



Chapter 4A

The Processor

Introduction


- CPU performance factors
 - Instruction count
 - Determined by ISA and compiler
 - CPI and Cycle time
 - Determined by CPU hardware
- We will examine two MIPS implementations
 - A simplified version
 - A more realistic pipelined version
- Simple subset, shows most aspects
 - Memory reference: `lw`, `sw`
 - Arithmetic/logical: `add`, `sub`, `and`, `or`, `slt`
 - Control transfer: `beq`, `j`

MIPS Microarchitecture

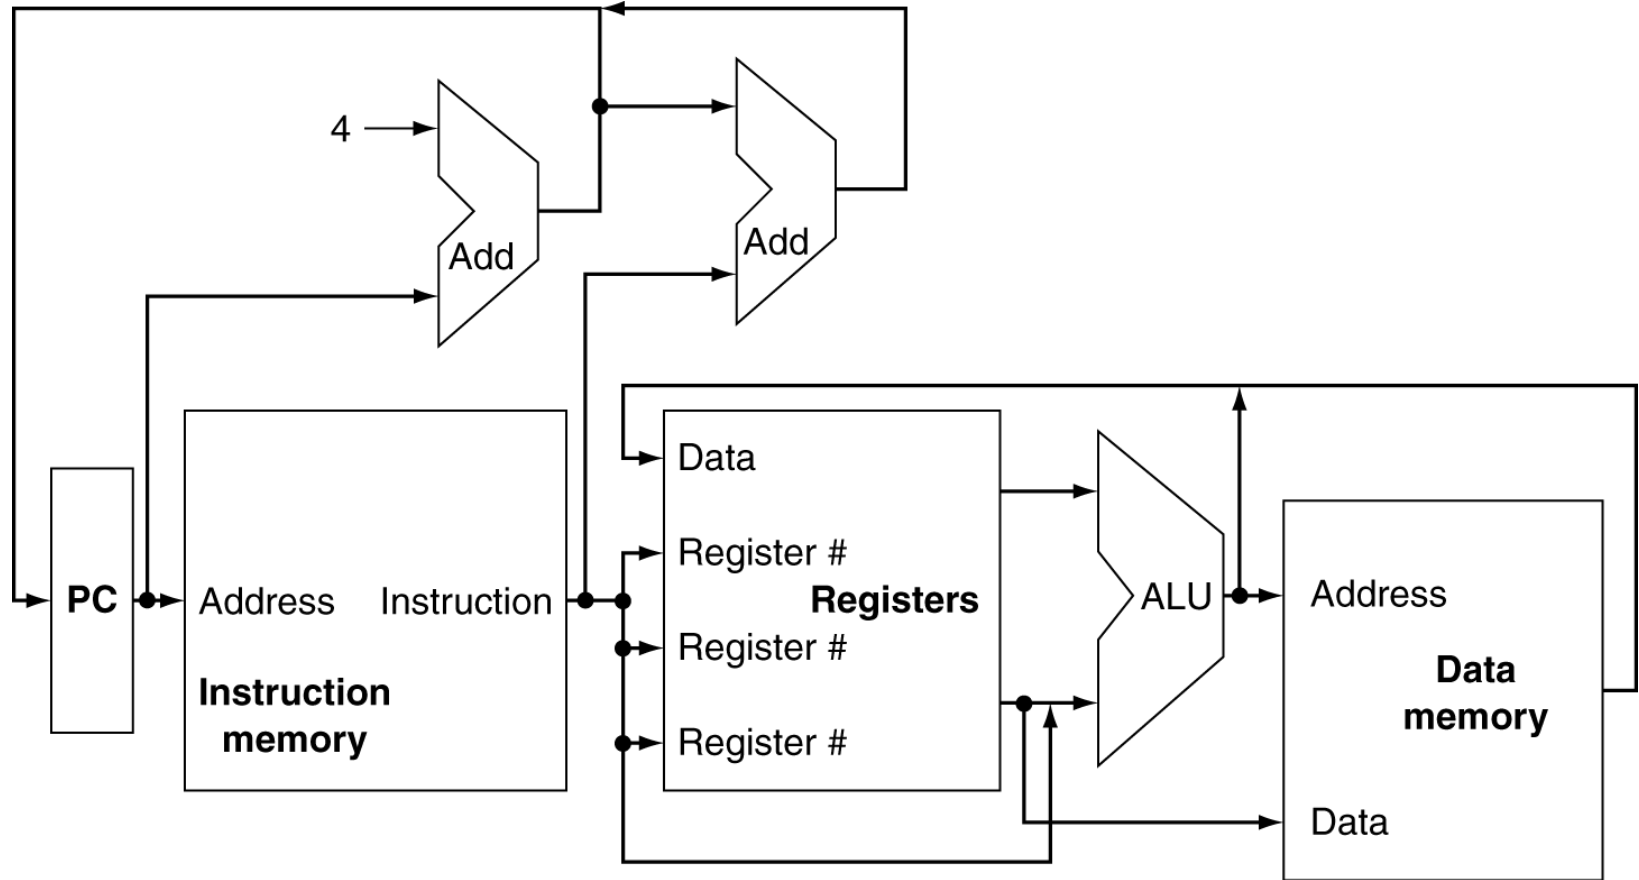
- Building blocks
 - Registers
 - Memories (instruction, data)
 - ALUs
 - Muxes
 - Wires

Instruction Execution

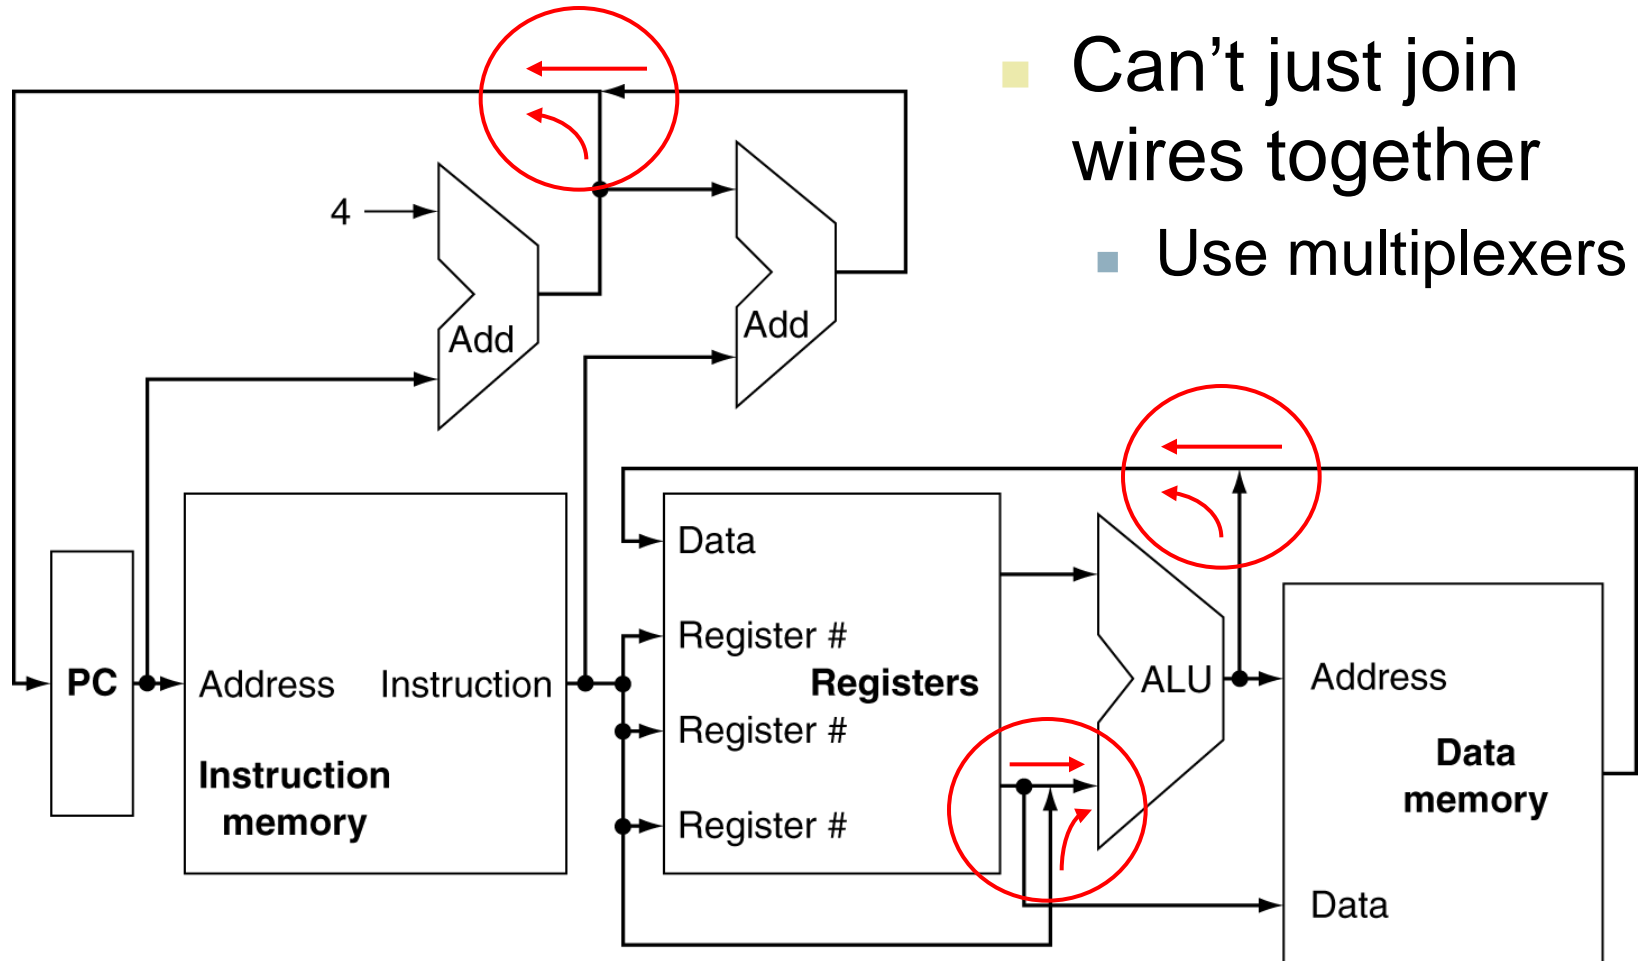
- PC → **instruction memory**, fetch instruction
- Register numbers → **register file**, read registers
- Depending on instruction class
 - Use **ALU** to calculate
 - Arithmetic result
 - Memory address for load/store
 - Branch target address
 - Access **data memory** for load/store
 - $PC \leftarrow \text{target address or } PC + 4$



CPU Overview



Multiplexers




Logic Design Basics

- Information encoded in binary
 - Low voltage = 0, High voltage = 1
 - One wire per bit
 - Multi-bit data encoded on multi-wire buses
- Combinational element
 - Operate on data
 - Output is a function of input
- State (sequential) elements
 - Store information

Combinational Elements


- AND-gate

- $Y = A \& B$




- Multiplexer

- $Y = S ? I1 : I0$




- Adder

- $Y = A + B$





- Arithmetic/Logic Unit

- $Y = F(A, B)$





Sequential Elements

- Register: stores data in a circuit
 - Uses a clock signal to determine when to update the stored value
 - Edge-triggered: update when Clk changes from 0 to 1




Sequential Elements

- Register with write control
 - **Only updates** on clock edge when write control input is 1
 - Used when stored value is required later



Clocking Methodology

- Combinational logic transforms data during clock cycles
 - Between clock edges
 - Input from state elements, output to state element
 - Longest delay determines clock period




Building a Datapath

- Datapath
 - Elements that process data and addresses in the CPU
 - Registers, ALUs, mux's, memories, ...
- We will build a MIPS datapath incrementally
 - Refining the overview design

Fetching Instructions

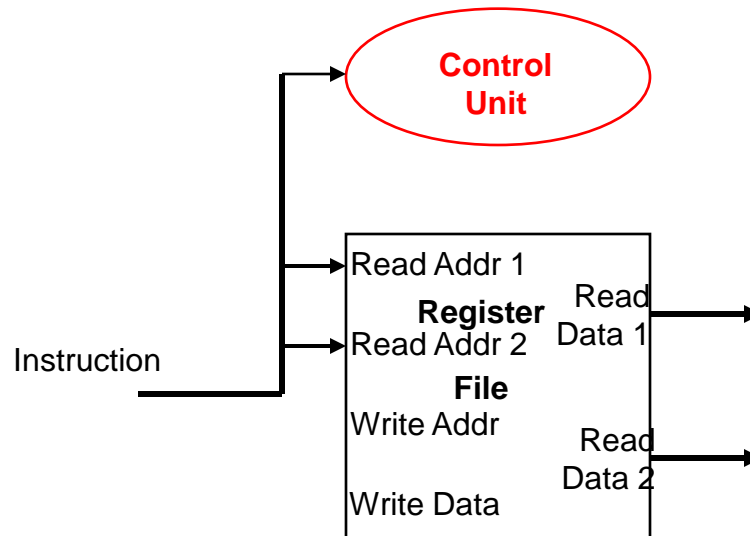
- Fetching instructions involves
 - reading the instruction from the Instruction Memory
 - updating the PC to hold the address of the next instruction



- PC is updated every cycle, so it does not need an explicit write control signal
- Instruction Memory is read every cycle, so it doesn't need an explicit read control signal

Decoding Instructions

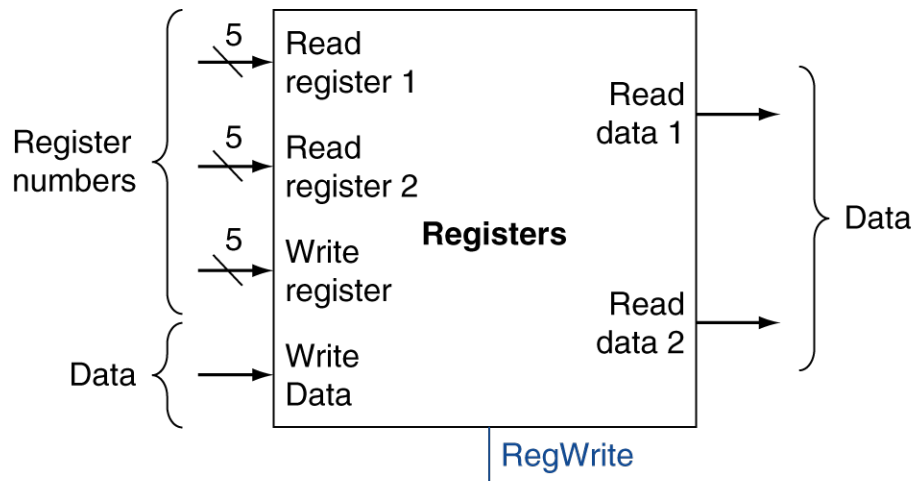
- Decoding instructions involves
 - sending the fetched instruction's opcode and function field bits to the control unit




- reading two values from the Register File
 - Register File addresses are contained in the instruction

R-Format Instructions

- Read two register operands
- Perform arithmetic/logical operation
- Write register result




a. Registers



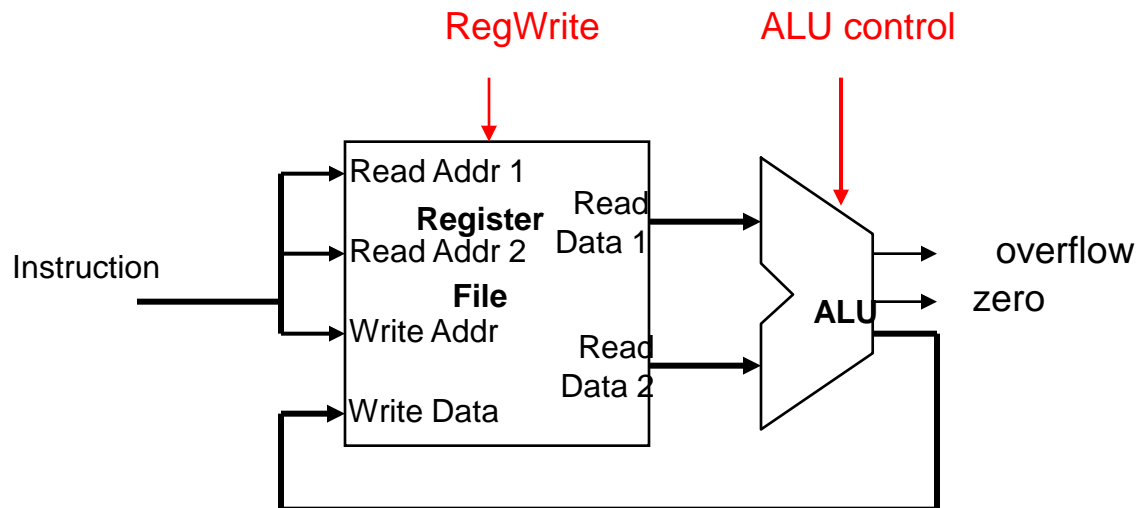
b. ALU

Executing R-type Operations

- R format operations (**add**, **sub**, **slt**, **and**, **or**)



