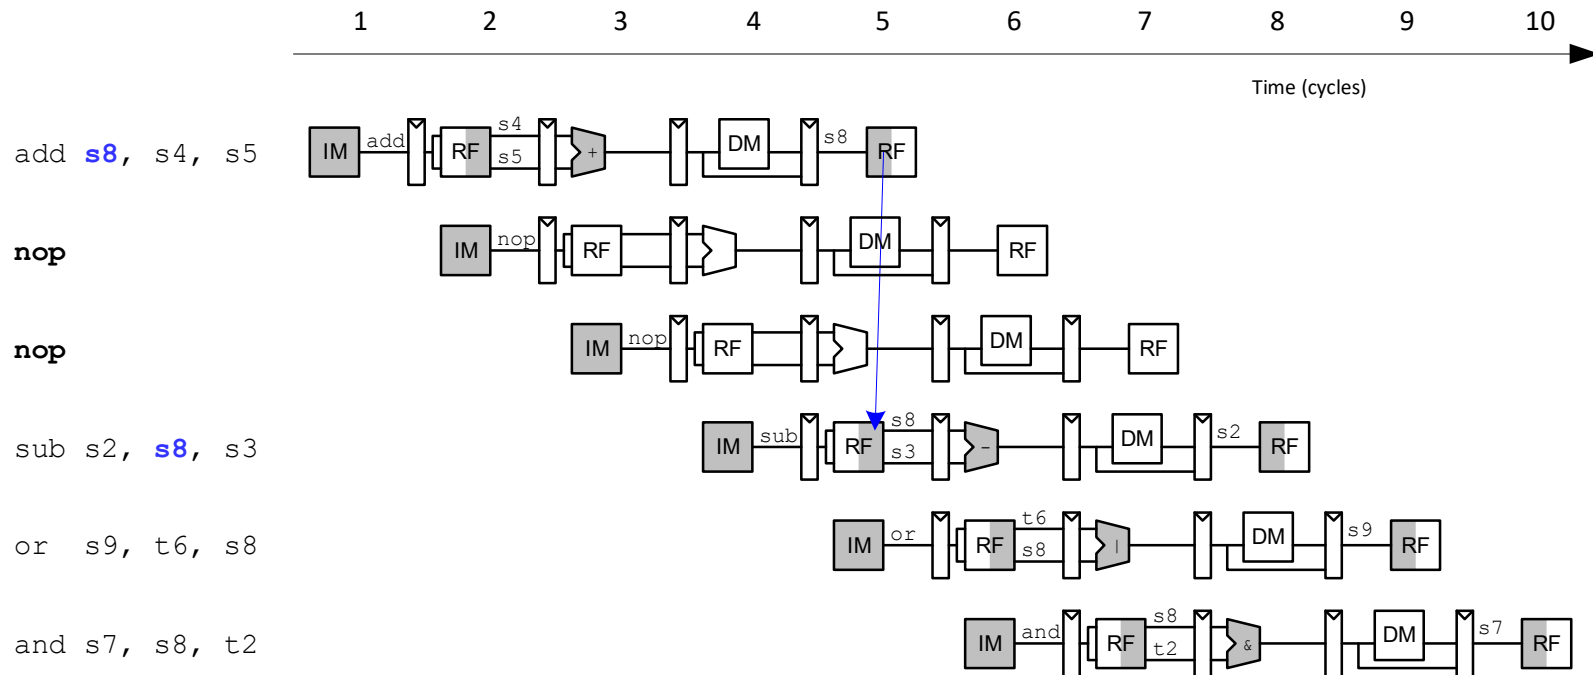


## Handling Data Hazards

- Insert **nops** in code at compile time
- **Rearrange** code at compile time
- **Forward** data at run time
- **Stall** the processor at run time

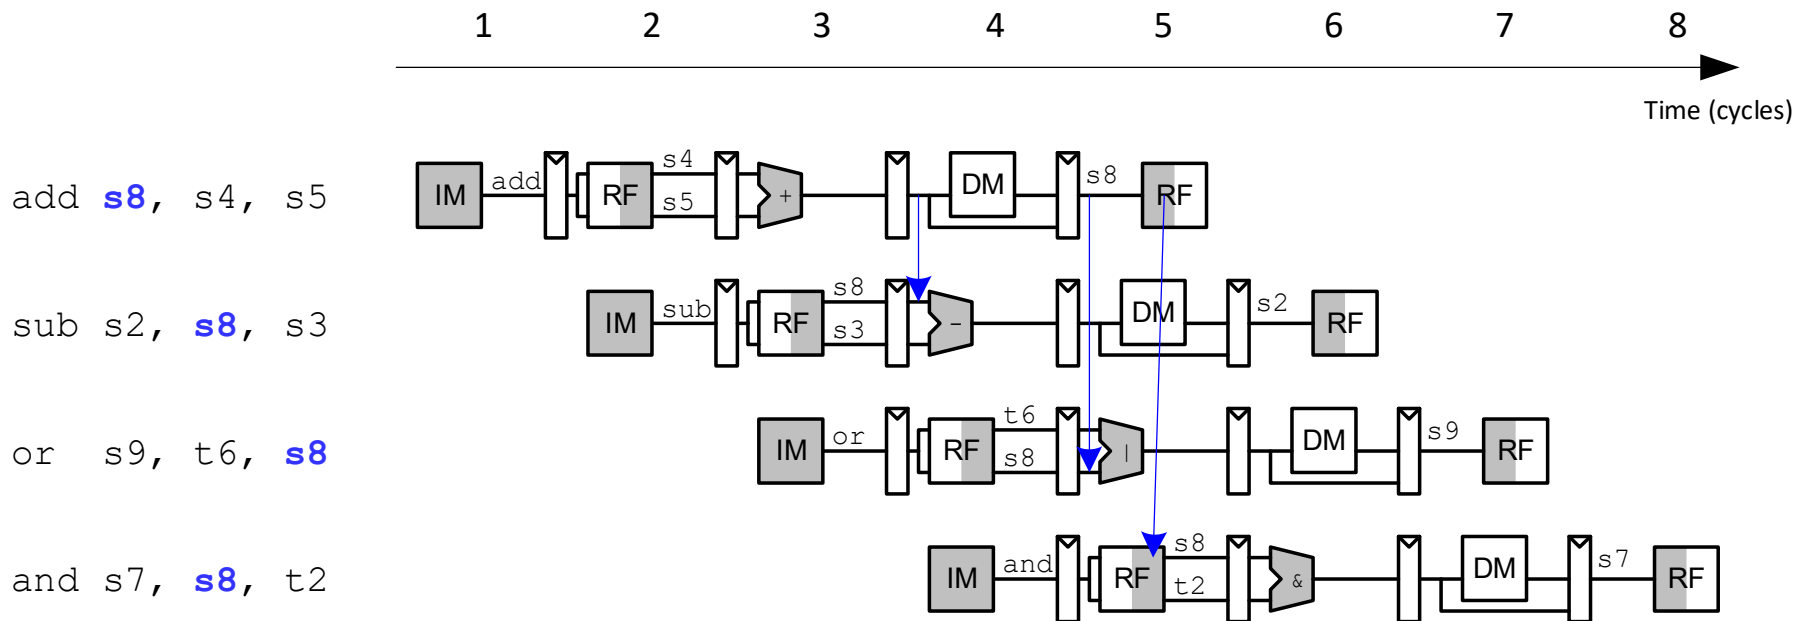
# Handling Data Hazards

- **Insert** enough **nops** for result to be ready
- Or move independent useful instructions forward