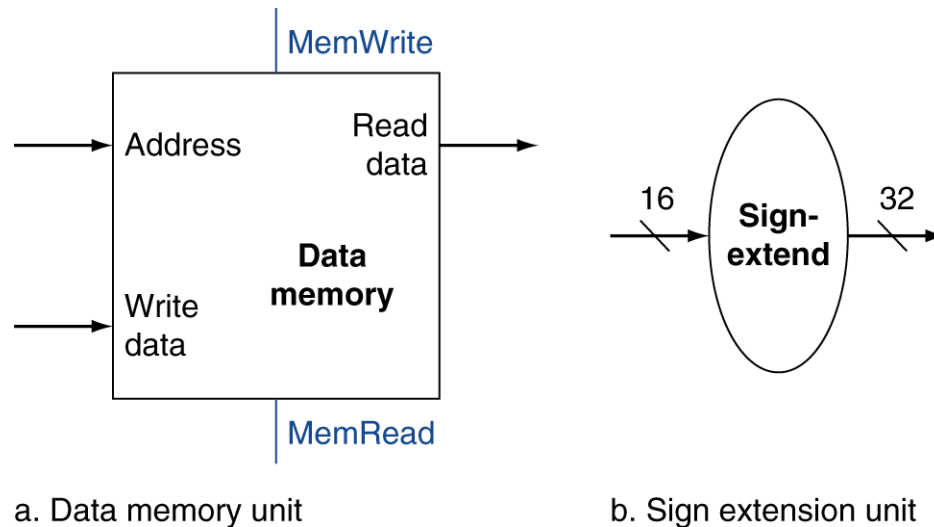
- perform the (**op** and **funct**) operation on values in **rs** and **rt**
- store the result back into the Register File (into location **rd**)



- The Register File is not written every cycle (e.g. **sw**), so we need an explicit write control signal for the Register File

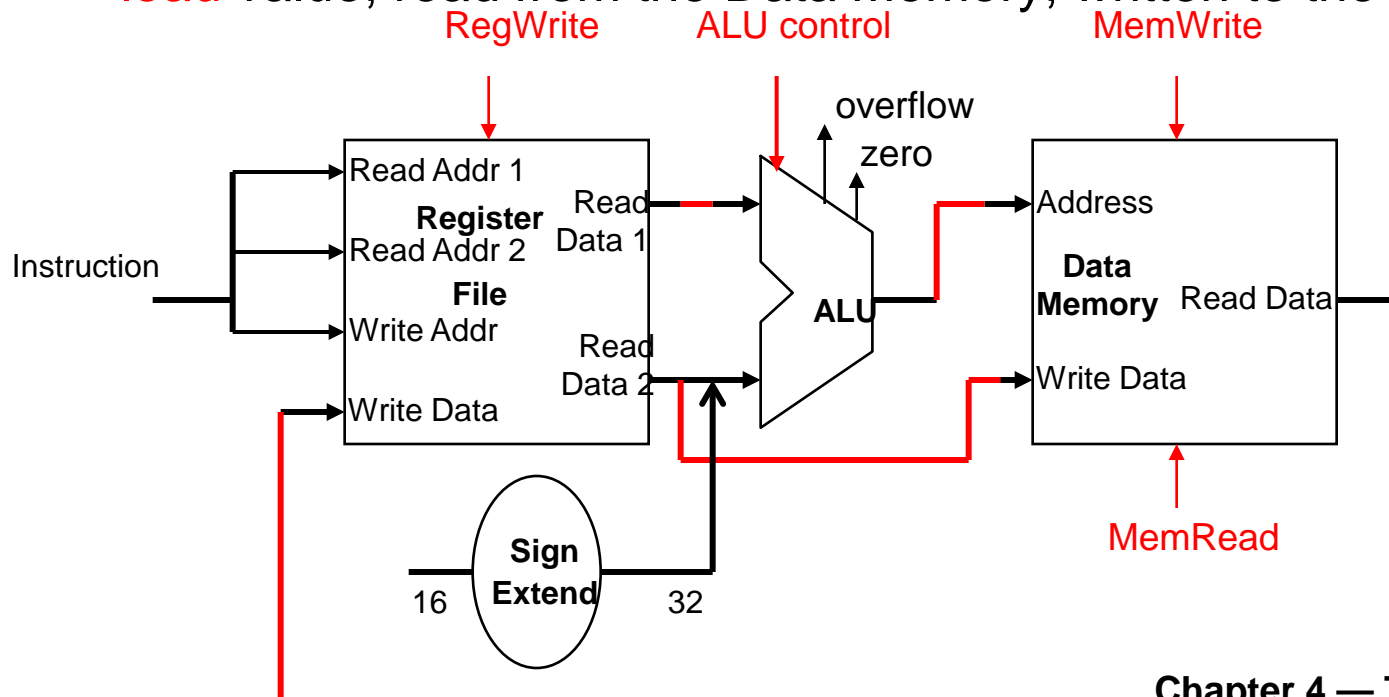
Load/Store Instructions

- Read register operands
- Calculate address using 16-bit offset
 - Use ALU, but sign-extend offset
- Load: Read memory and update register
- Store: Write register value to memory



Executing lw/sw Operations

- Load and store operations involves
 - compute memory address by adding the base register (read from the Register File during decode) to the 16-bit signed-extended offset field in the instruction
 - **store** value (read from the Register File during decode) written to the Data Memory
 - **load** value, read from the Data Memory, written to the Register File

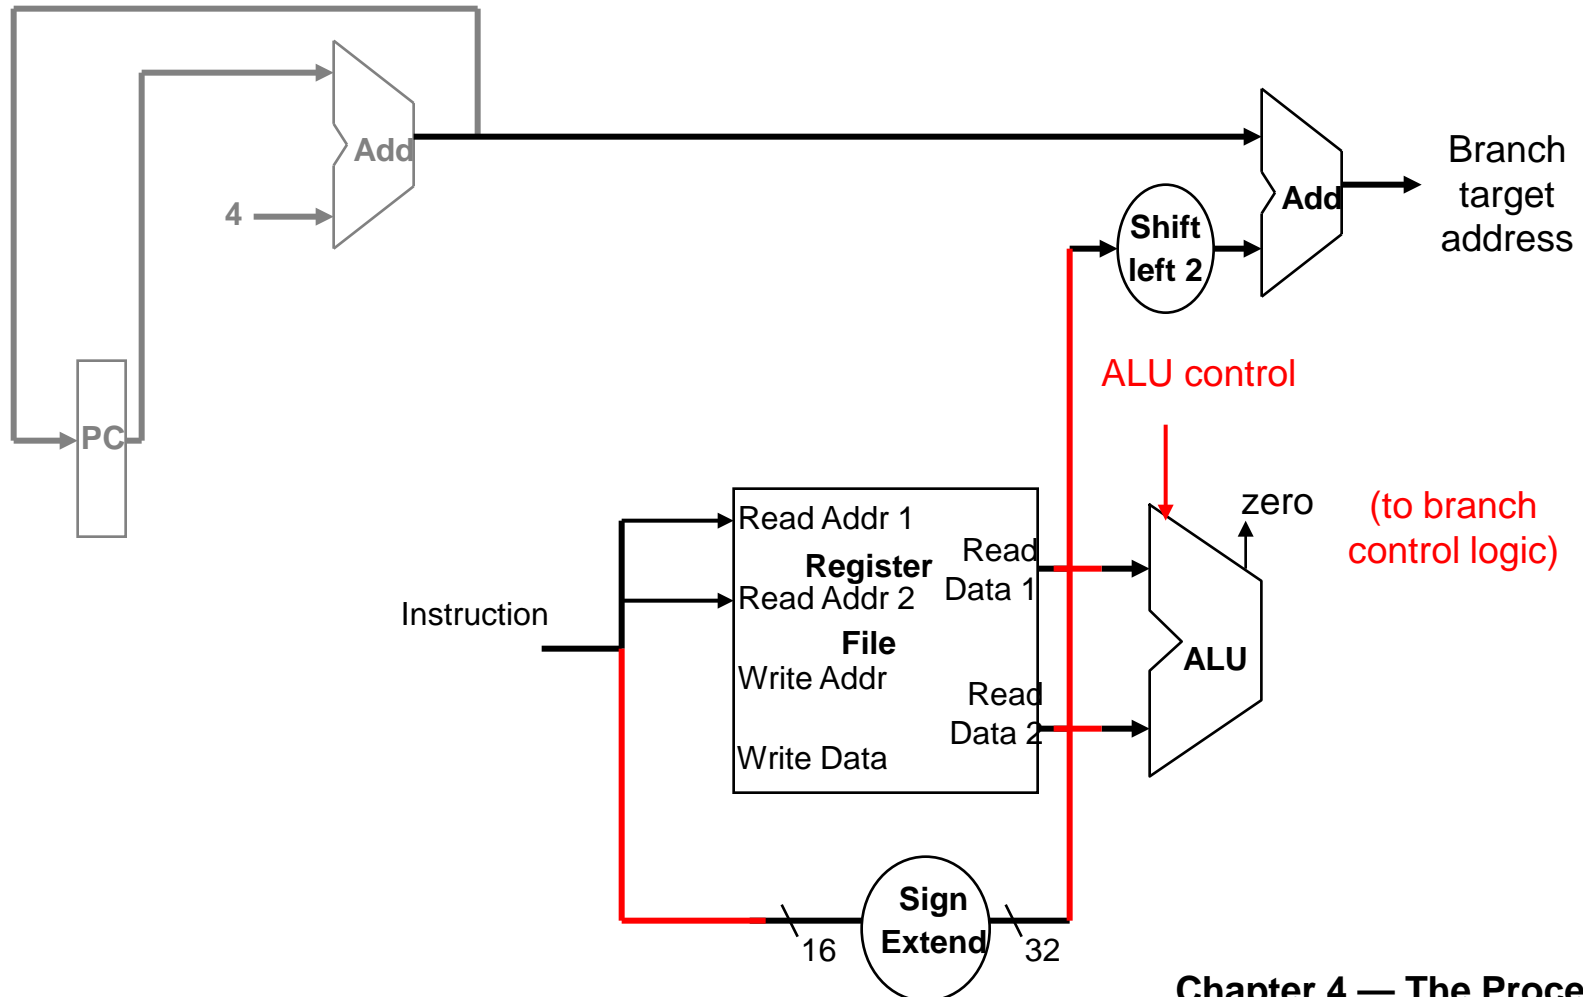


Branch Instructions

- Read register operands
- Compare operands
 - Use ALU, subtract and check Zero output
- Calculate target address
 - Sign-extend displacement
 - Shift left 2 places (word displacement)
 - Add to PC + 4
 - Already calculated by instruction fetch

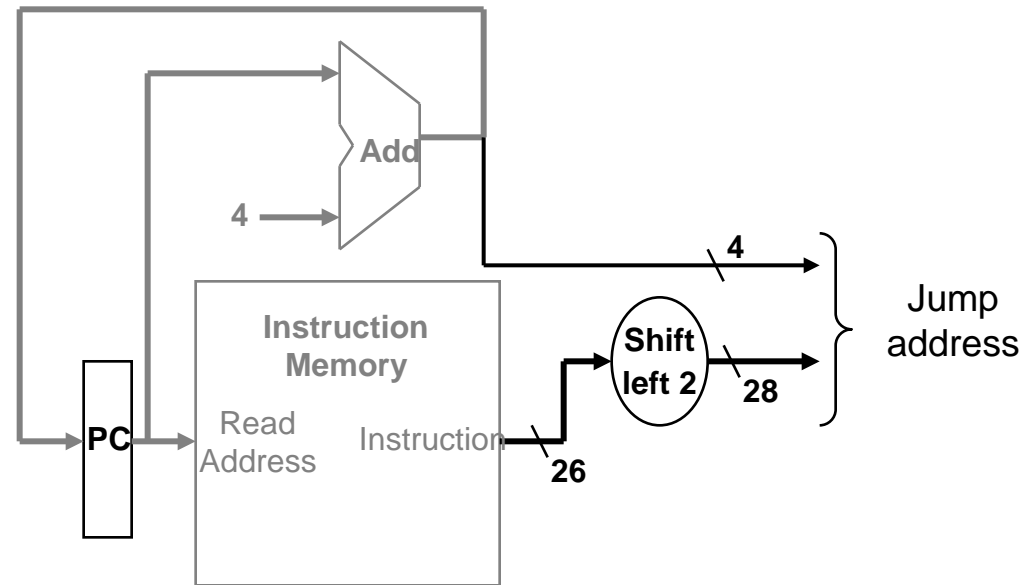
Executing Branch Operations

- Branch operations involves
 - compare registers + compute branch target



Executing Jump Operations

- Jump operation involves
 - replace the lower 28 bits of the PC with the lower 26 bits of the fetched instruction shifted left by 2 bits



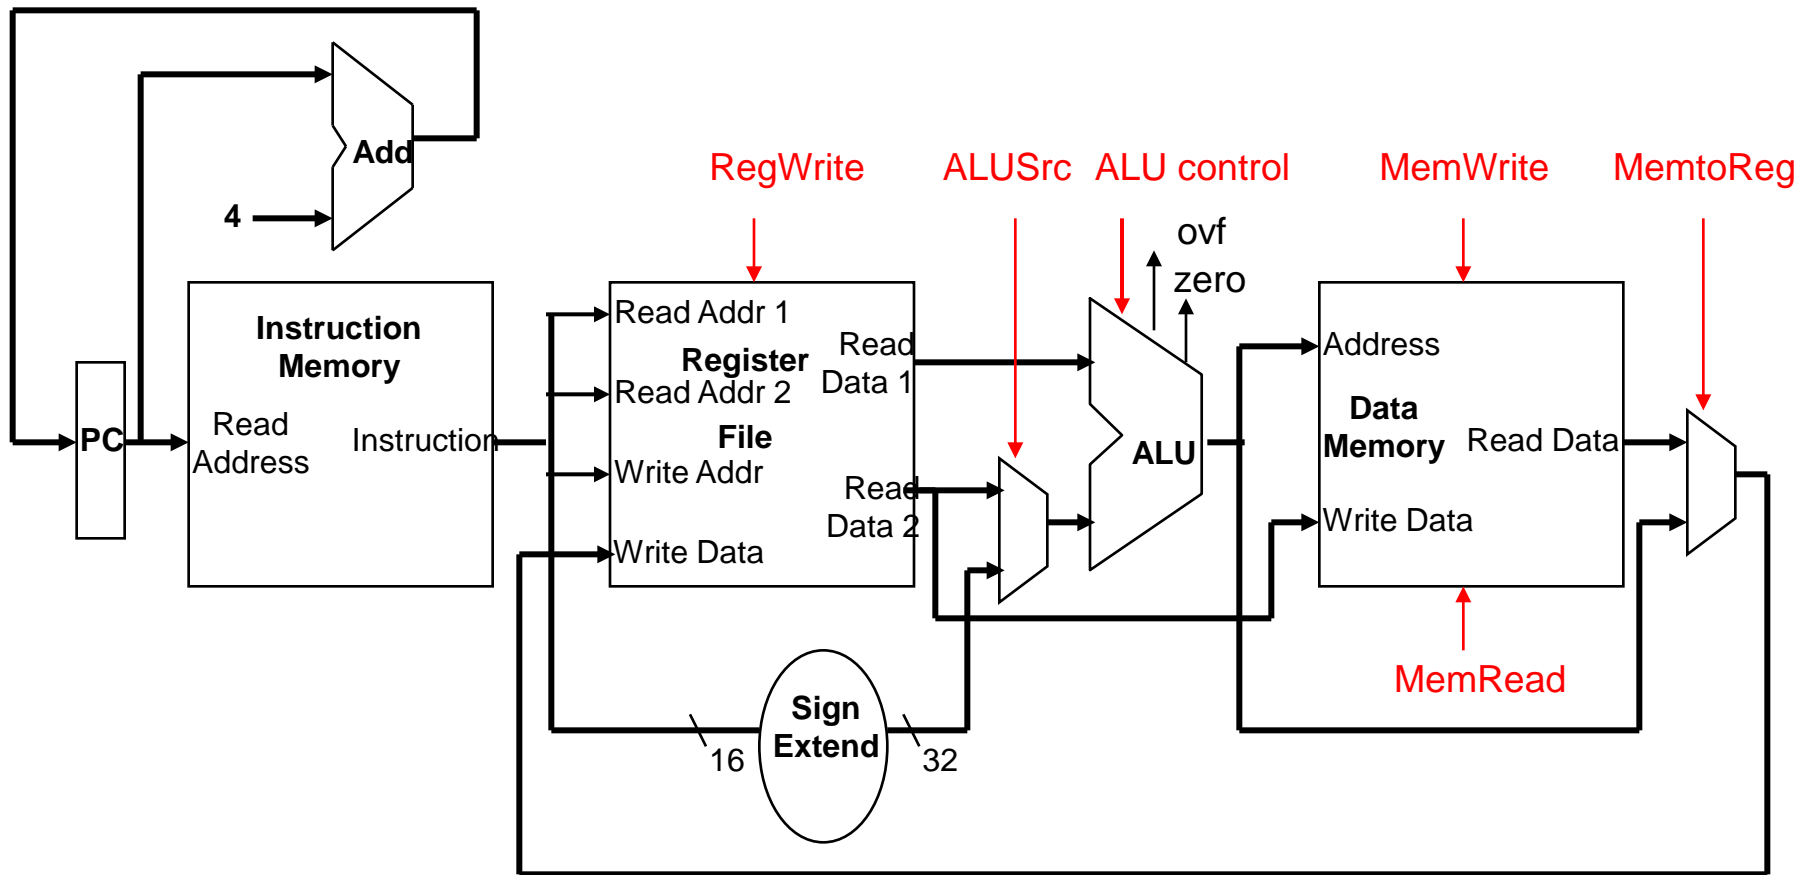
Composing the Elements

- First-cut data path does an instruction in one clock cycle
 - Each datapath element can only do one function at a time
 - Hence, we need separate instruction and data memories
- Use **multiplexers** where alternate data sources are used for different instructions

Assembling the Parts

- Assemble the datapath segments and add control lines and multiplexors as needed
- **Single cycle** design – fetch, decode and execute each instructions in **one** clock cycle
 - no datapath resource can be used more than once per instruction, so some must be duplicated (e.g., separate Instruction Memory and Data Memory, several adders)
 - **multiplexors** needed at the input of shared elements with control lines to do the selection
 - write signals to control writing to the Register File and Data Memory

Fetching, Register, and Memory Access Portions



ALU Control

- ALU used for
 - Load/Store: F = add
 - Branch: F = subtract
 - R-type: F depends on funct field

ALU control	Function
0000	AND
0001	OR
0010	add
0110	subtract
0111	set-on-less-than
1100	NOR

ALU Control

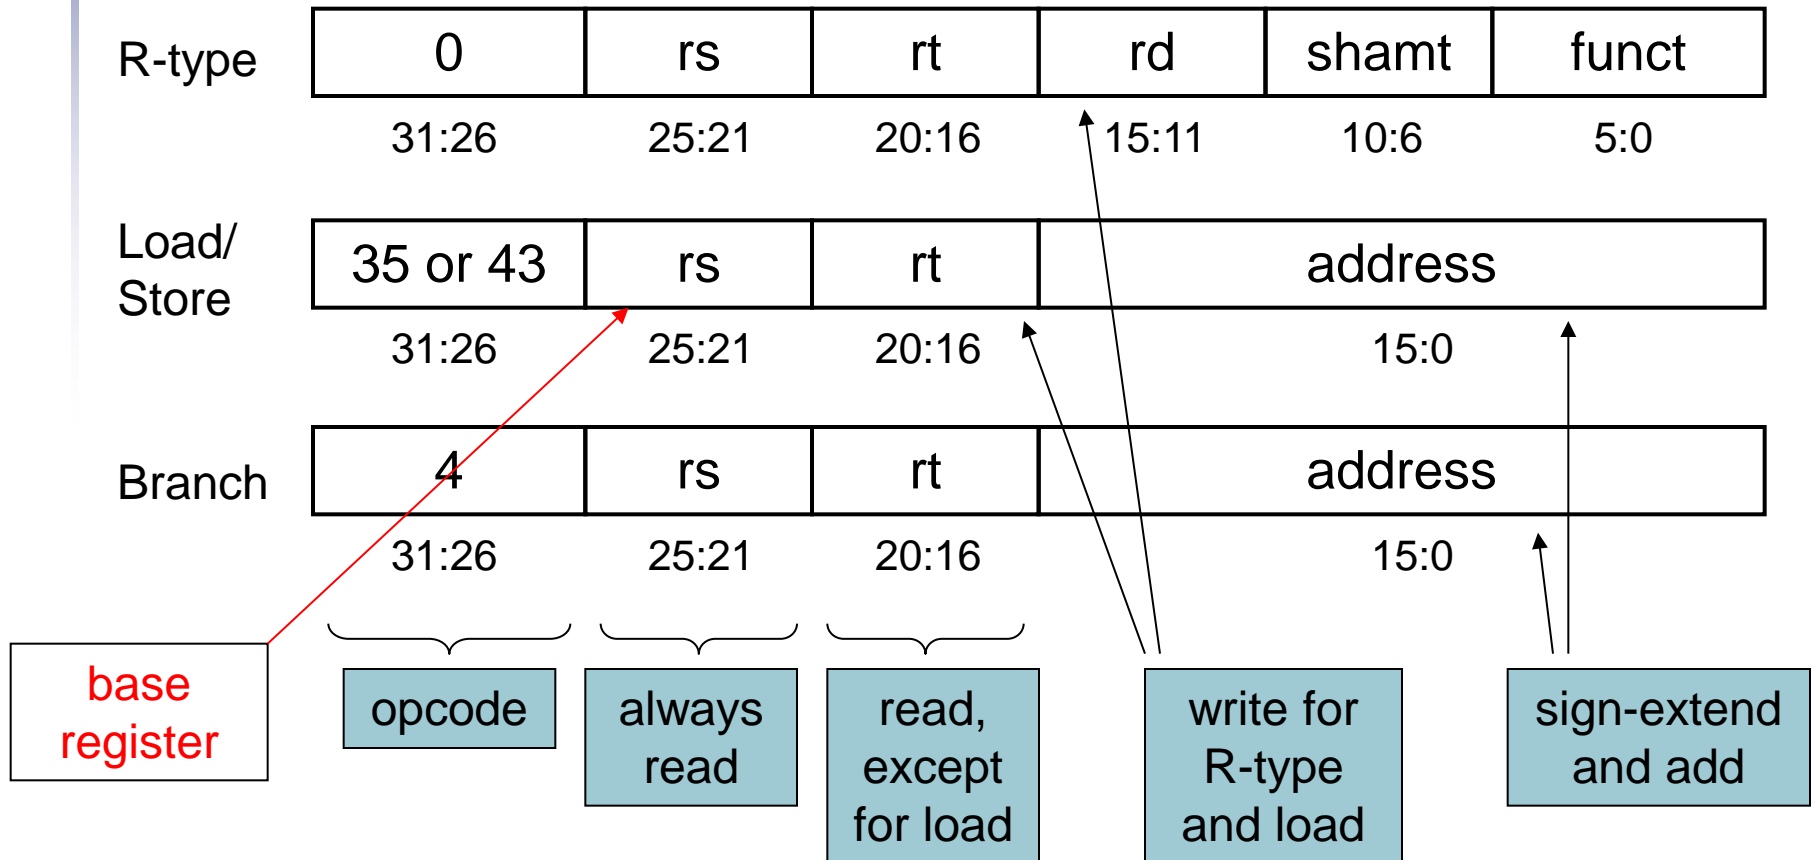
- Assume 2-bit ALUOp derived from opcode
 - Combinational logic derives ALU control

opcode	ALUOp	Operation	funct	ALU function	ALU control
lw	00	load word	XXXXXX	add	0010
sw	00	store word	XXXXXX	add	0010
beq	01	branch equal	XXXXXX	subtract	0110
R-type	10	add	100000	add	0010
		subtract	100010	subtract	0110
		AND	100100	AND	0000
		OR	100101	OR	0001
		set-on-less-than	101010	set-on-less-than	0111

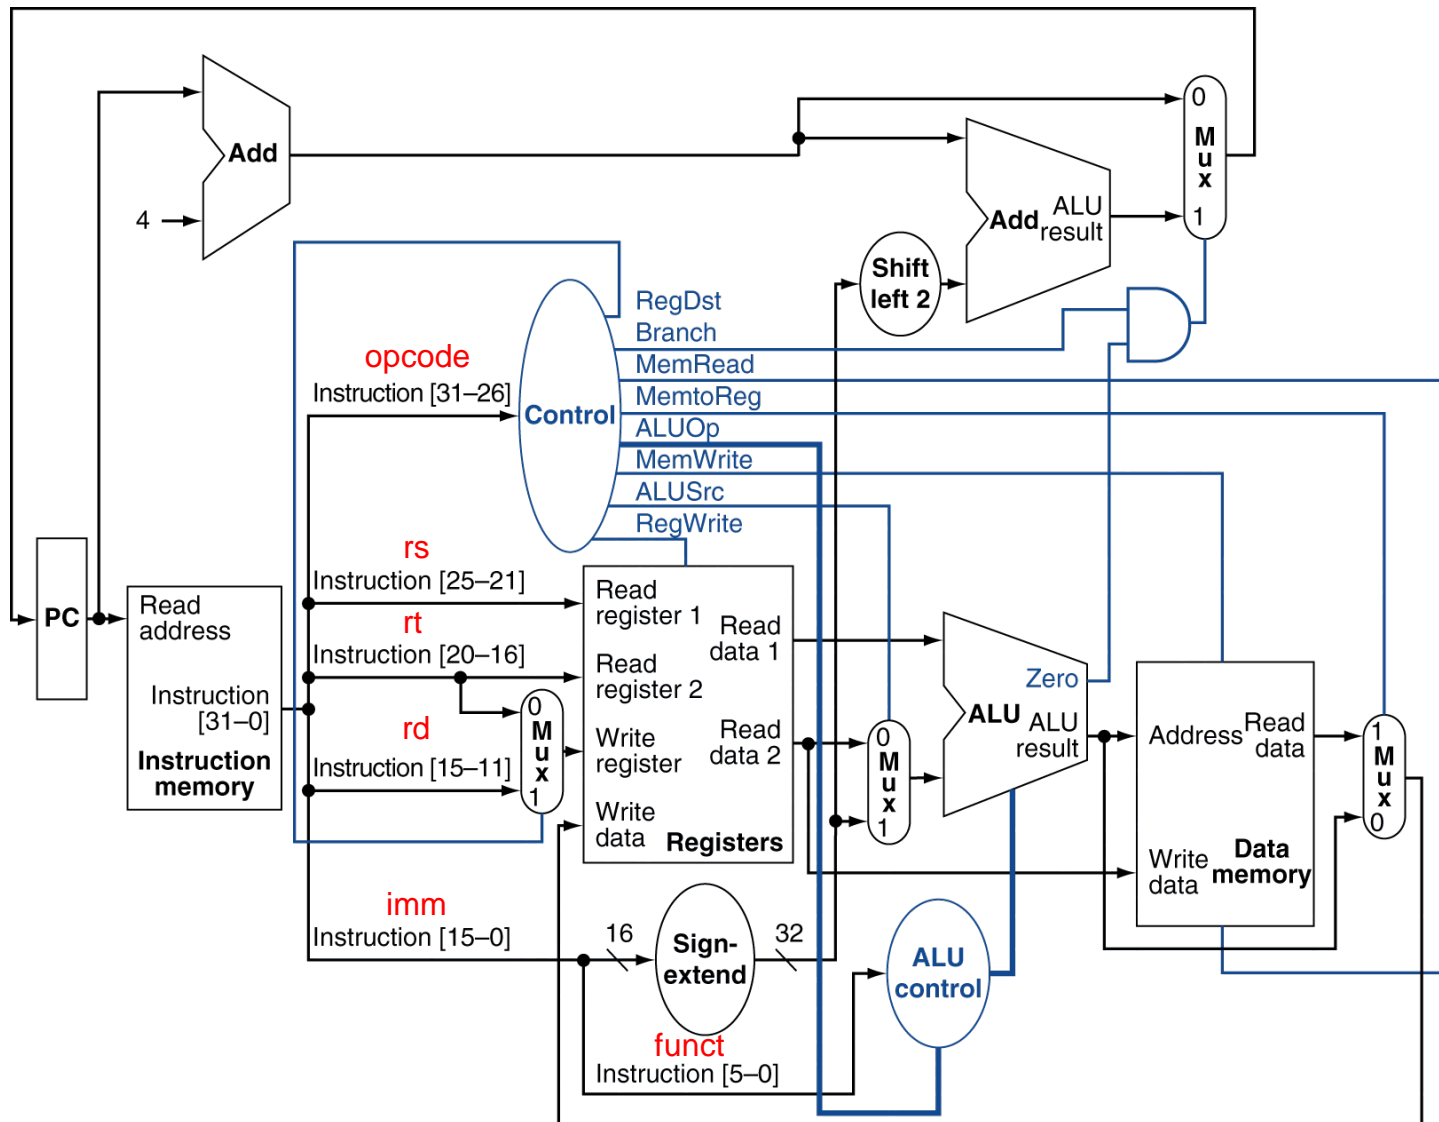
The Main Control Unit



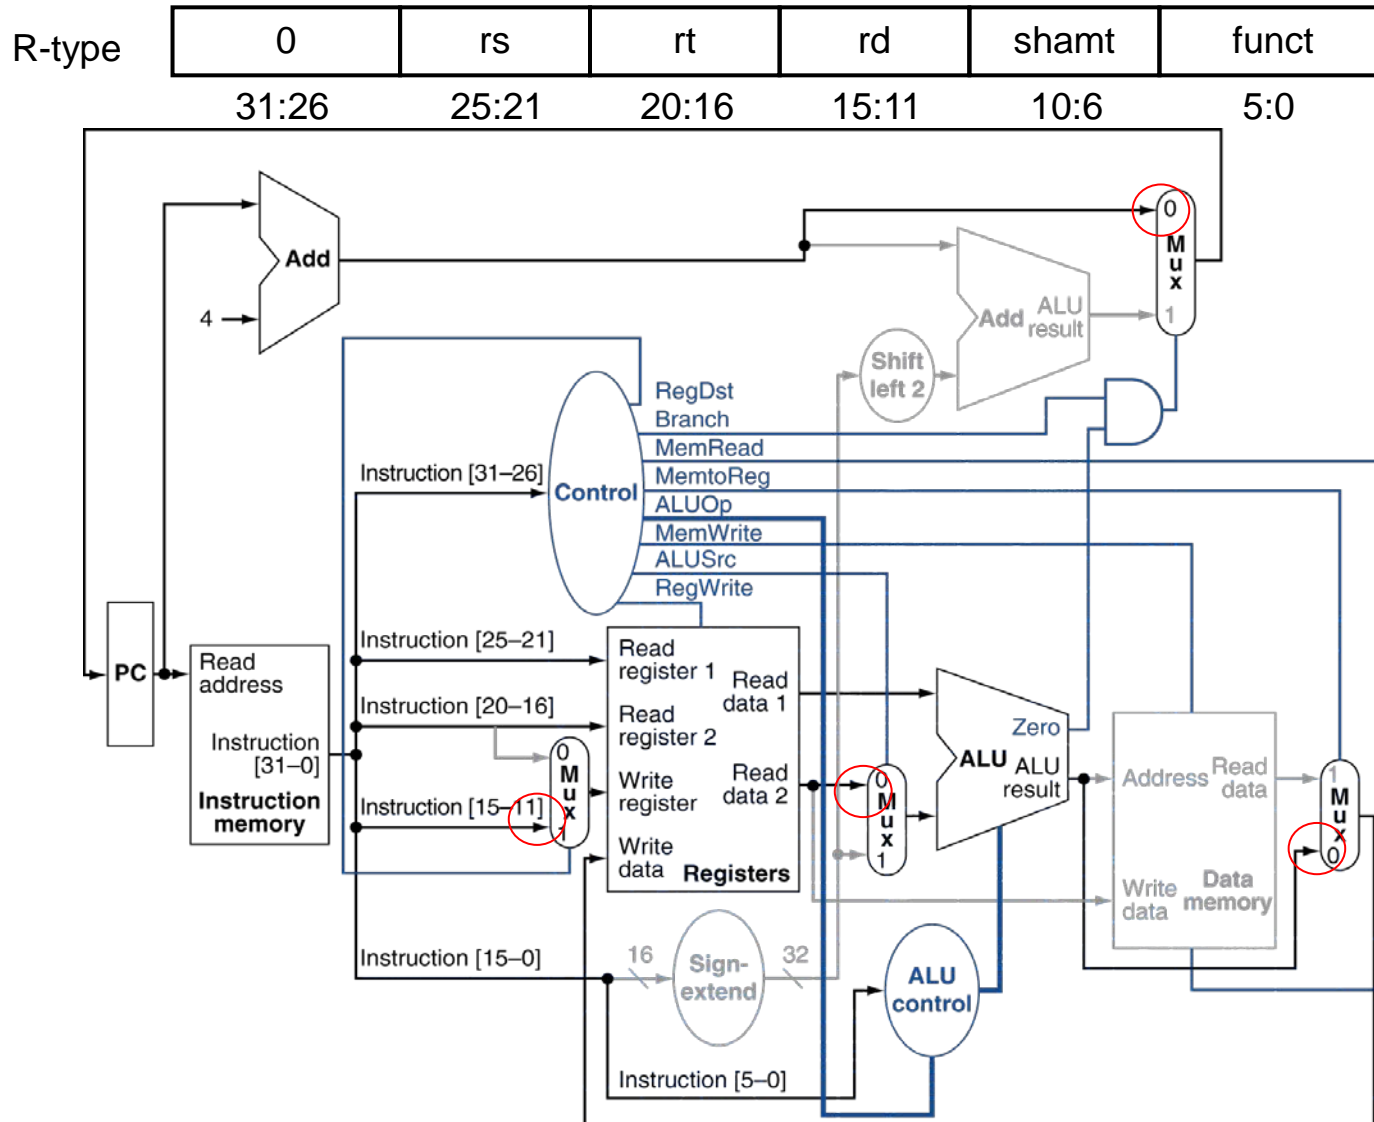
- Control signals derived from instruction