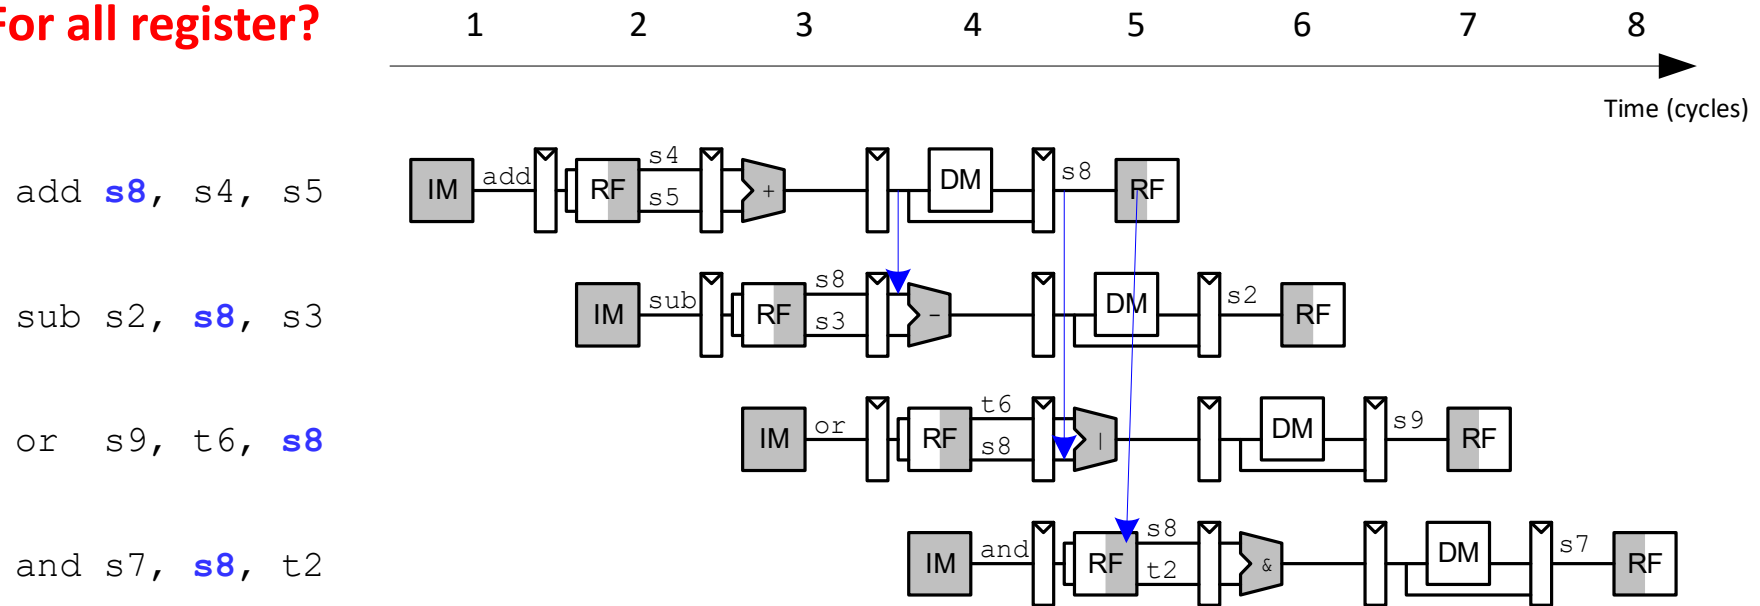
# Data Forwarding

- Data is **available on internal busses** before it is written back to the register file (RF).
- **Forward data** from internal busses **to Execute stage**.

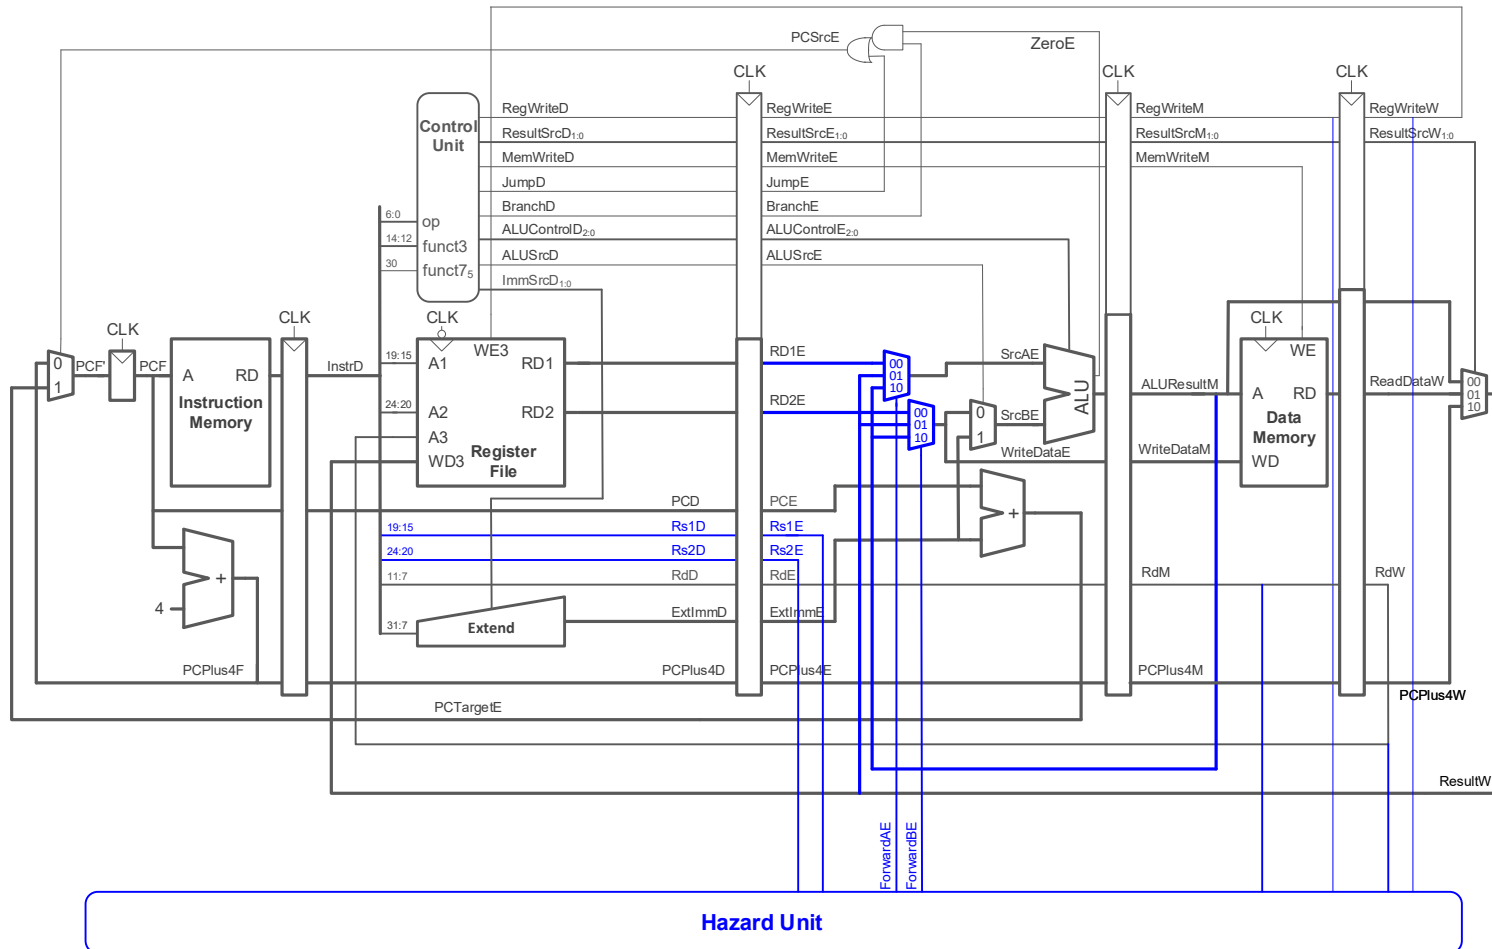


# Data Forwarding

- Check if source register **in Execute stage matches** destination register of instruction in **Memory or Writeback stage**.
- If so, forward result.
- **For all register?**



# Data Forwarding: Hazard Unit



## Data Forwarding

- **Case 1: Execute** stage  $Rs1$  or  $Rs2$  matches **Memory** stage  $Rd$ ?  
Forward from Memory stage
- **Case 2: Execute** stage  $Rs1$  or  $Rs2$  matches **Writeback** stage  $Rd$ ?  
Forward from Writeback stage
- **Case 3:** Otherwise use value read from register file (as usual)

- **Equations for ForwardAE -  $Rs1$ :**

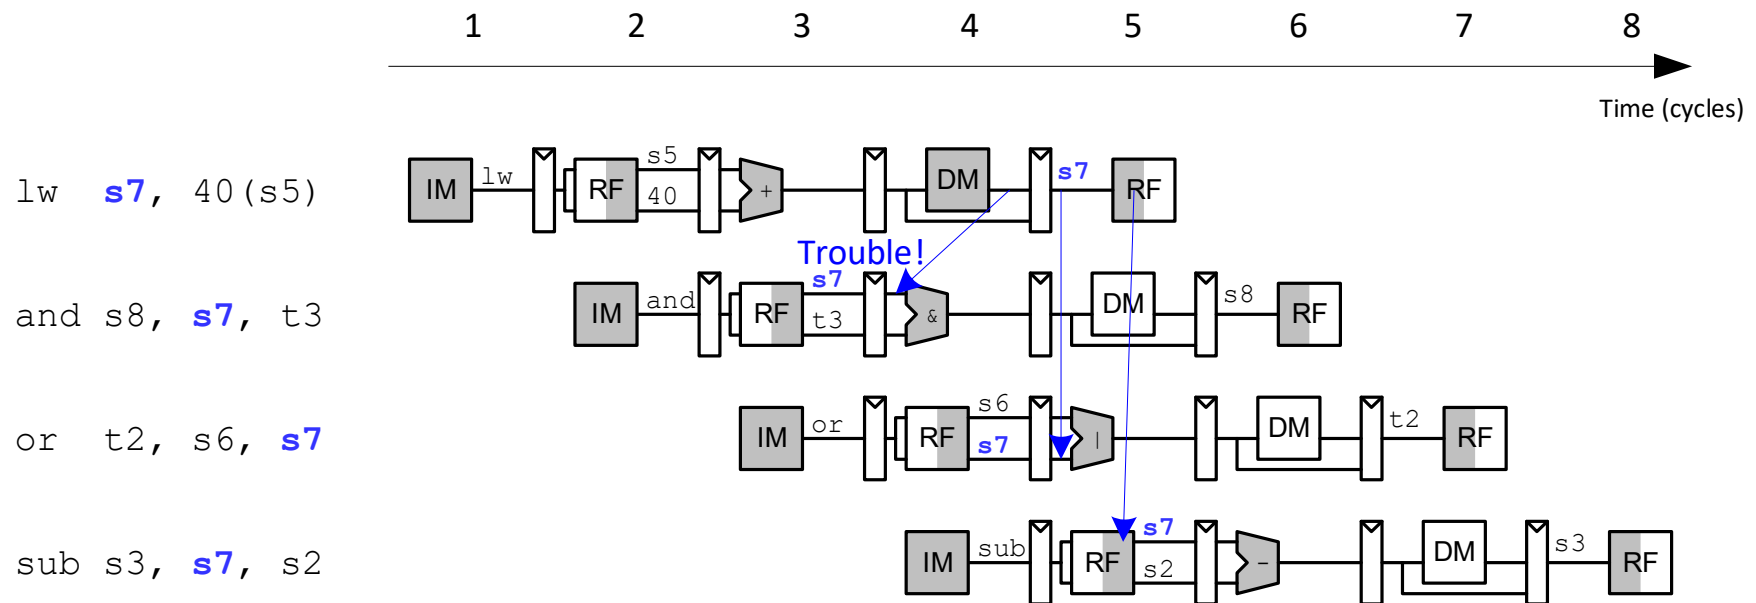
```
if    (( $Rs1E == RdM$ ) AND  $RegWriteM$ ) AND ( $Rs1E != 0$ ) // Case 1
      ForwardAE = 10
else if (( $Rs1E == RdW$ ) AND  $RegWriteW$ ) AND ( $Rs1E != 0$ ) // Case 2
      ForwardAE = 01
else      ForwardAE = 00 // Case 3
```

- **ForwardBE** equations are similar (replace  $Rs1E$  with  $Rs2E$ )