Datapath With Control

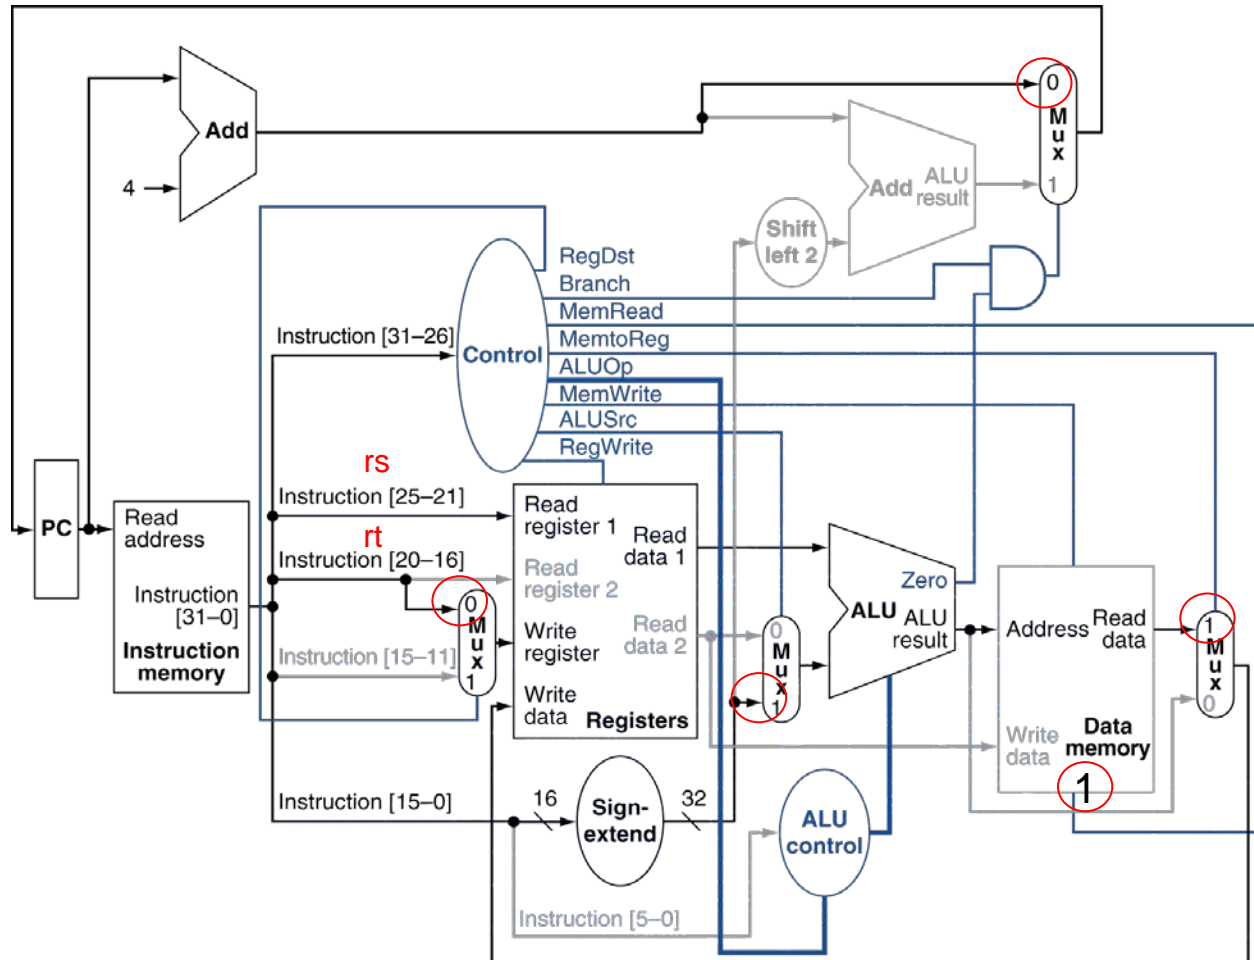


R-Type Instruction



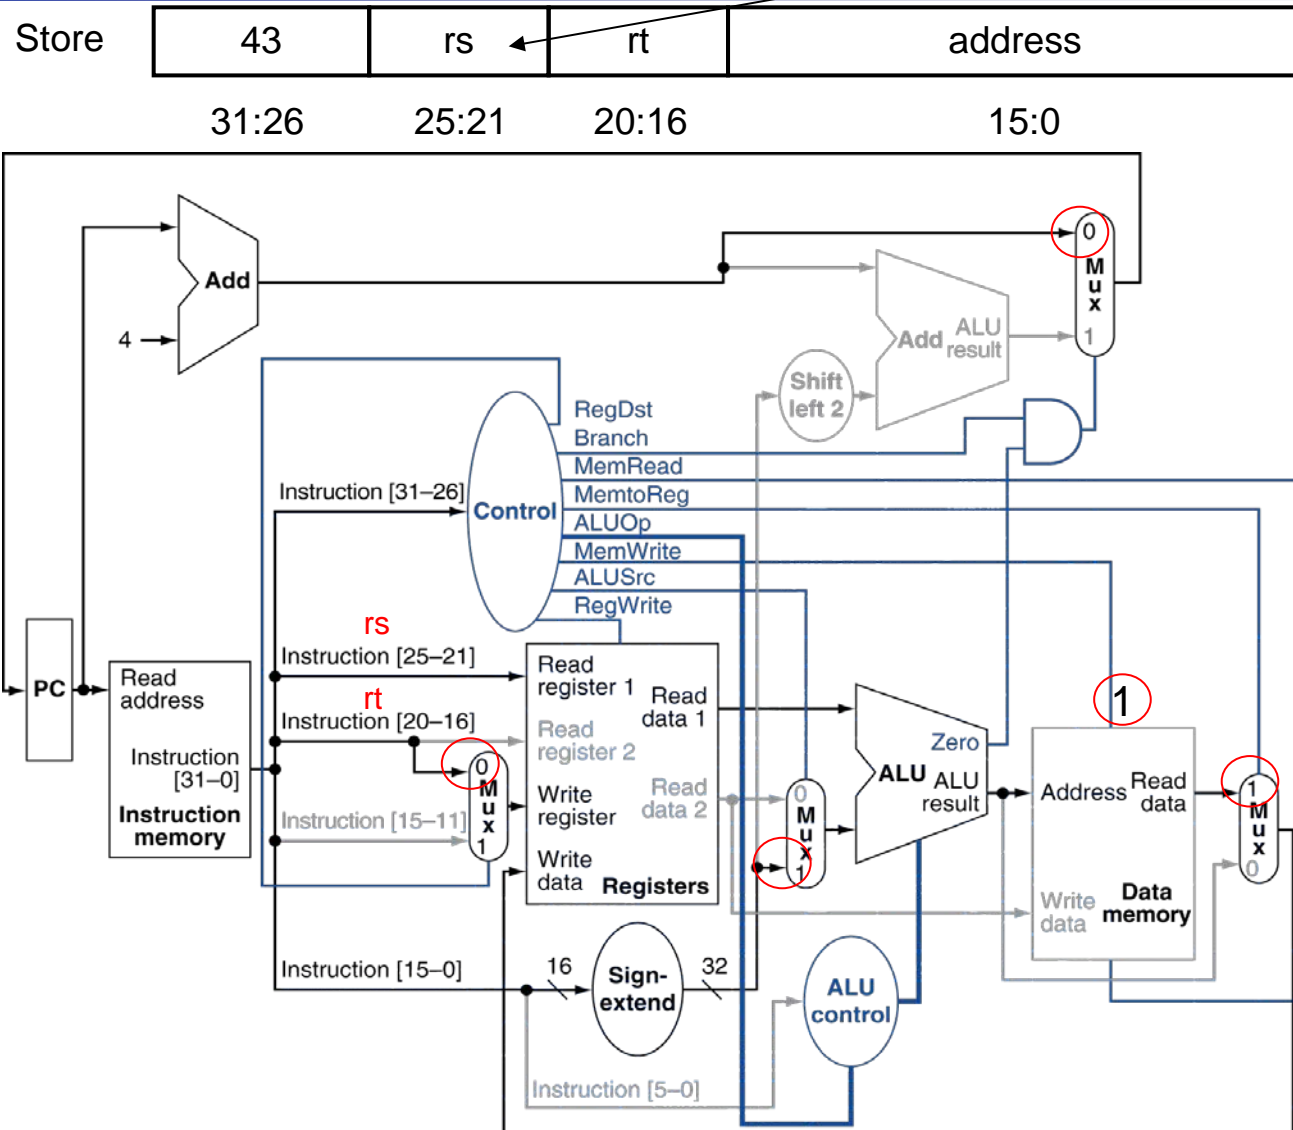
Load Instruction

base register

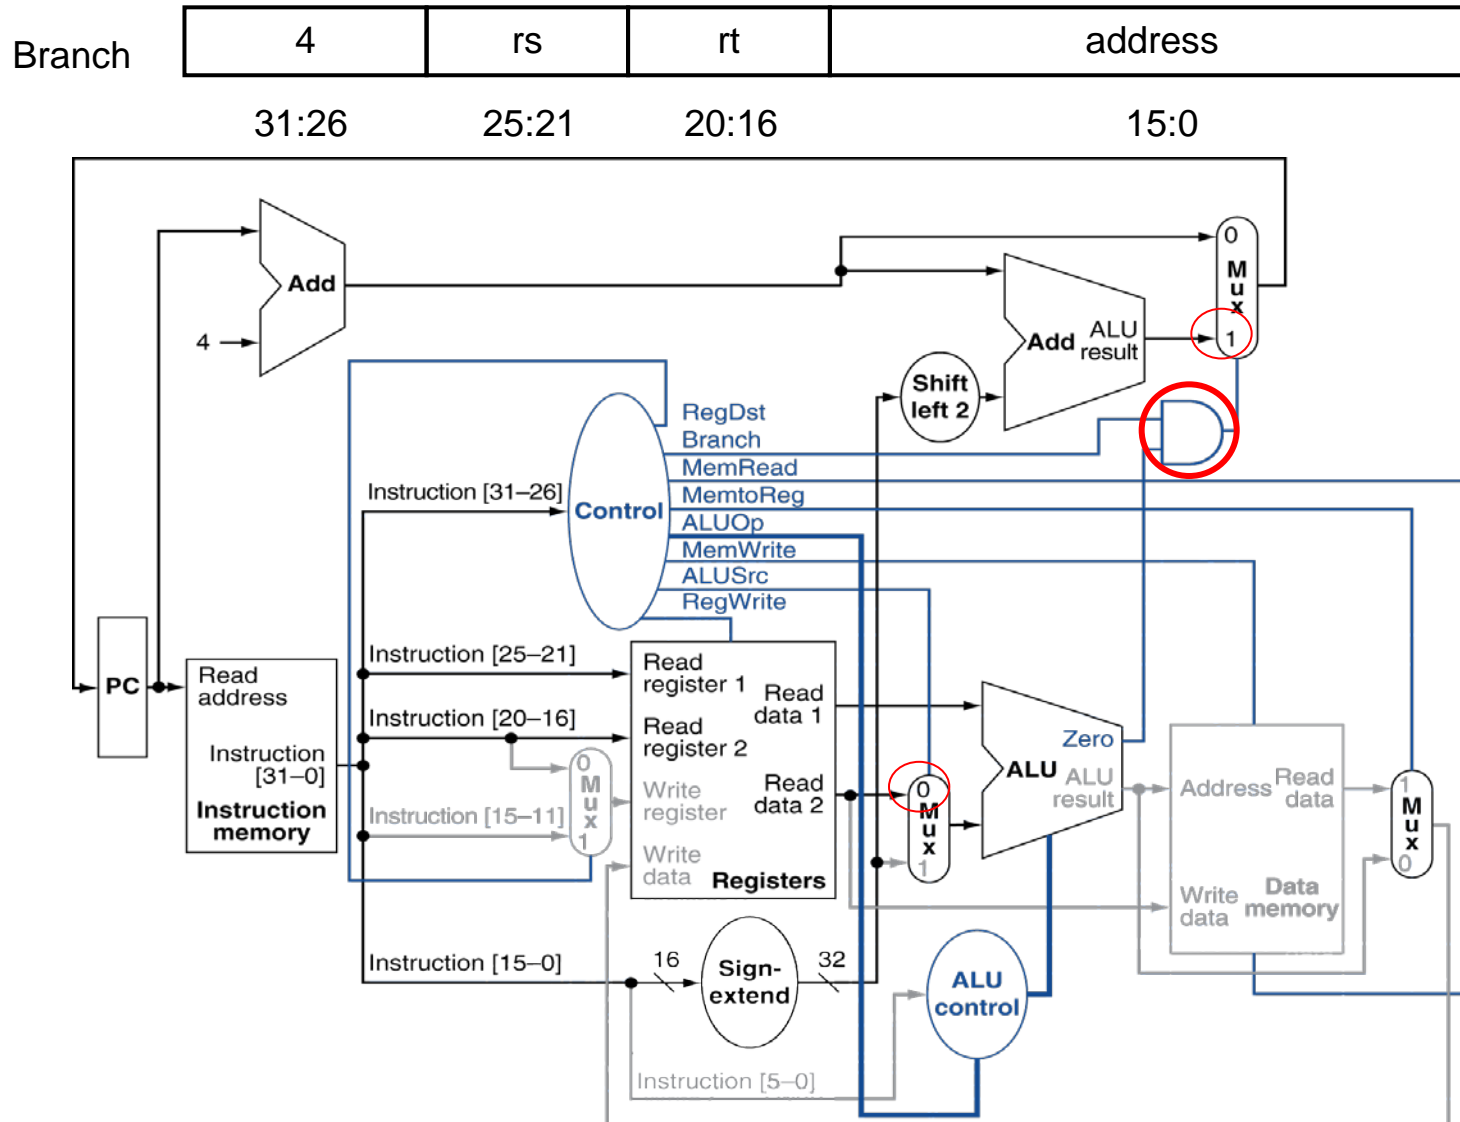


Store Instruction


base register



Branch-on-Equal Instruction

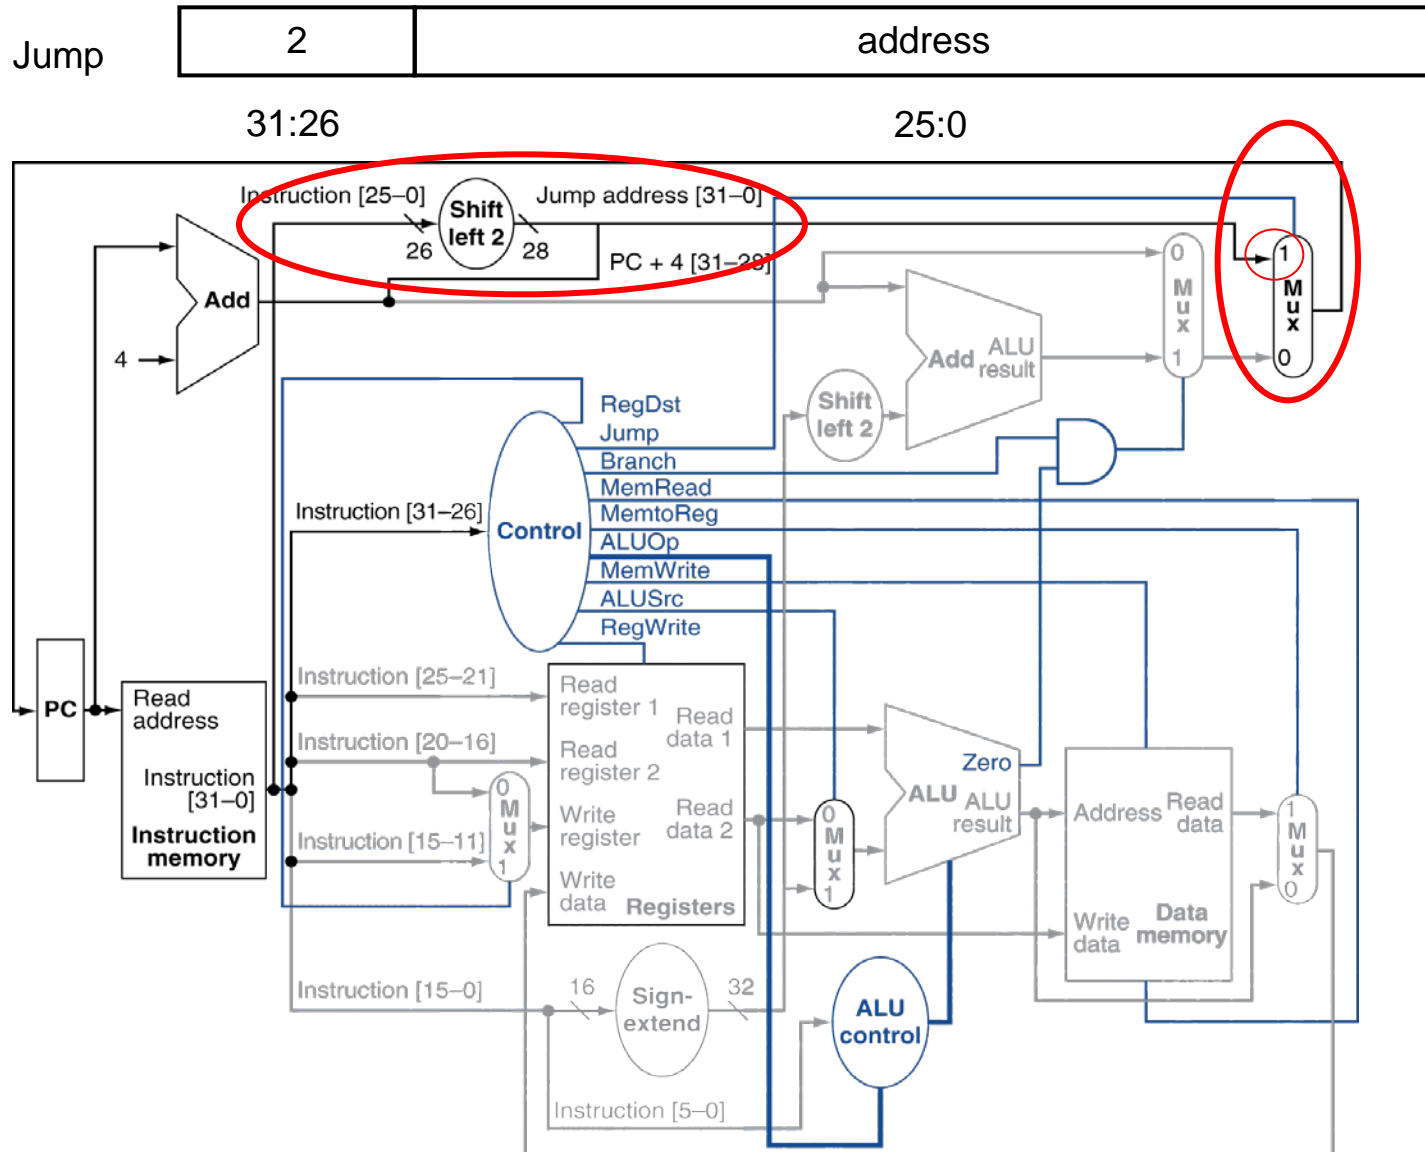


Implementing Jumps



- Jump uses word address
- Update PC with concatenation of
 - Top 4 bits of old PC
 - 26-bit jump address
 - 00
- Need an extra control signal decoded from opcode

Datapath With Jumps Added





Performance Issues

- Longest delay determines clock period
 - Critical path: load instruction
 - **lw**: Instruction memory → register file → ALU
→ data memory → register file
- Not feasible to vary period for different instructions
- Violates design principle
 - Making the common case fast
- We will improve performance by **pipelining**

Pipelining Analogy

- Pipelined laundry: overlapping execution
 - Parallelism improves performance



- Four loads:

- Speedup
 $= 8 / 3.5 = 2.3$

- Non-stop:

- Speedup
 $= 2n / (0.5n + 1.5) \approx 4$
 $= \text{number of stages}$

MIPS Pipeline

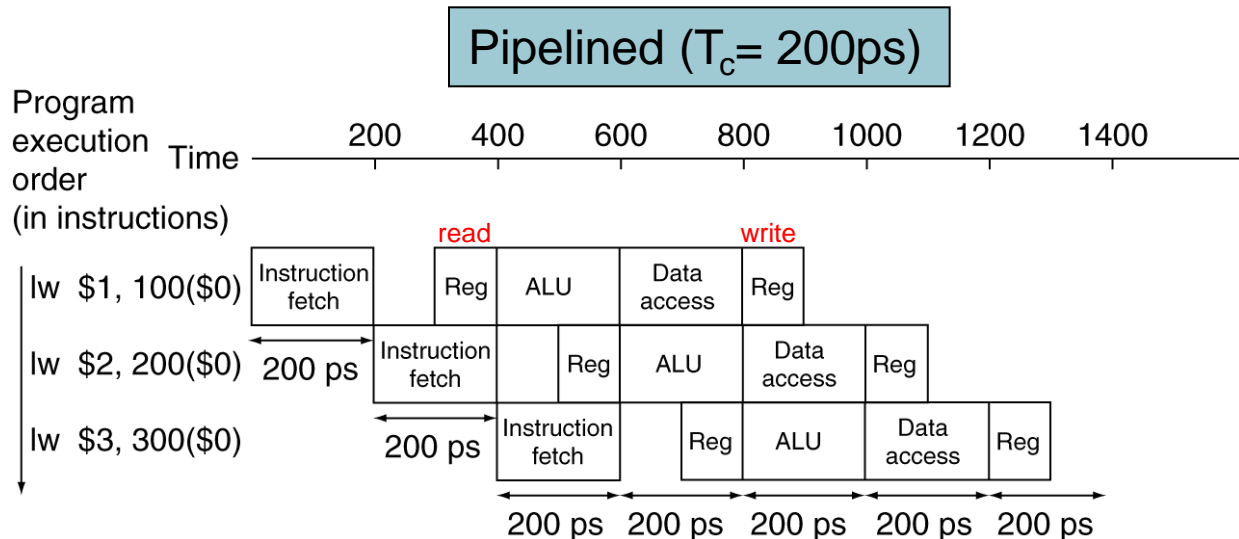
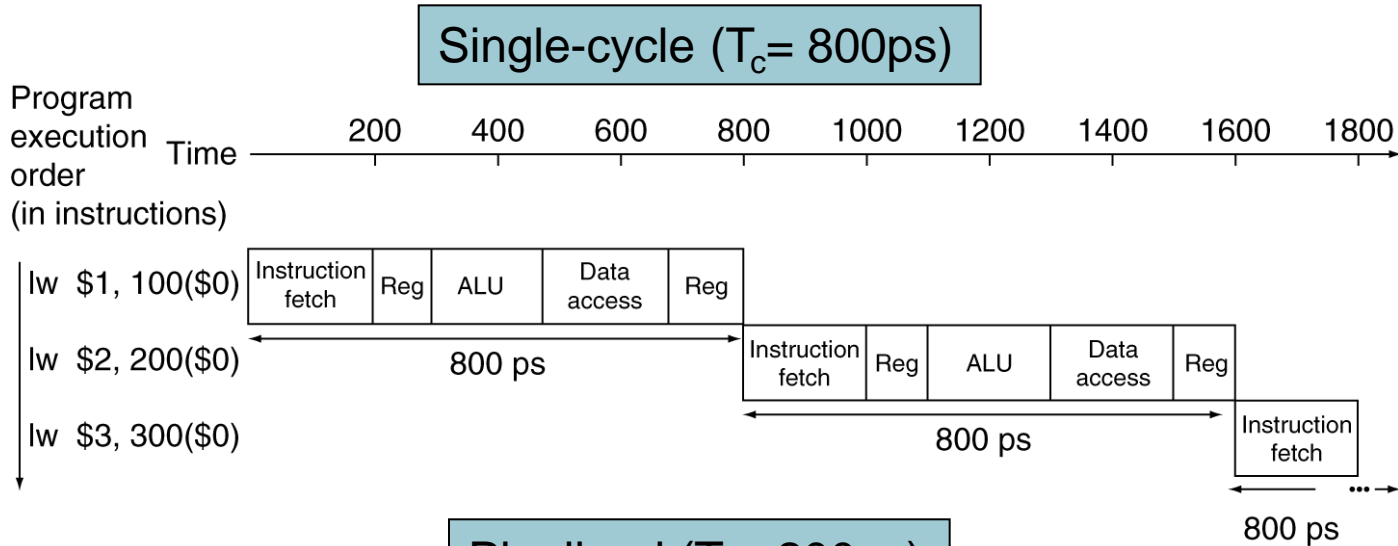
- Five stages, one step per stage
 1. IF: Instruction fetch from memory
 2. ID: Instruction decode & register read
 3. EX: Execute operation or calculate address
 4. MEM: Access memory operand
 5. WB: Write result back to register

Pipeline Performance

- Assume time for stages is
 - 100ps for register read or write
 - 200ps for other stages
- Compare pipelined datapath with single-cycle datapath

Instr	Instr fetch	Register read	ALU op	Memory access	Register write	Total time