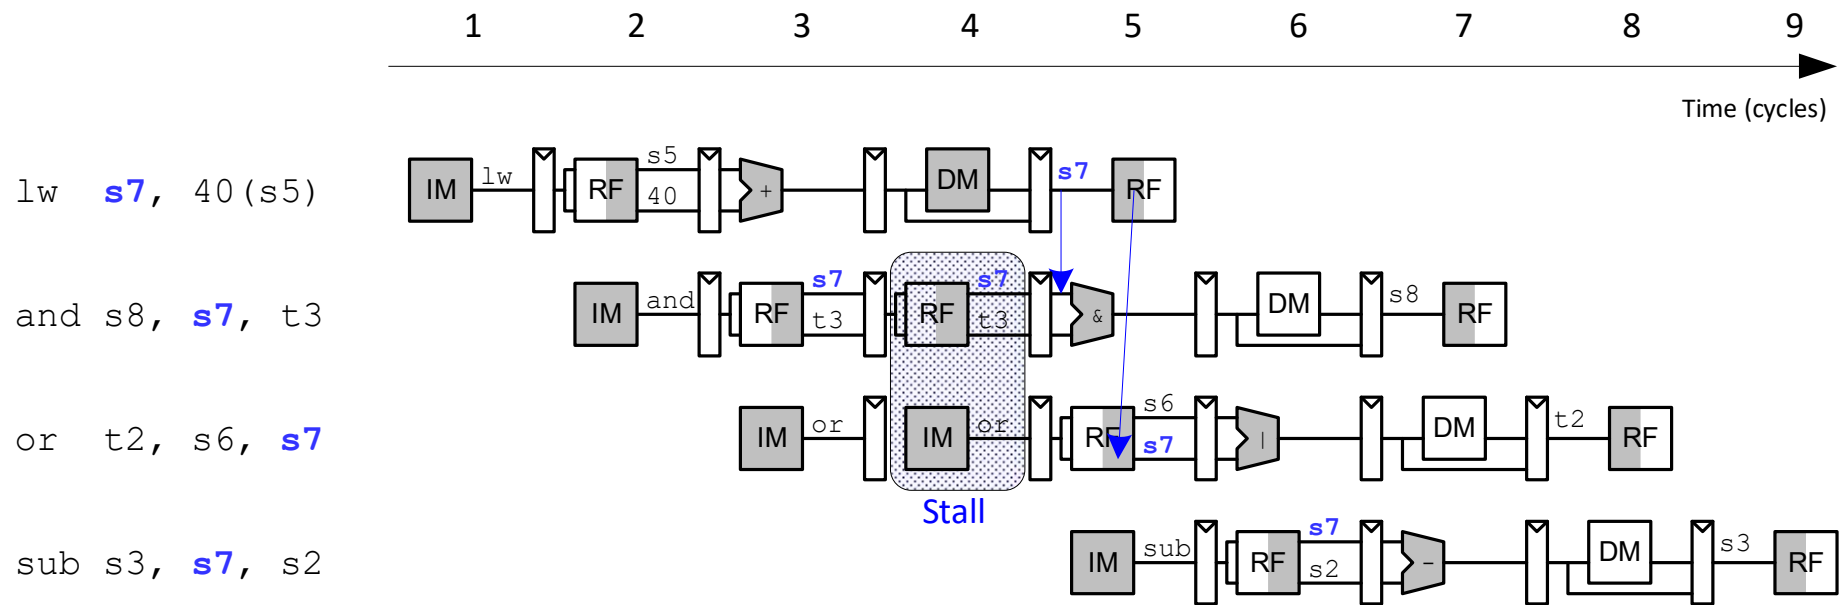




# Data Hazard due to $lw$ Dependency



# Data Hazard due to $lw$ Dependency



## Stalling Logic

- Is either **source register in the Decode stage** the same as the **destination register in the Execute stage**?

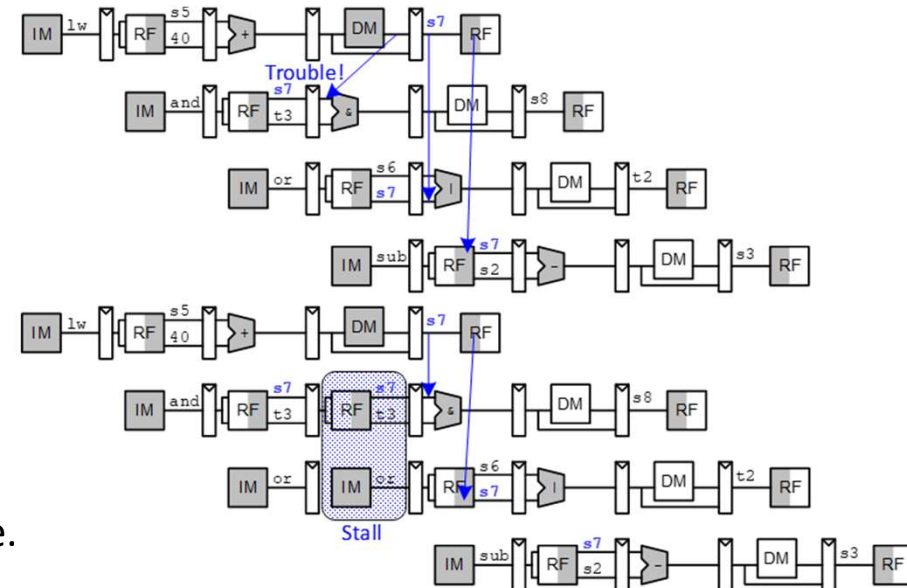
**AND**

- Is the instruction in the **Execute stage a lw**?

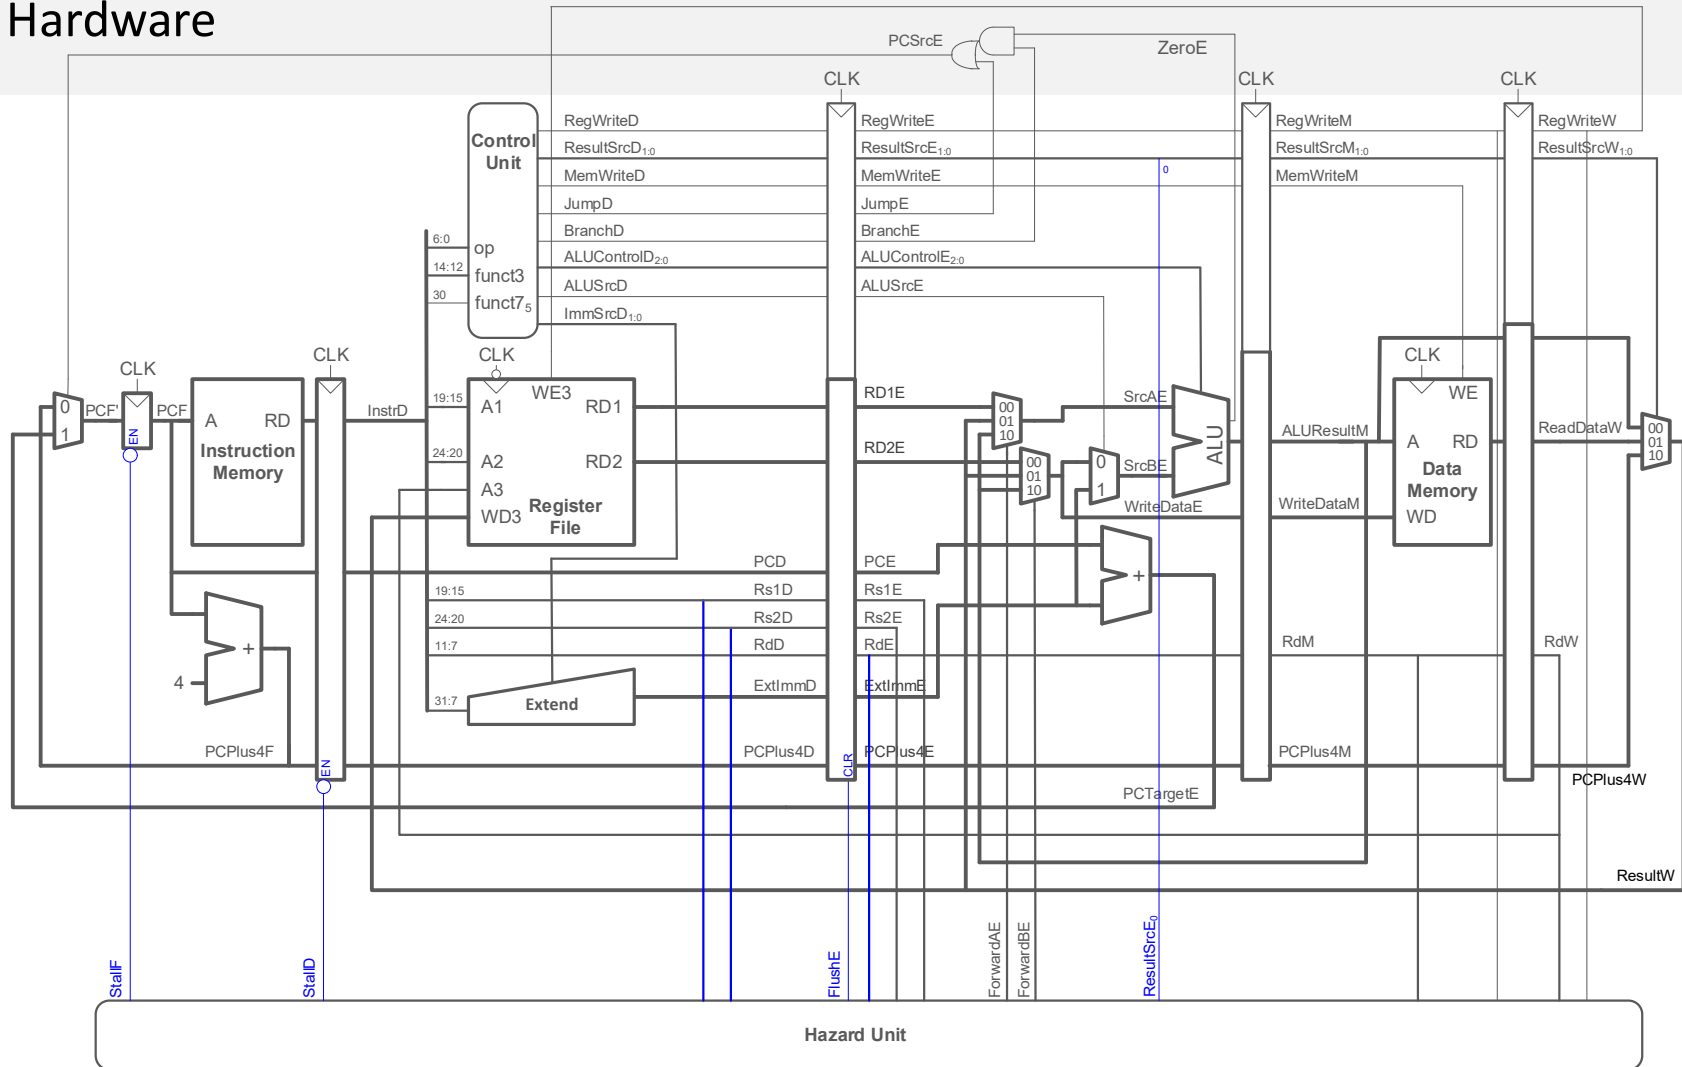
$lwStall = ((Rs1D == RdE) \text{ OR } (Rs2D == RdE)) \text{ AND } ResultSrcE_0$

$StallF = StallD = FlushE = lwStall$

(Stall the Fetch and Decode stages, and flush the Execute stage.)



# Stalling Hardware



# Pipelined Processor Control Hazards

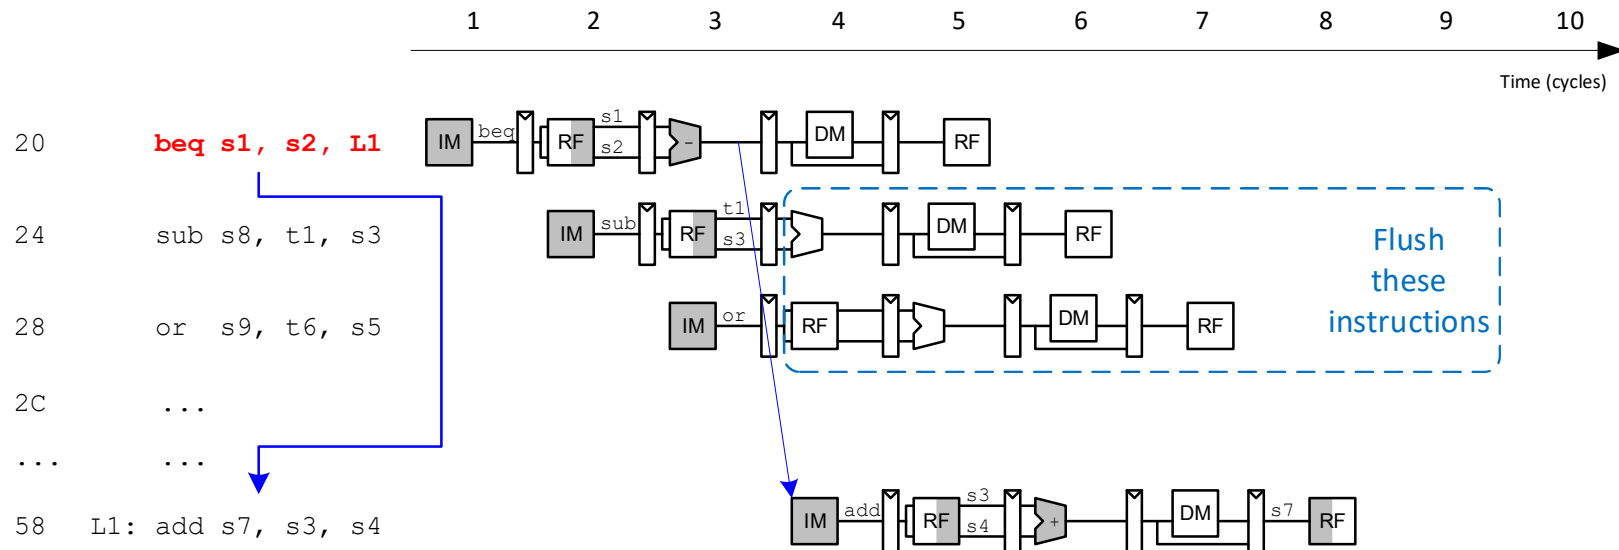
---

DDCA Ch7 - Part 15: Pipelined Processor Control Hazards <https://www.youtube.com/watch?v=VcnwVxD4LAc>

# Control Hazards

- **beq:**

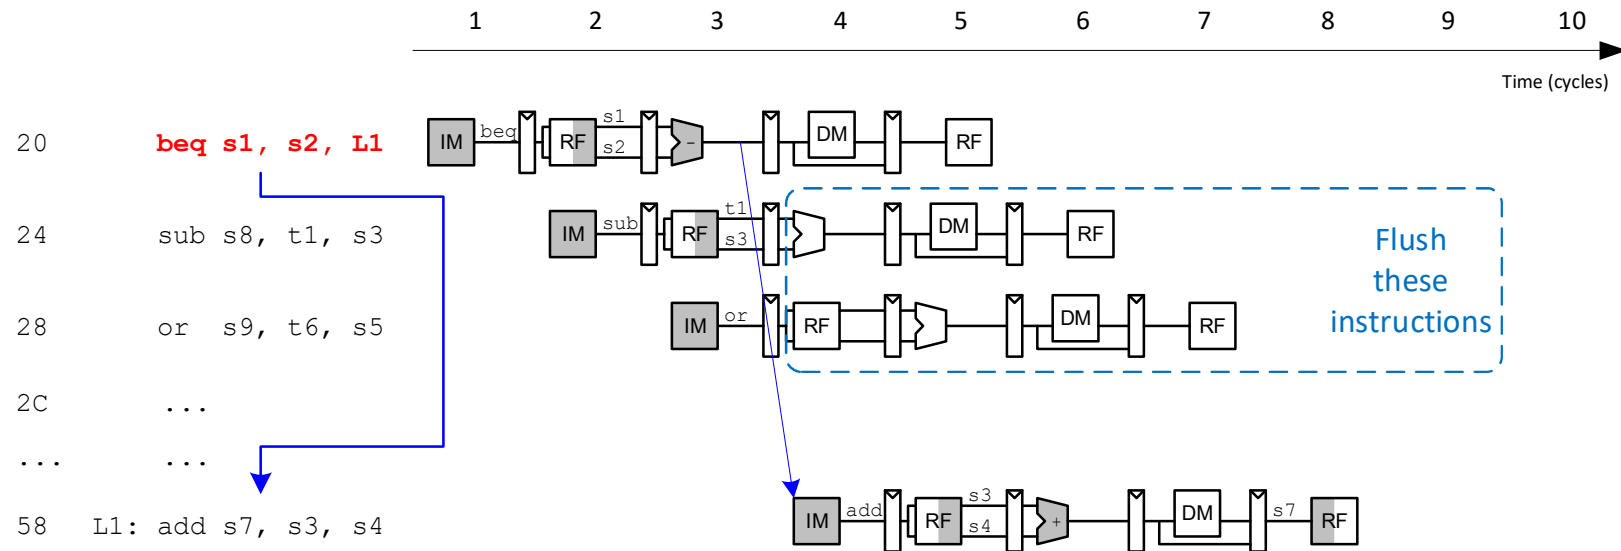
- Branch **not determined** until the **Execute** stage of pipeline
- **Instructions** after branch  **fetched** before branch occurs
- These **2 instructions must be flushed** if branch happens