lw	200ps	100 ps	200ps	200ps	100 ps	800ps
sw	200ps	100 ps	200ps	200ps		700ps
R-format	200ps	100 ps	200ps		100 ps	600ps
beq	200ps	100 ps	200ps			500ps

Pipeline Performance



Pipeline Speedup


- If all stages are balanced
 - i.e., all take the same time
 - Time between instructions_{pipelined}
= $\frac{\text{Time between instructions}_{\text{nonpipelined}}}{\text{Number of stages}}$
- If not balanced, speedup is less
- Speedup due to increased throughput
 - Latency (time for each instruction) does not decrease

Pipelining and ISA Design


- MIPS ISA designed for pipelining
 - All instructions are 32-bits
 - Easier to fetch and decode in one cycle
 - c.f. x86: 1- to 17-byte instructions
 - Few and regular instruction formats
 - Can decode and read registers in one step
 - Load/store addressing
 - Can calculate address in 3rd stage, access memory in 4th stage
 - **Alignment** of memory operands
 - Memory access takes only one cycle

Other Pipeline Structures

- What about the (slow) multiply operation?
 - Make the clock twice as slow or ...
 - let it take two cycles (since it doesn't use the DM stage)




- What if the data memory access is twice as slow as the instruction memory?
 - make the clock twice as slow or ...
 - let data memory access take two cycles (and keep the same clock rate)




Sample Pipeline Alternatives

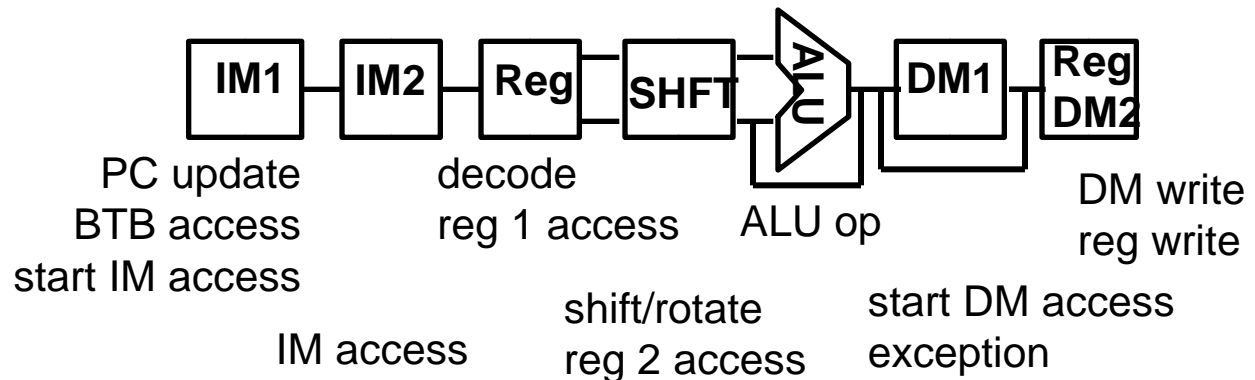
■ ARM7



■ StrongARM-1

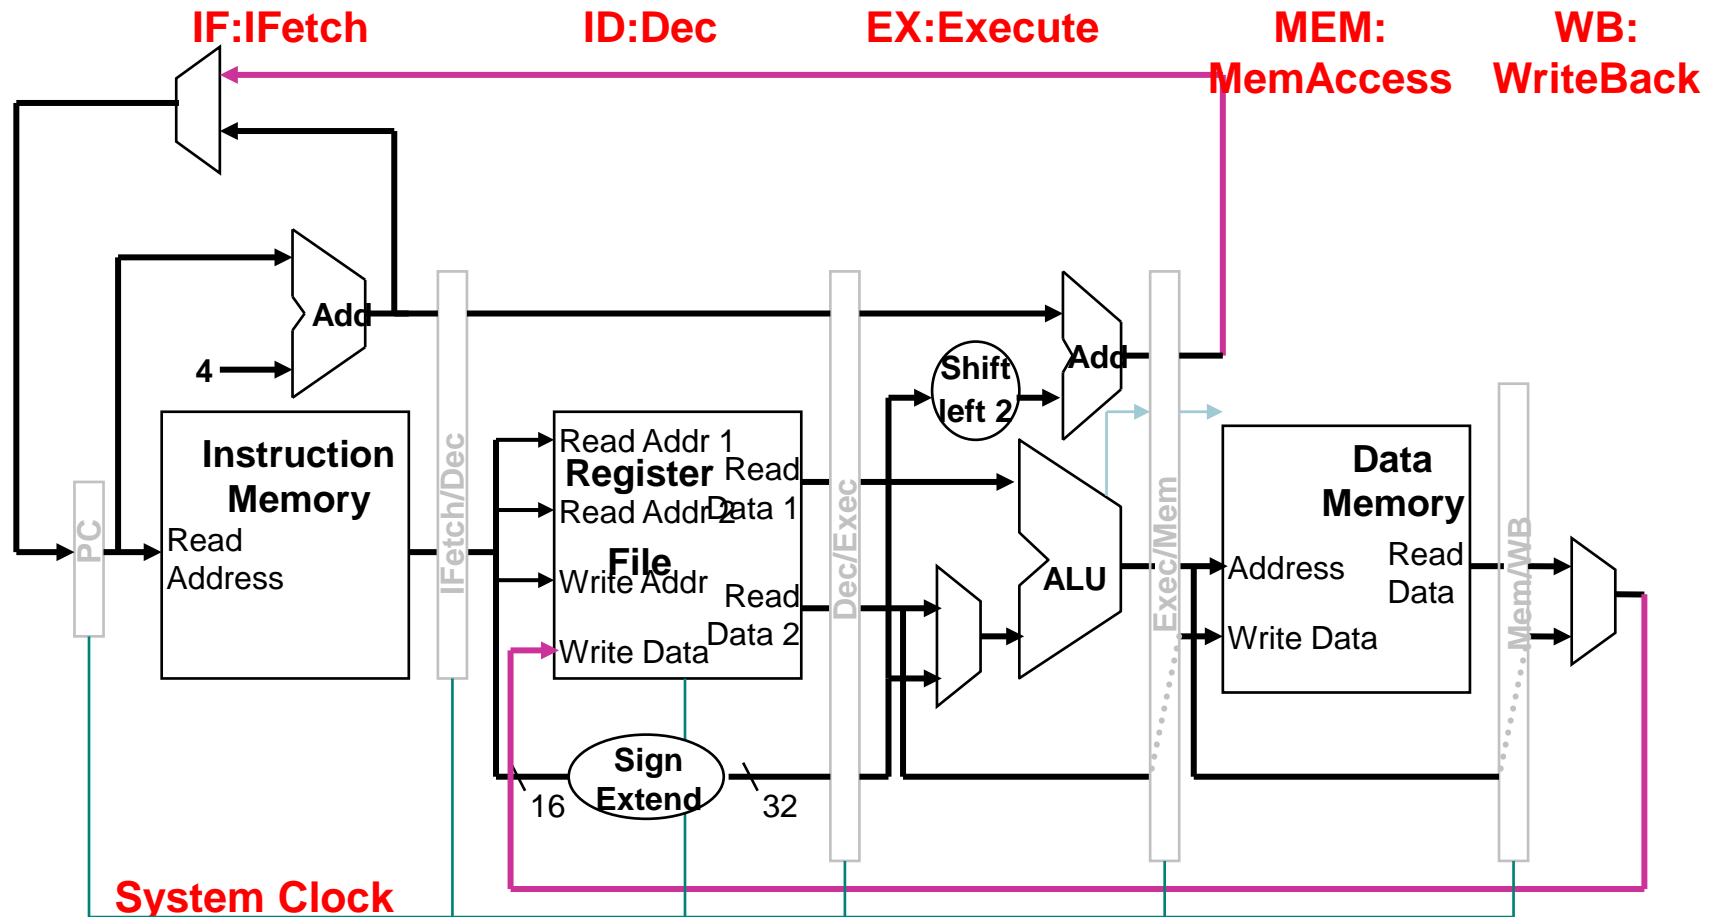


■ XScale



MIPS Pipeline Datapath

- What do we need to add/modify in our MIPS datapath?
 - State registers between each pipeline stage to **isolate** them



What is wrong here?