# Control Hazards

## Branch misprediction penalty:

- The number of instructions flushed when a branch is taken (in this case, 2 instructions)





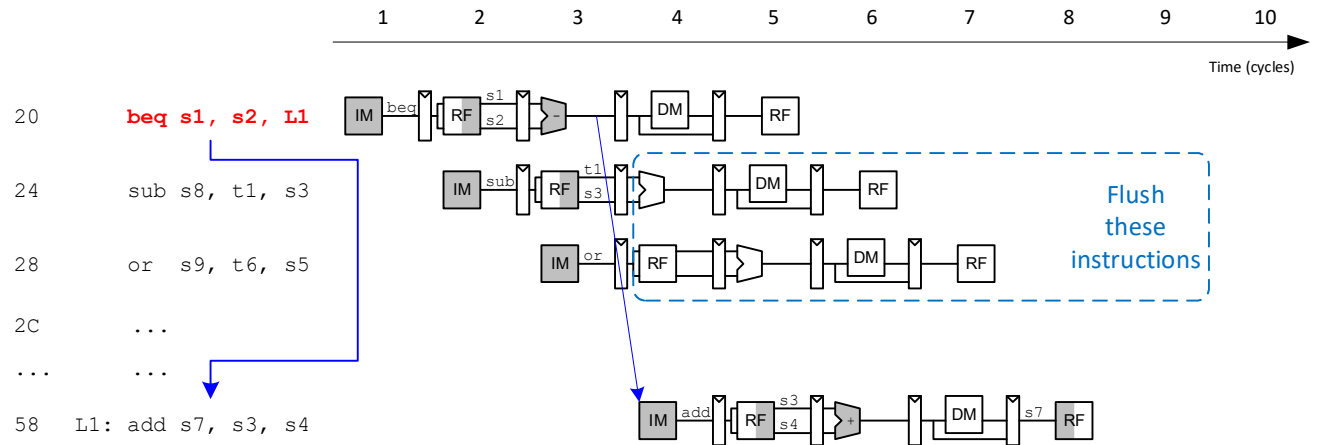
# Control Hazards: Flushing Logic

- If branch is taken in execute stage, need to flush the instructions in the Fetch and Decode stages
  - Do this by clearing Decode and Execute Pipeline registers using *FlushD* and *FlushE*

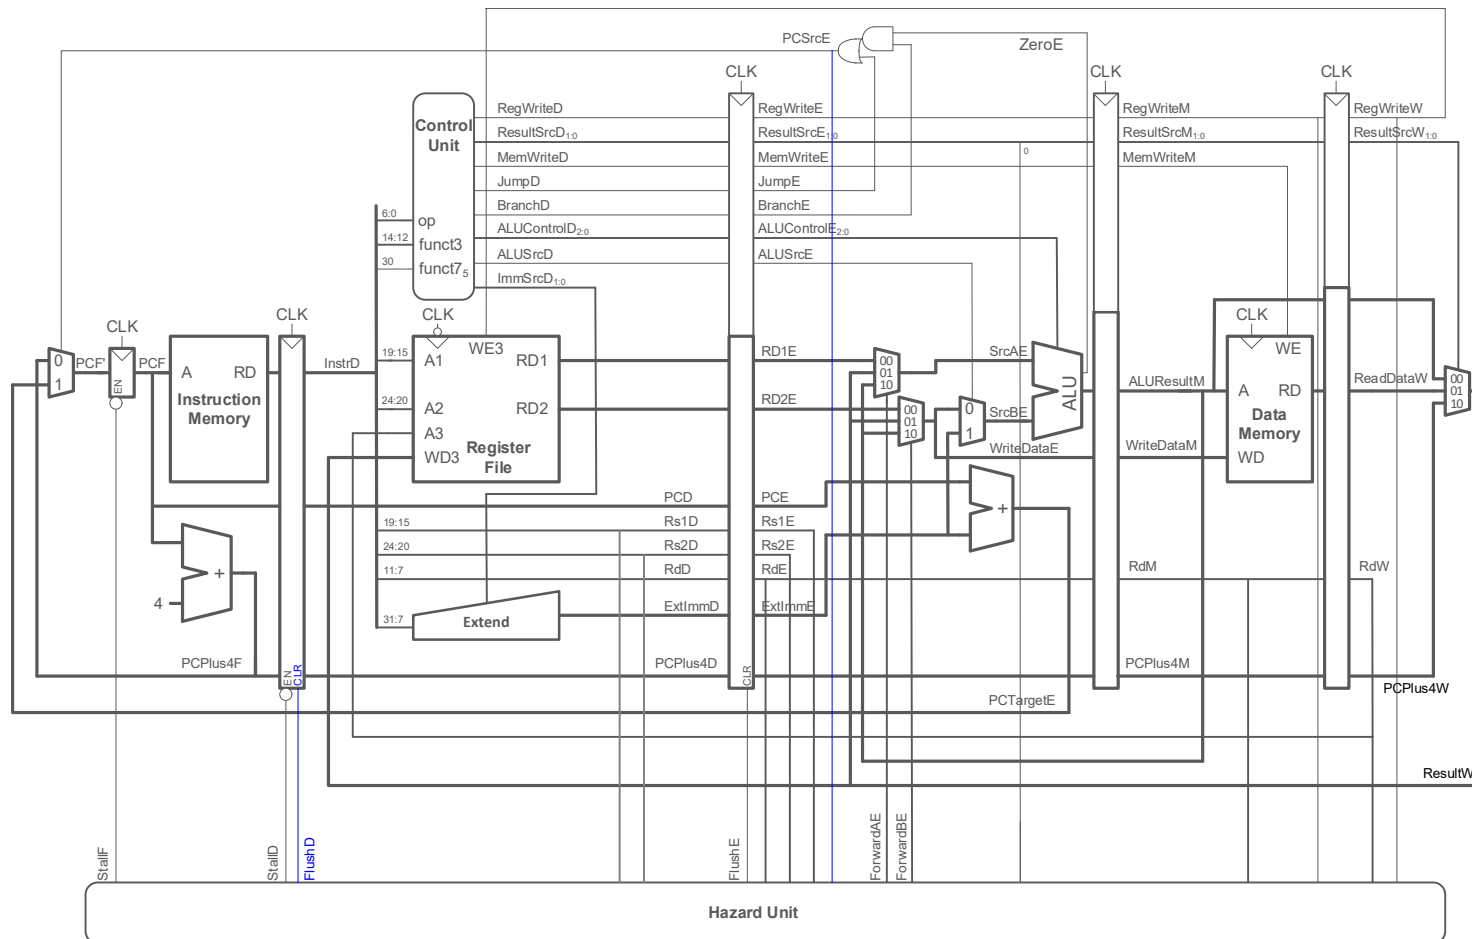
## Equations:

$$FlushD = PCSrcE$$

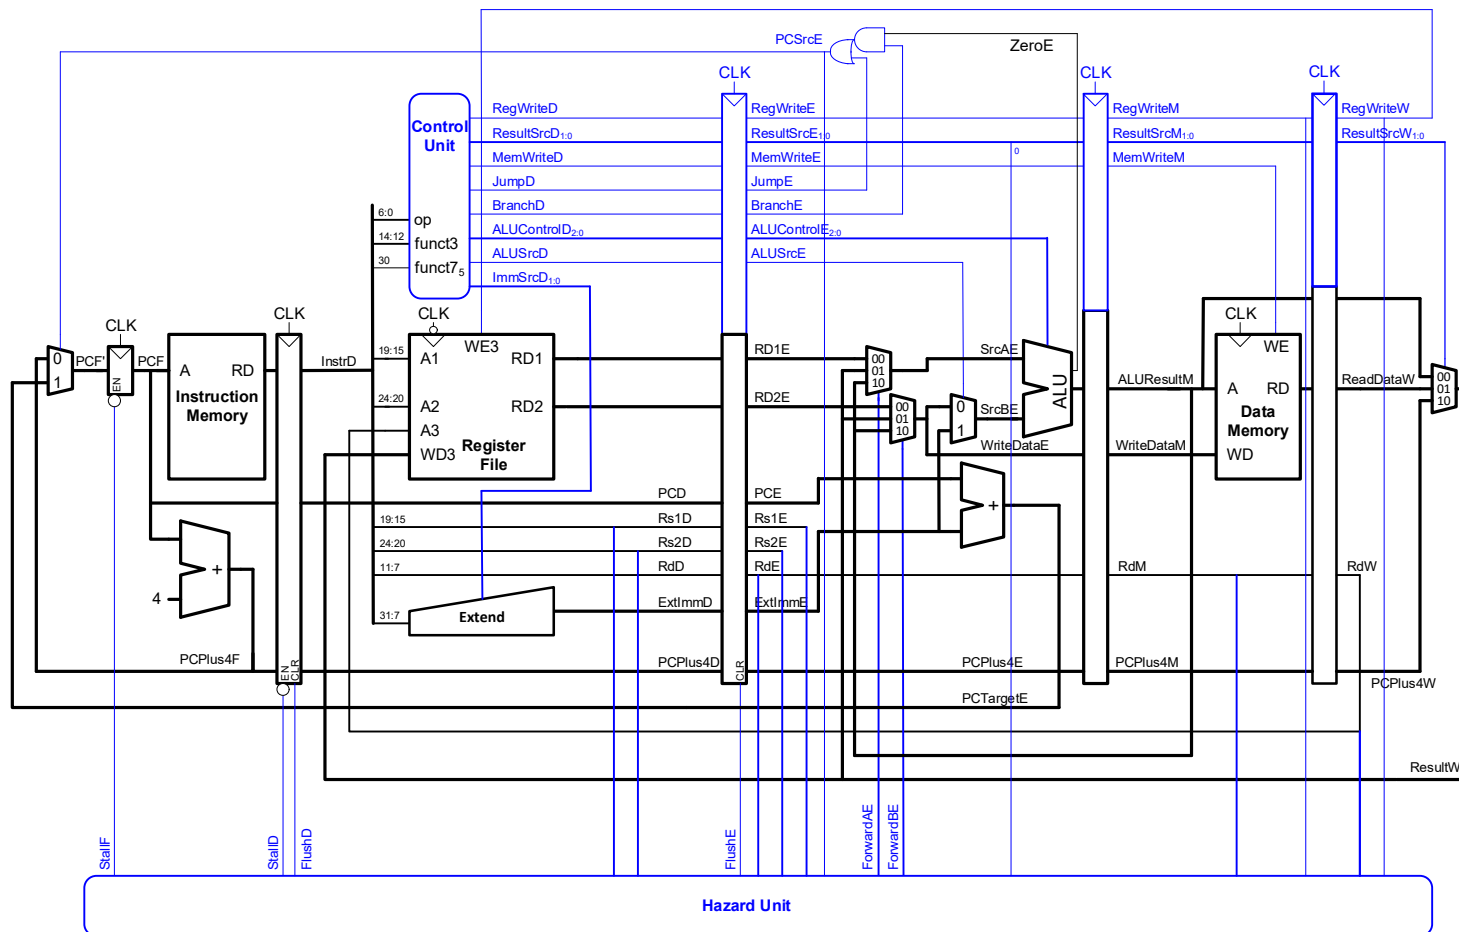
$$FlushE = lwStall \text{ OR } PCSrcE$$



# Control Hazards: Flushing Hardware



# RISC-V Pipelined Processor with Hazard Unit



# Summary of Hazard Logic

## Data hazard logic (shown for SrcA of ALU):

if  $((Rs1E == RdM) \text{ AND } RegWriteM) \text{ AND } (Rs1E != 0)$  // Case 1

$ForwardAE = 10$

else if  $((Rs1E == RdW) \text{ AND } RegWriteW) \text{ AND } (Rs1E != 0)$  // Case 2

$ForwardAE = 01$

else // Case 3

$ForwardAE = 00$

## Load word stall logic:

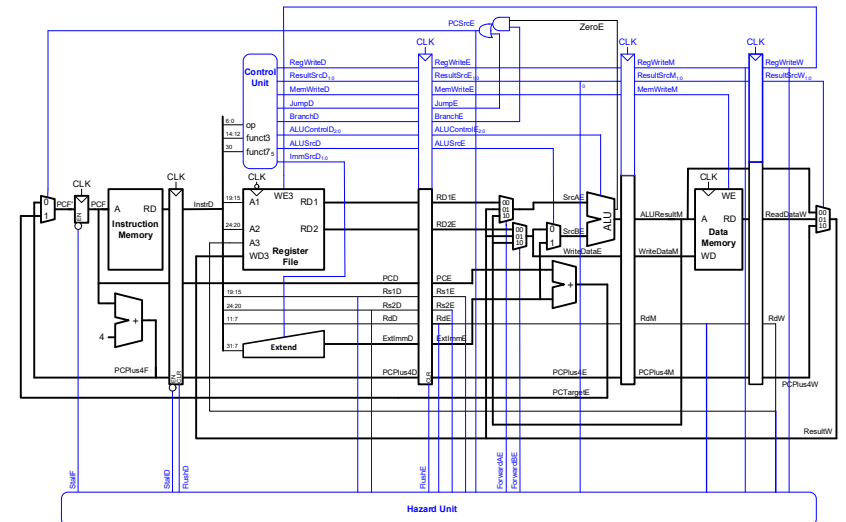
$lwStall = ((Rs1D == RdE) \text{ OR } (Rs2D == RdE)) \text{ AND } ResultSrcE_0$

$StallF = StallD = lwStall$

## Control hazard flush:

$FlushD = PCSrcE$

$FlushE = lwStall \text{ OR } PCSrcE$



# Pipelined Performance

---

DDCA Ch7 - Part 14: Pipelined Processor Data Hazards <https://www.youtube.com/watch?v=zuegcg6ZSFQ>

## Pipelined Processor Performance Example

- **SPECINT2000 benchmark:**

- 25% loads
- 10% stores
- 13% branches
- 52% R-type

- **Suppose:**

- 40% of loads used by next instruction
- 50% of branches mispredicted

- **What is the average CPI?**

(Ideally it's 1, but...)