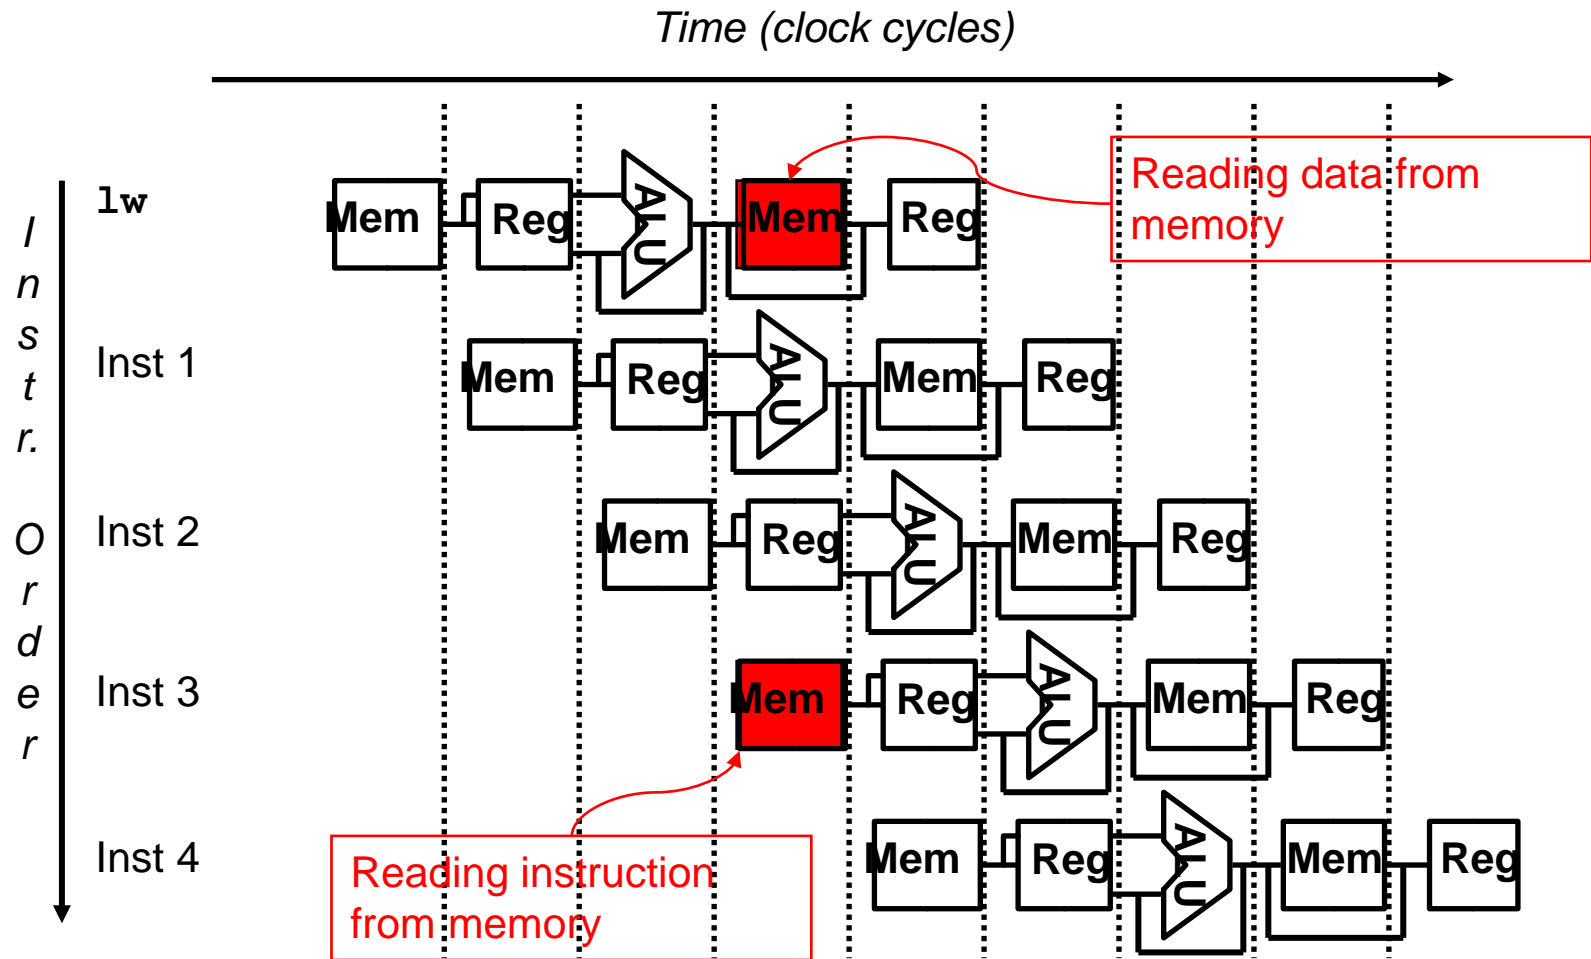
(Pipeline) Hazards

- Situations that prevent starting the next instruction in the next cycle
- **Structure** hazards
 - A required resource is busy
- **Data** hazard
 - Need to wait for previous instruction to complete its data read/write
- **Control** (branch) hazard
 - Deciding on control action depends on previous instruction

Structure Hazards

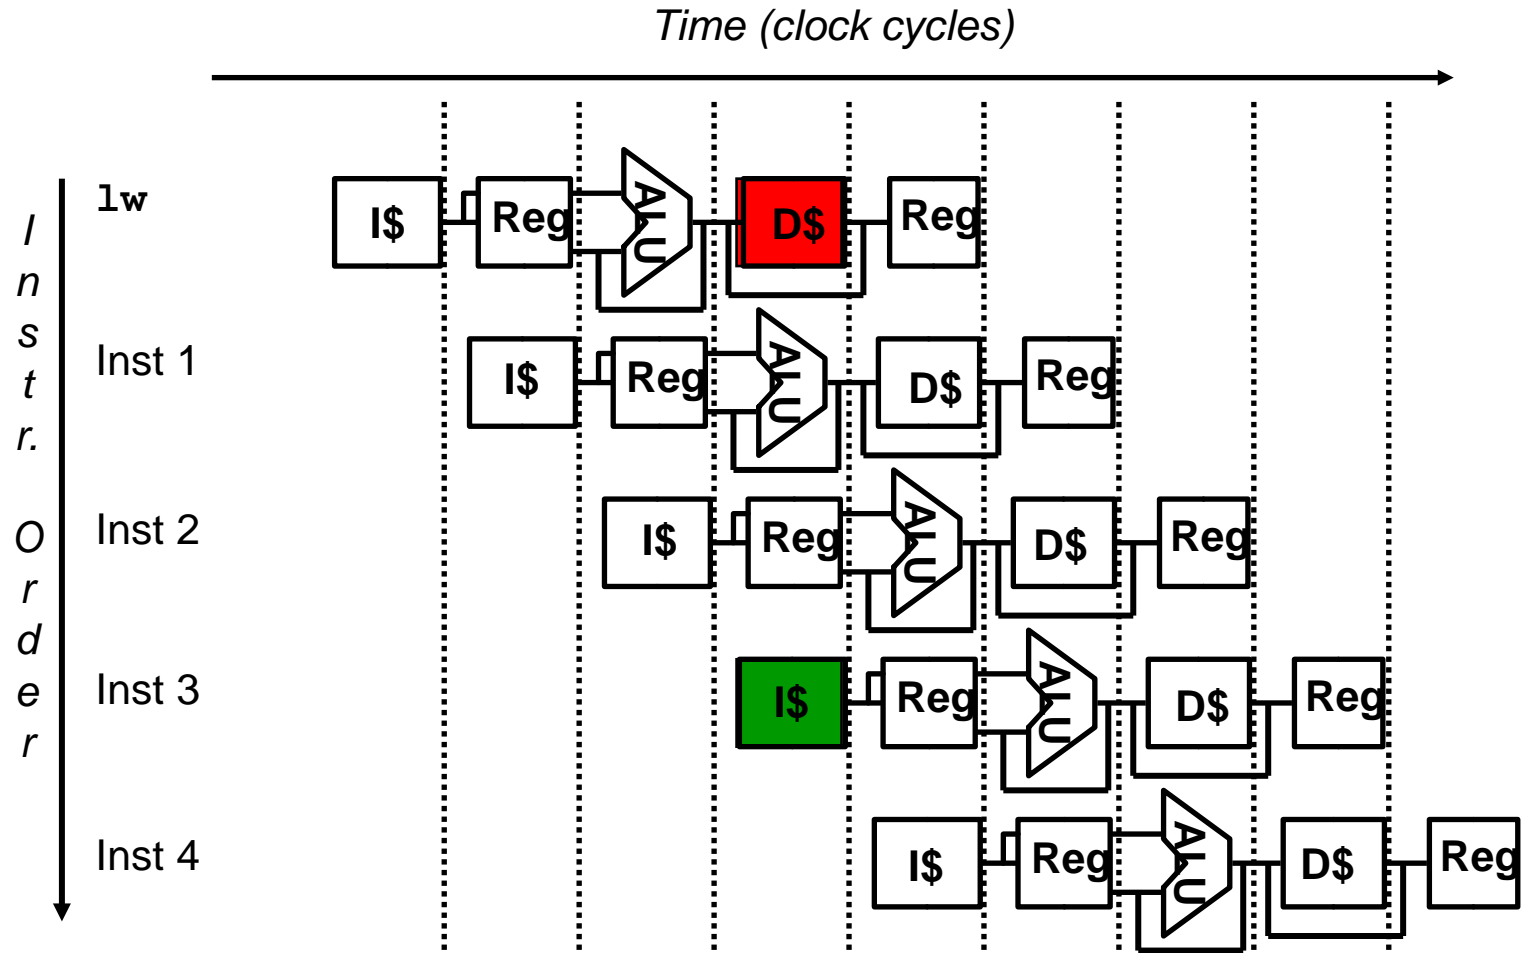
- Conflict for use of a resource
- In MIPS pipeline with a single memory
 - Load/store requires data access
 - Instruction fetch would have to *stall* for that cycle
 - Would cause a pipeline “bubble”
- Hence, pipelined datapaths require separate instruction/data memories
 - Or separate instruction/data **caches**

A Single Memory ?!



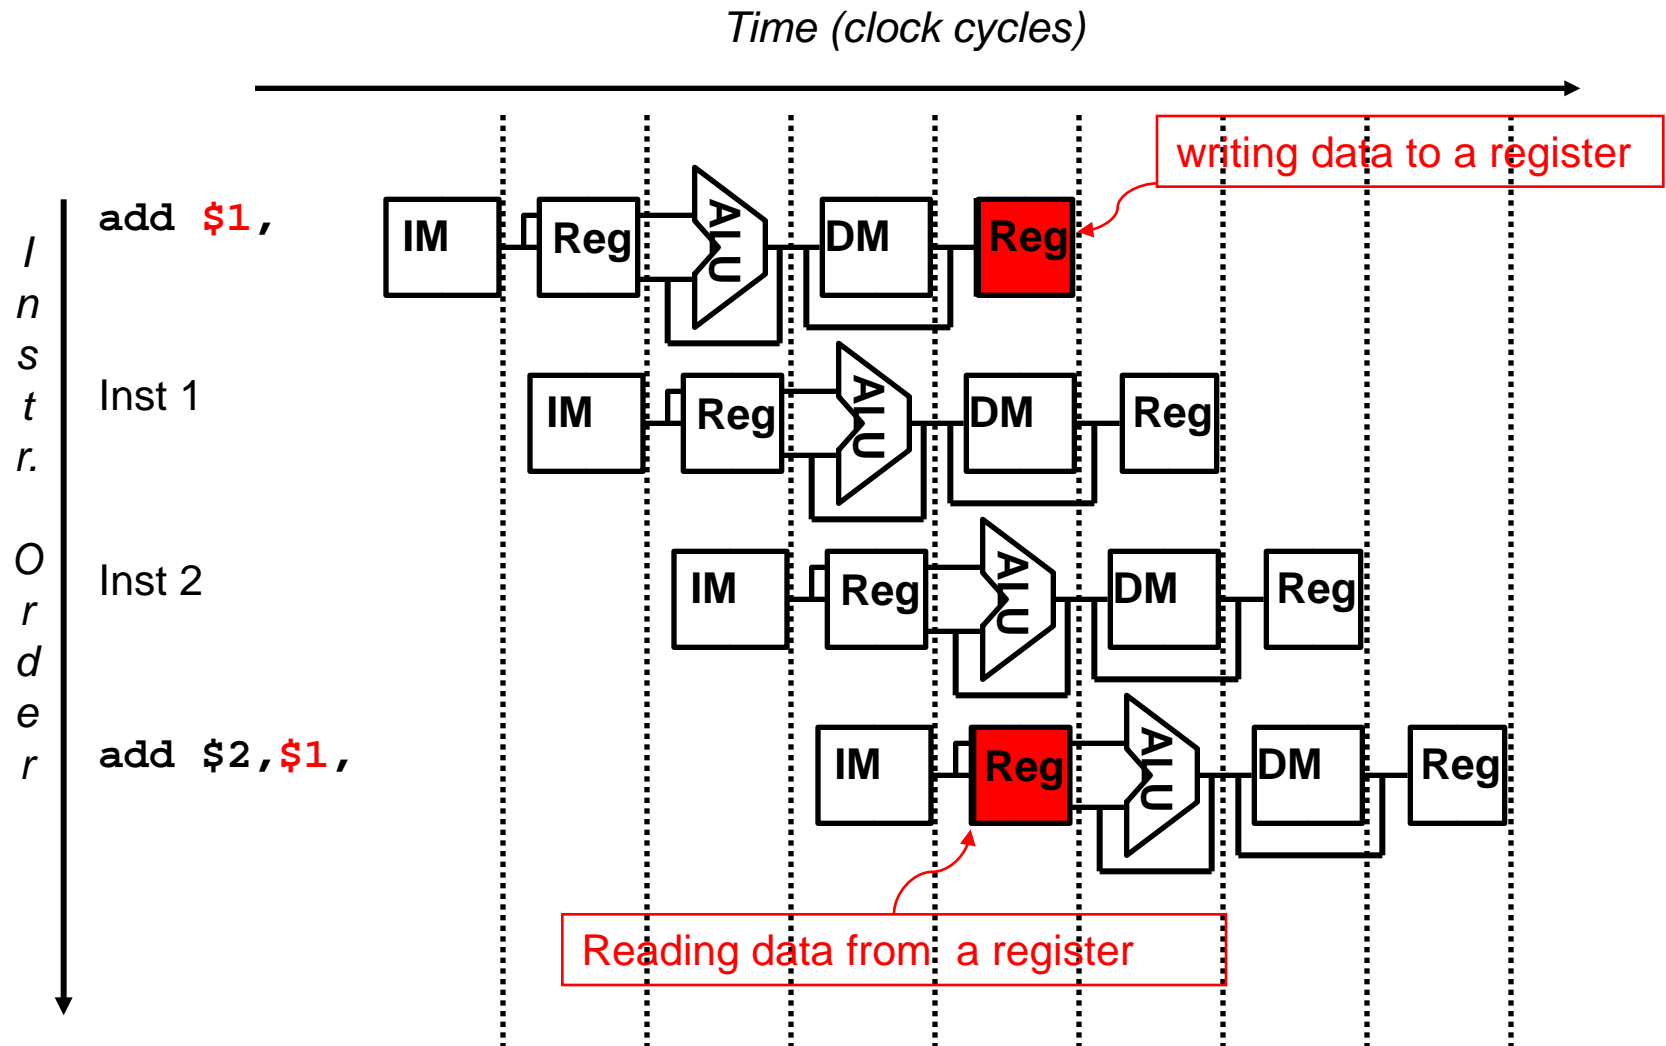
- A structure hazard!

Dual Memory: I\$, D\$!



- Fix with separate instruction and data memories (I\$ and D\$)

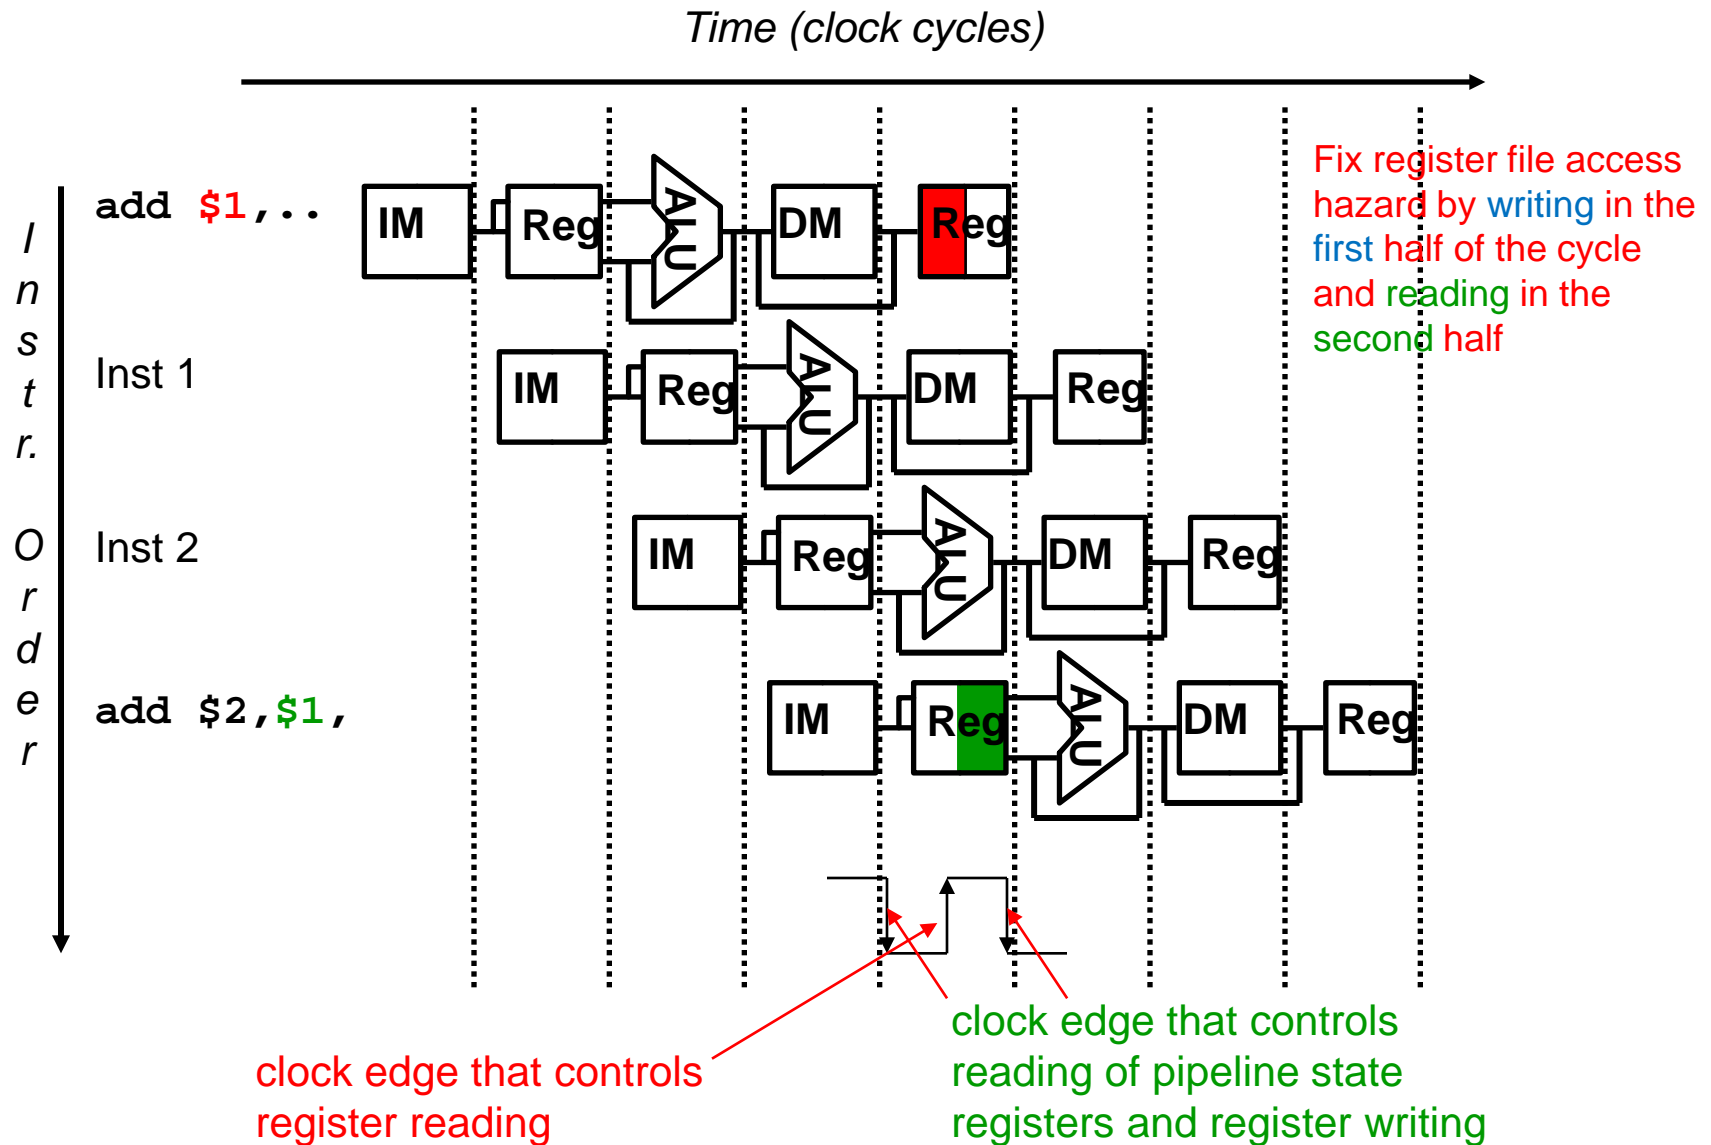
How About Register File Access?



- Another structure hazard!

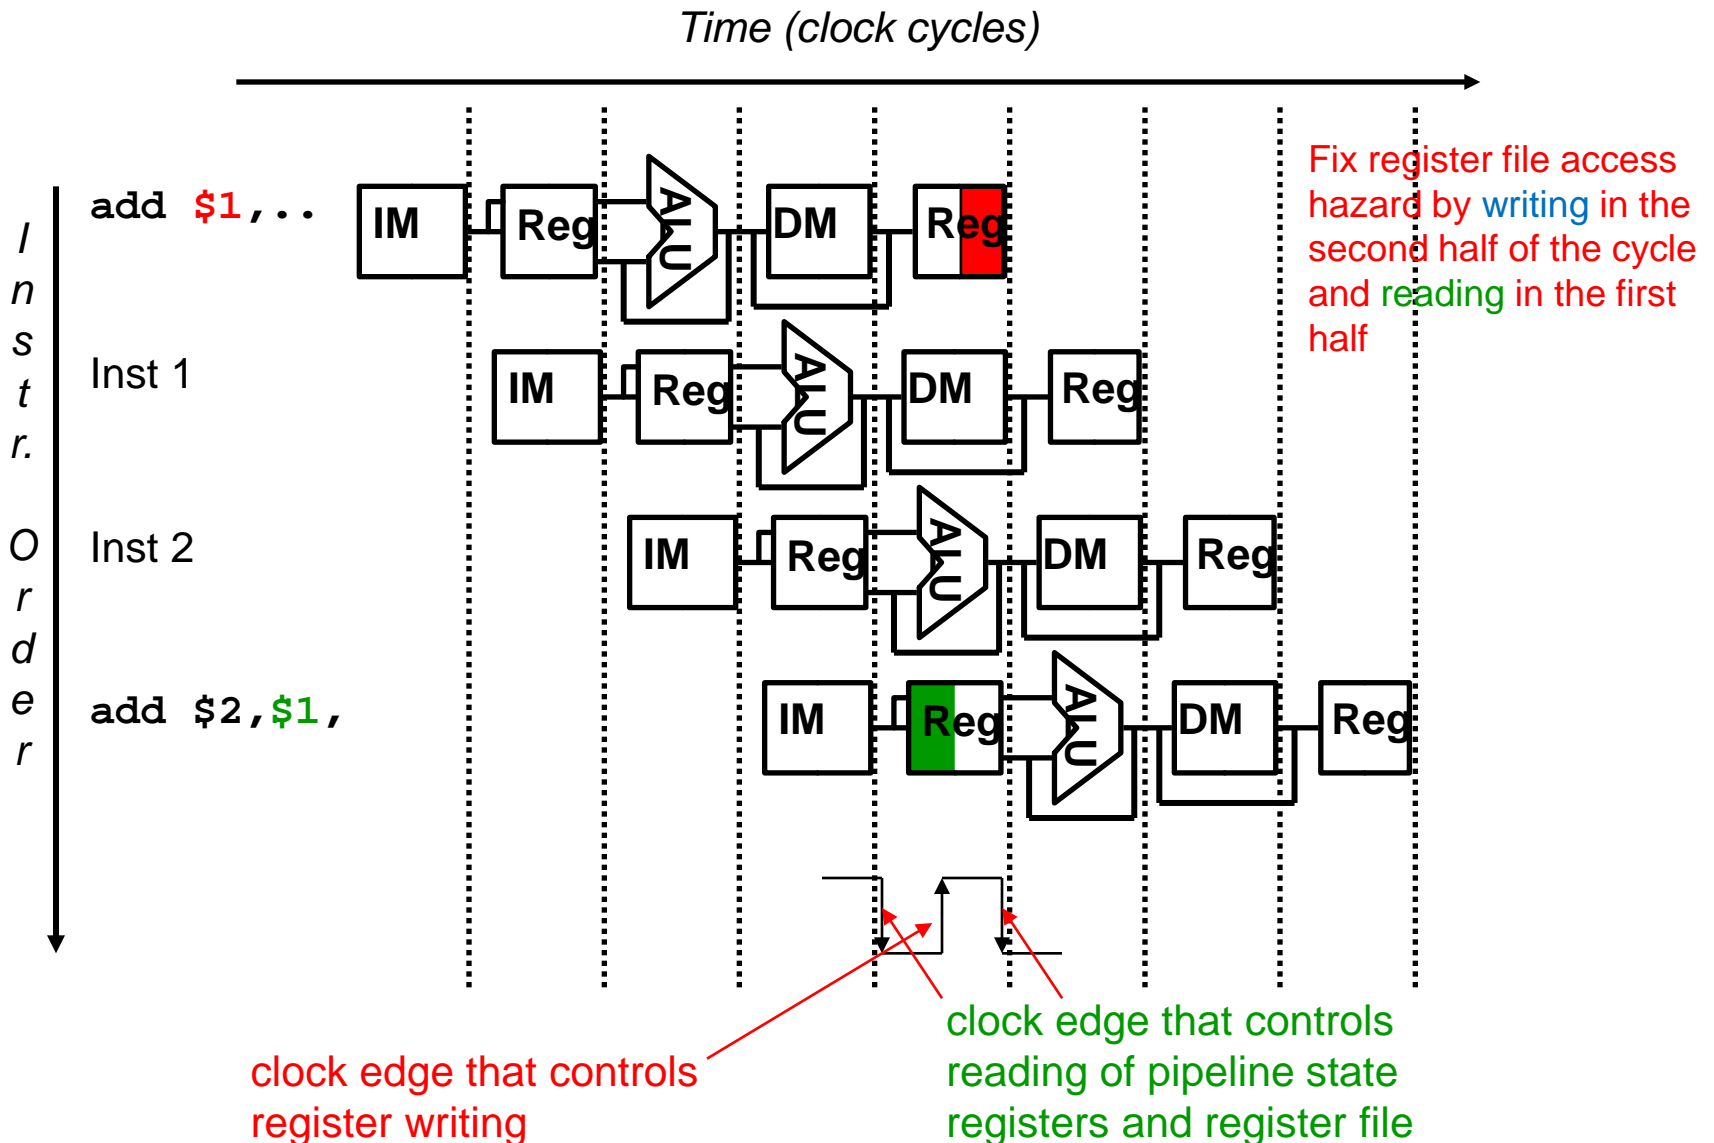
A solution!

Is it OK?



Another solution!

Is it OK?

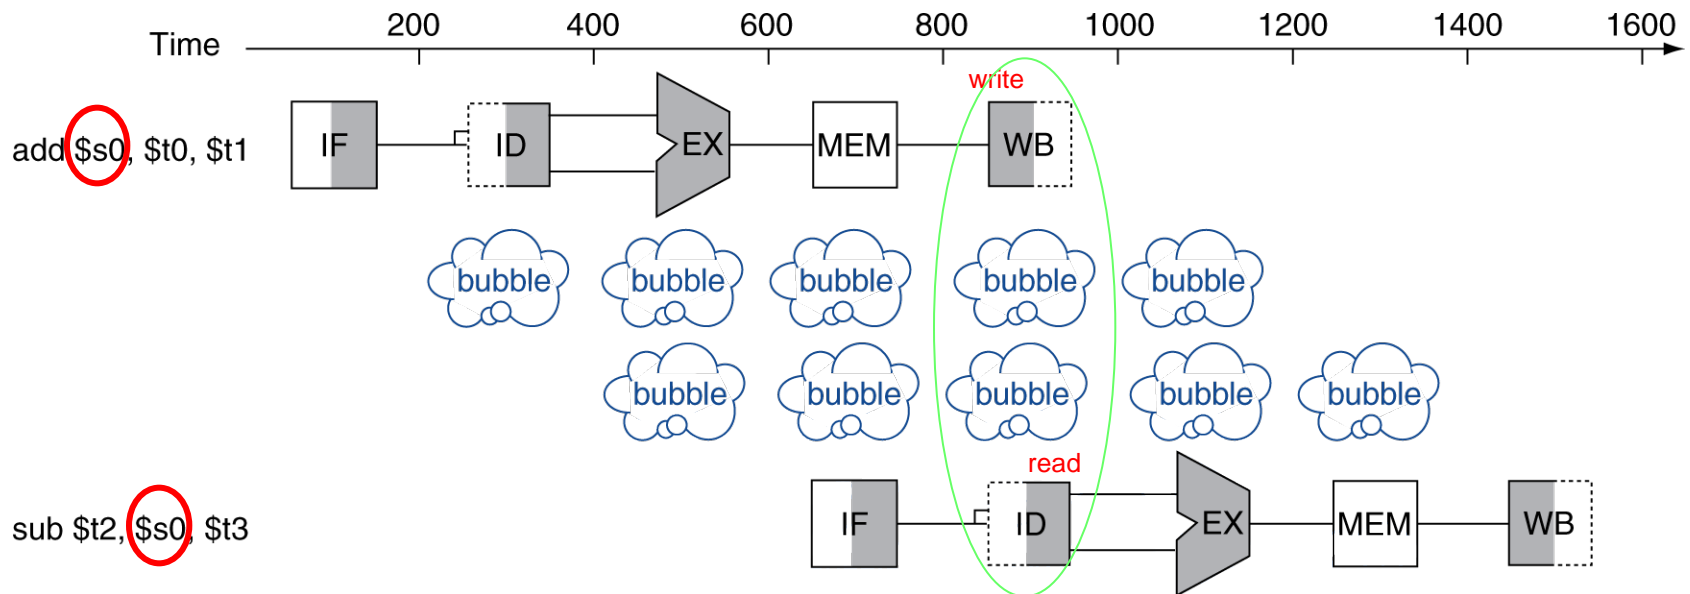


Data Hazards

- An instruction depends on completion of data access by a previous instruction

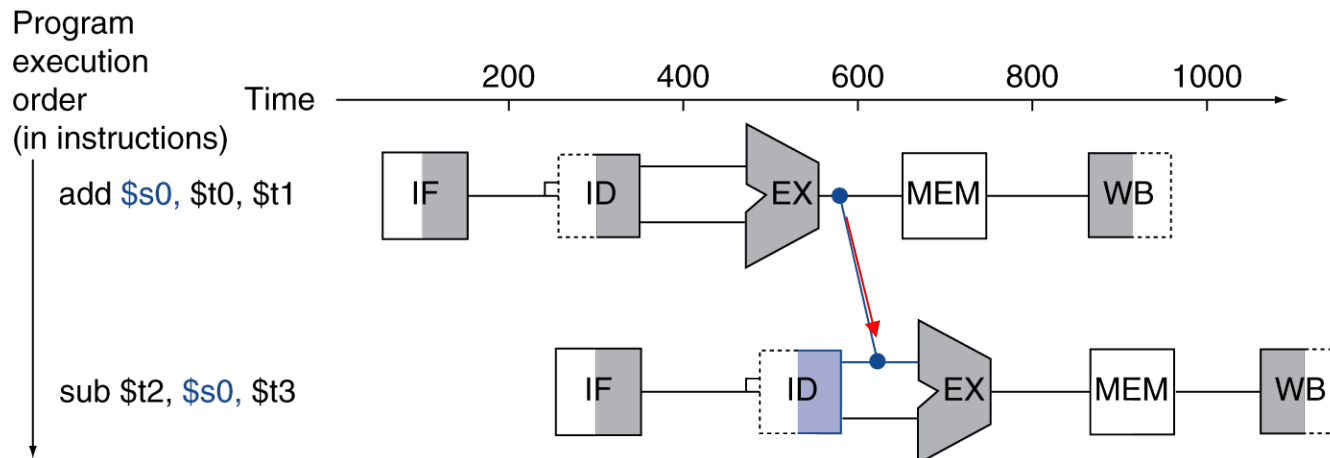
- add **\$s0**, \$t0, \$t1

- sub \$t2, **\$s0**, \$t3



Forwarding (aka Bypassing)