- Load CPI = 1 when not stalling, 2 when stalling
- Branch CPI = 1 when not stalling, 3 when stalling

$$\rightarrow \text{So, } \text{CPI}_{\text{lw}} = 1(0.6) + 2(0.4) = 1.4$$

$$\rightarrow \text{So, } \text{CPI}_{\text{beq}} = 1(0.5) + 3(0.5) = 2$$

- **Average CPI =  $(0.25)(1.4) + (0.1)(1) + (0.13)(2) + (0.52)(1) = 1.23$**

## Pipelined Processor Performance Example

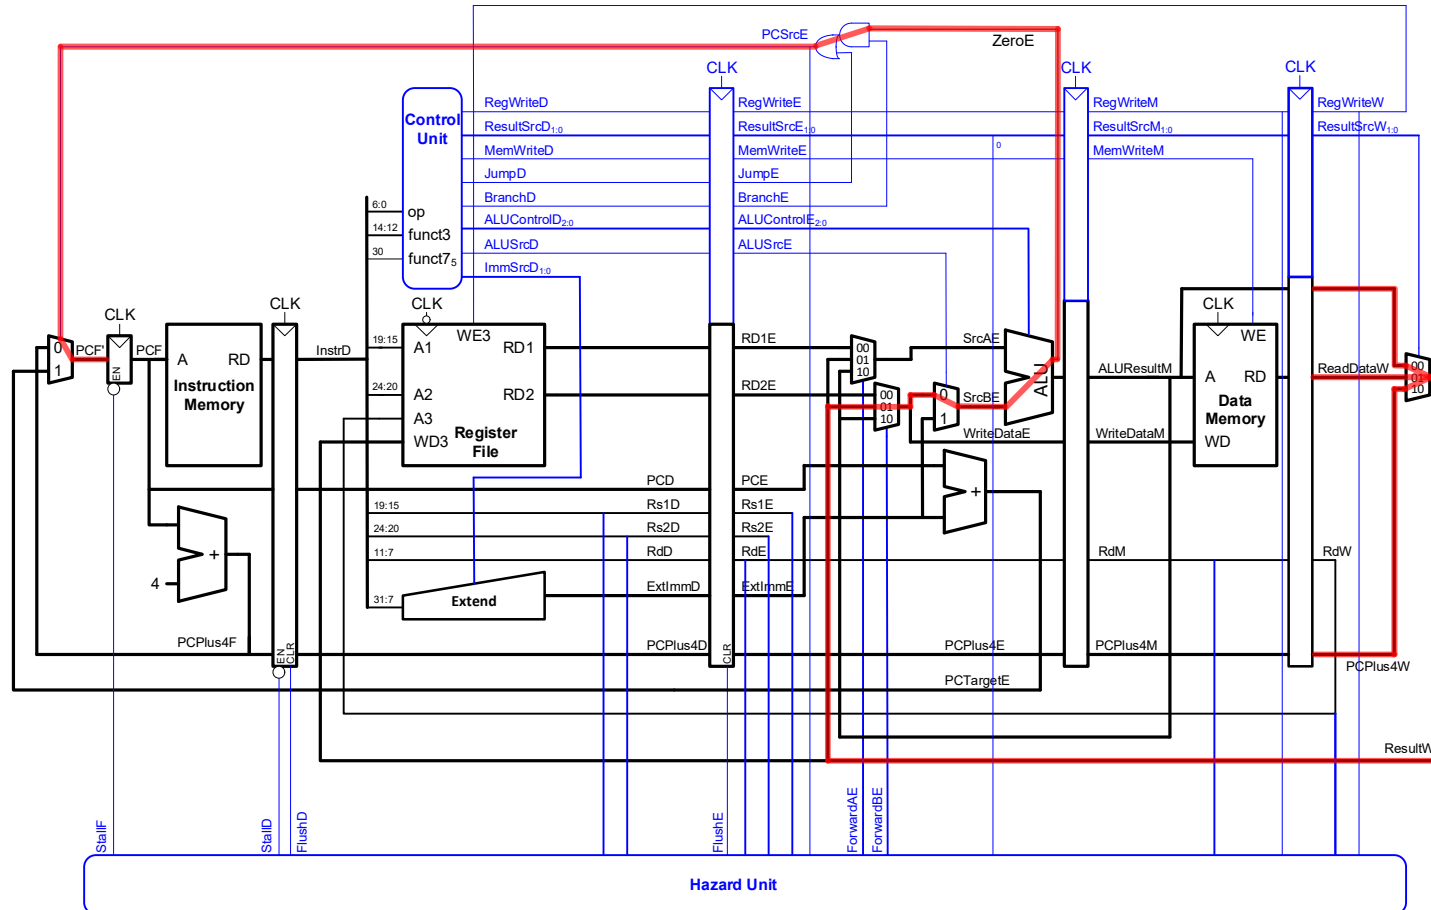
- Pipelined processor critical path:

$$T_{c\_pipelined} = \max \left[ \begin{array}{l} t_{pcq} + t_{mem} + t_{setup} \\ 2(t_{RFread} + t_{setup}) \\ t_{pcq} + 4t_{mux} + t_{ALU} + t_{AND-OR} + t_{setup} \\ t_{pcq} + t_{mem} + t_{setup} \\ 2(t_{pcq} + t_{mux} + t_{RFwrite}) \end{array} \right]$$

Fetch  
Decode  
Execute  
Memory  
Writeback

- Decode and Writeback stages **both use the register file** in each cycle
- So each stage gets half of the cycle time ( $T_c/2$ ) to do their work
- Or, stated a different way, **2x of their work** must fit in a cycle ( $T_c$ )

# Pipelined Critical Path: Execute Stage





## Pipelined Performance Example

Element	Parameter	Delay (ps)
Register clock-to-Q	$t_{pcq}$ PC	40
Register setup	$t_{setup}$	50
Multiplexer	$t_{mux}$	30
AND-OR gate	$t_{AND-OR}$	20
ALU	$t_{ALU}$	120
Decoder (Control Unit)	$t_{dec}$	25
Extend unit	$t_{dec}$	35
Memory read	$t_{mem}$	200
Register file read	$t_{RFread}$	100
Register file setup	$t_{RFsetup}$	60

$$\begin{aligned} T_{c\_pipelined} &= t_{pcq} + 4t_{mux} + t_{ALU} + t_{AND-OR} + t_{setup} \\ &= (40 + 4*30 + 120 + 20 + 50) \text{ ps} = \mathbf{350 \text{ ps}} \end{aligned}$$

## Pipelined Performance Example

Program with 100 billion instructions

$$\begin{aligned}\text{Execution Time} &= (\# \text{ instructions}) \times \text{CPI} \times T_c \\ &= (100 \times 10^9)(1.23)(350 \times 10^{-12}) \\ &= \mathbf{43 \text{ seconds}}\end{aligned}$$

## Processor Performance Comparison

Processor	Execution Time (seconds)	Speedup (single-cycle as baseline)
Single-cycle	75	1
Multi-cycle	155	0.5
Pipelined	43	1.7