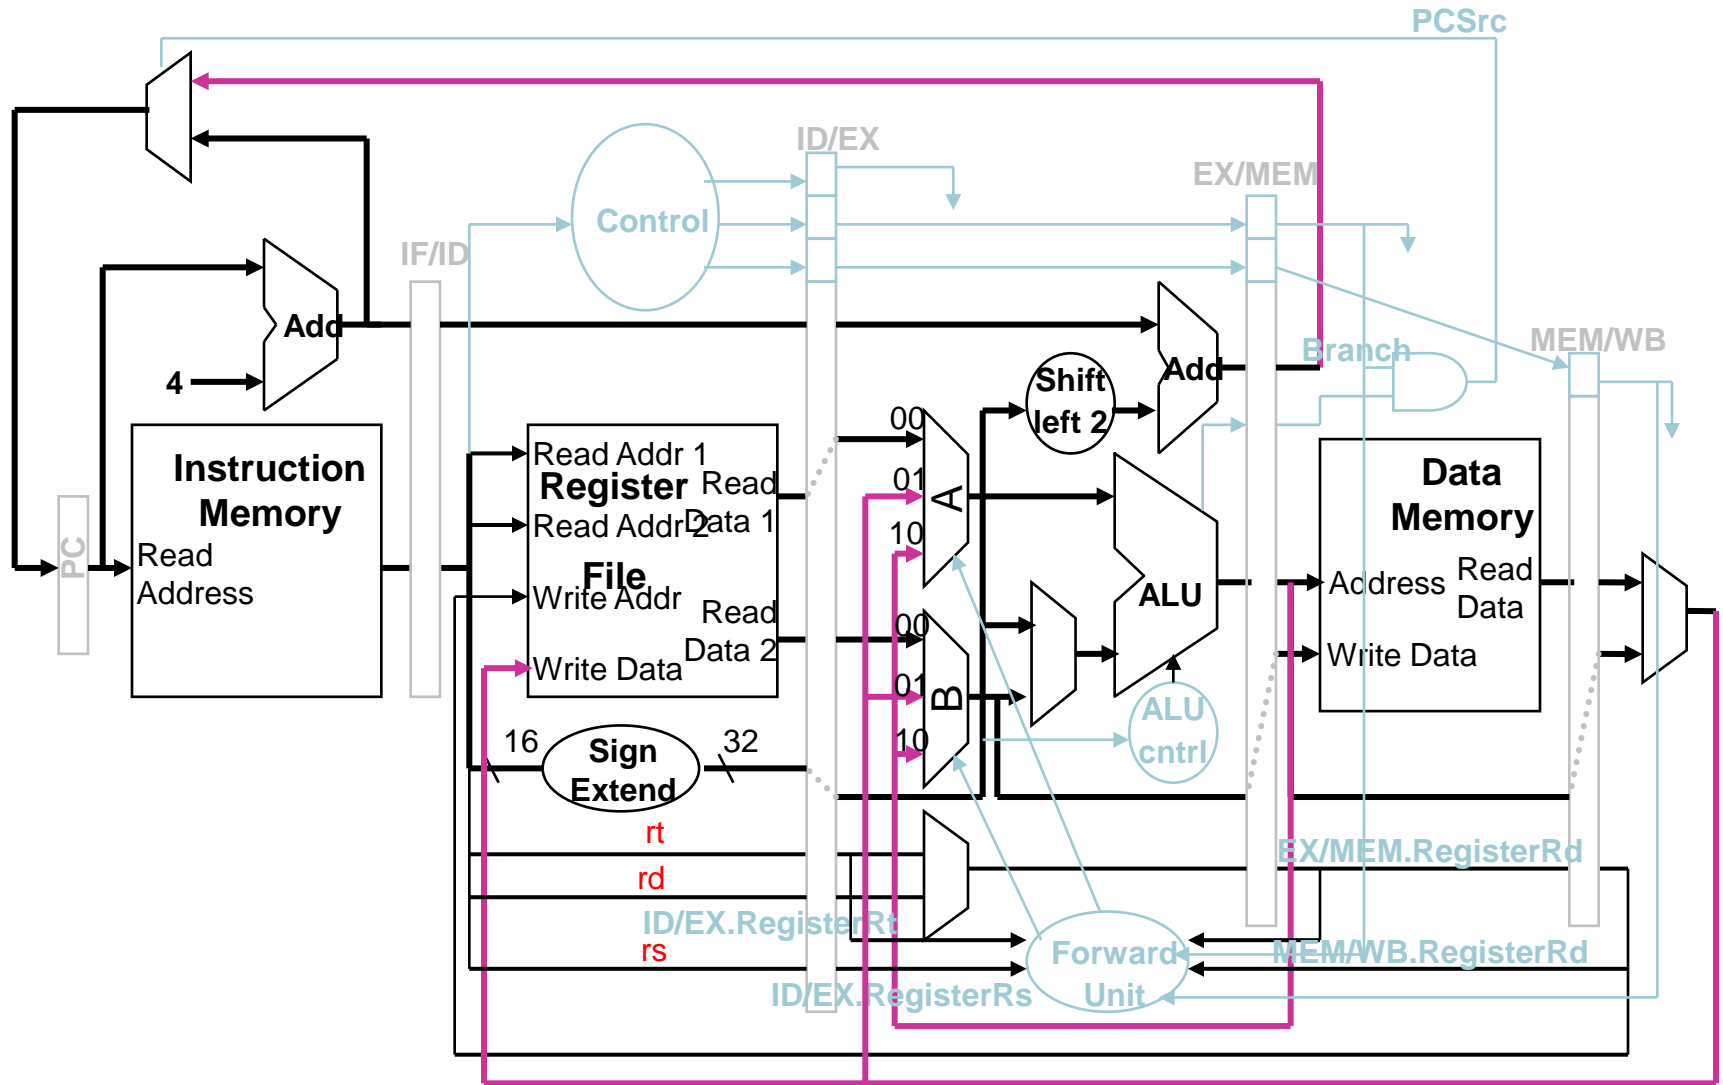
- Use result when it is computed
 - Don't wait for it to be stored in a register
 - Requires extra connections in the datapath



Data Forwarding (Bypassing)

- Take the result from the earliest point that it exists in **any** of the pipeline state registers and forward it to the functional units (e.g., the ALU) that need it that cycle
- For ALU functional unit: the inputs can come from **any** pipeline register rather than just from **ID/EX** by
 - adding multiplexors to the inputs of the ALU
 - connecting the Rd write data in EX/MEM or MEM/WB to either (or both) of the EX's stage Rs and Rt ALU mux inputs
 - adding the proper control hardware to control the new muxes
- Other functional units may need similar forwarding logic (e.g., the DM)
- With forwarding can achieve a CPI of 1 even in the presence of data dependencies

Forwarding Hardware



Forwarding Control Conditions

□ EX hazard:

```
if (EX/MEM.RegWrite
and (EX/MEM.RegisterRd != 0)
and (EX/MEM.RegisterRd = ID/EX.RegisterRs))
    ForwardA = 10
if (EX/MEM.RegWrite
and (EX/MEM.RegisterRd != 0)
and (EX/MEM.RegisterRd = ID/EX.RegisterRt))
    ForwardB = 10
```

Forwards the
result from the
previous instr.
to either input
of the ALU

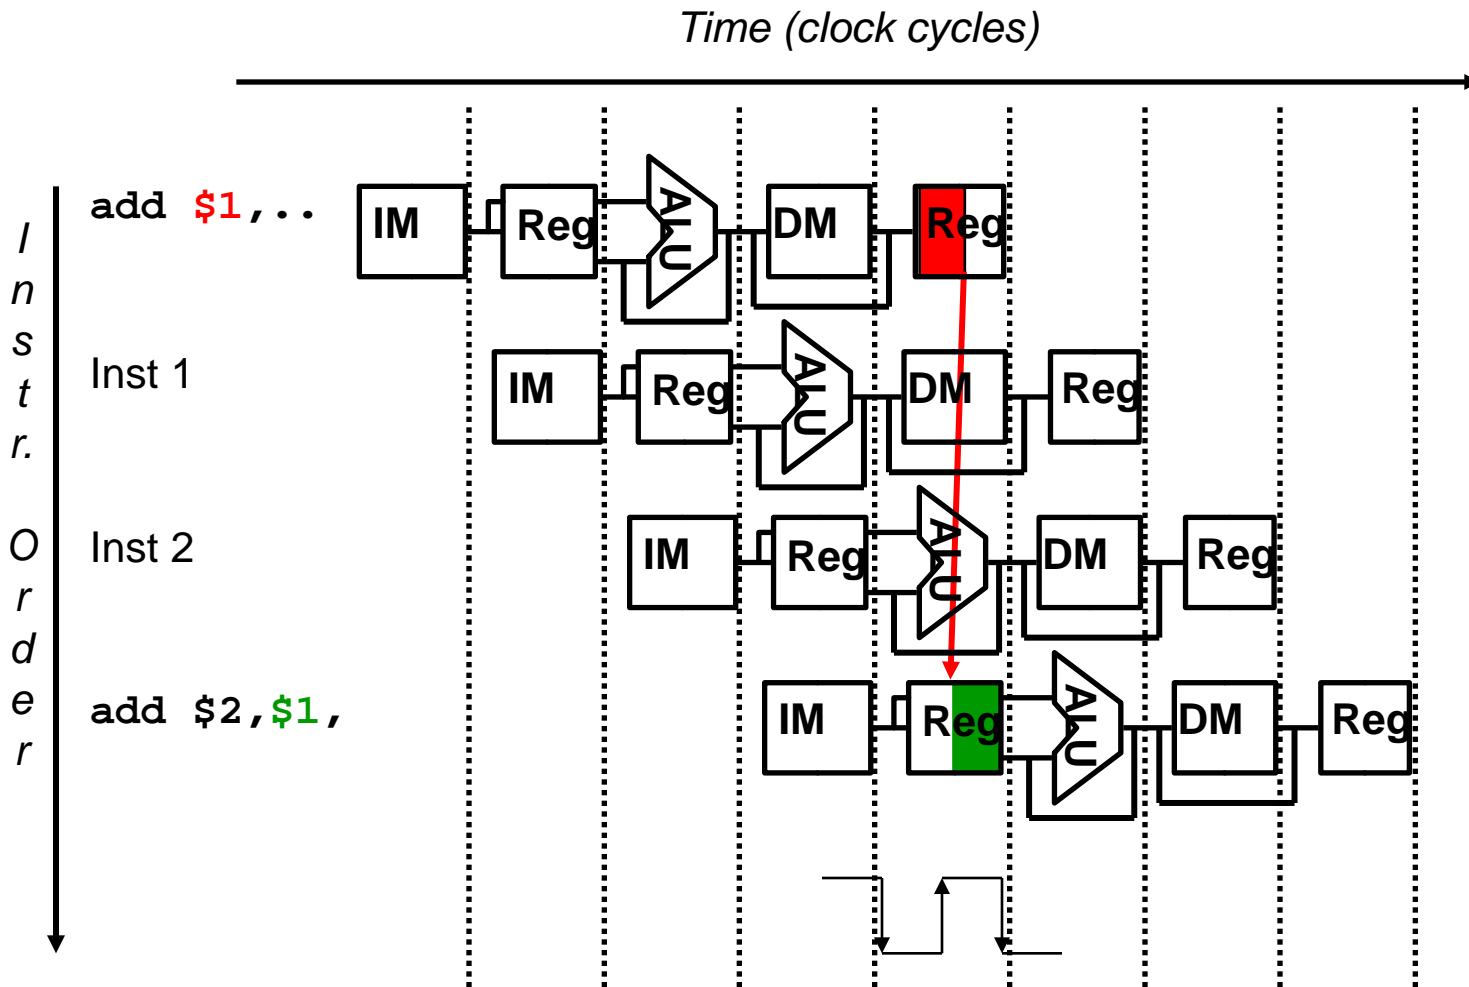
Will be explained again later!

□ MEM hazard:

```
if (MEM/WB.RegWrite
and (MEM/WB.RegisterRd != 0)
and (MEM/WB.RegisterRd = ID/EX.RegisterRs))
    ForwardA = 01
if (MEM/WB.RegWrite
and (MEM/WB.RegisterRd != 0)
and (MEM/WB.RegisterRd = ID/EX.RegisterRt))
    ForwardB = 01
```

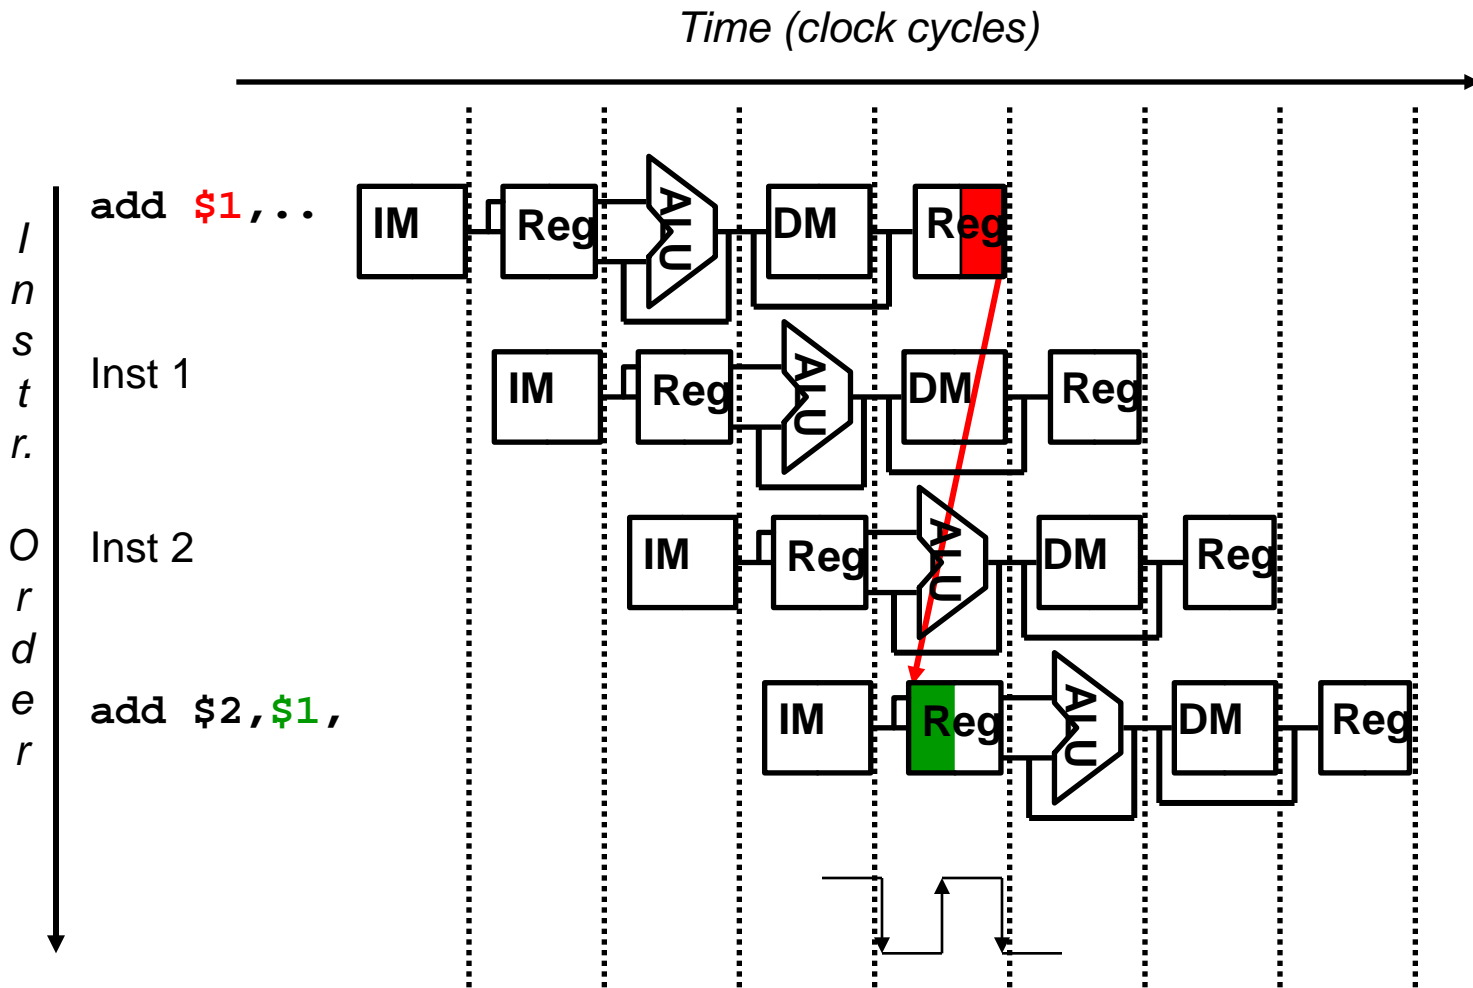
Forwards the
result from the
second
previous instr.
to either input
of the ALU

What if both instructions are accessing the same register?



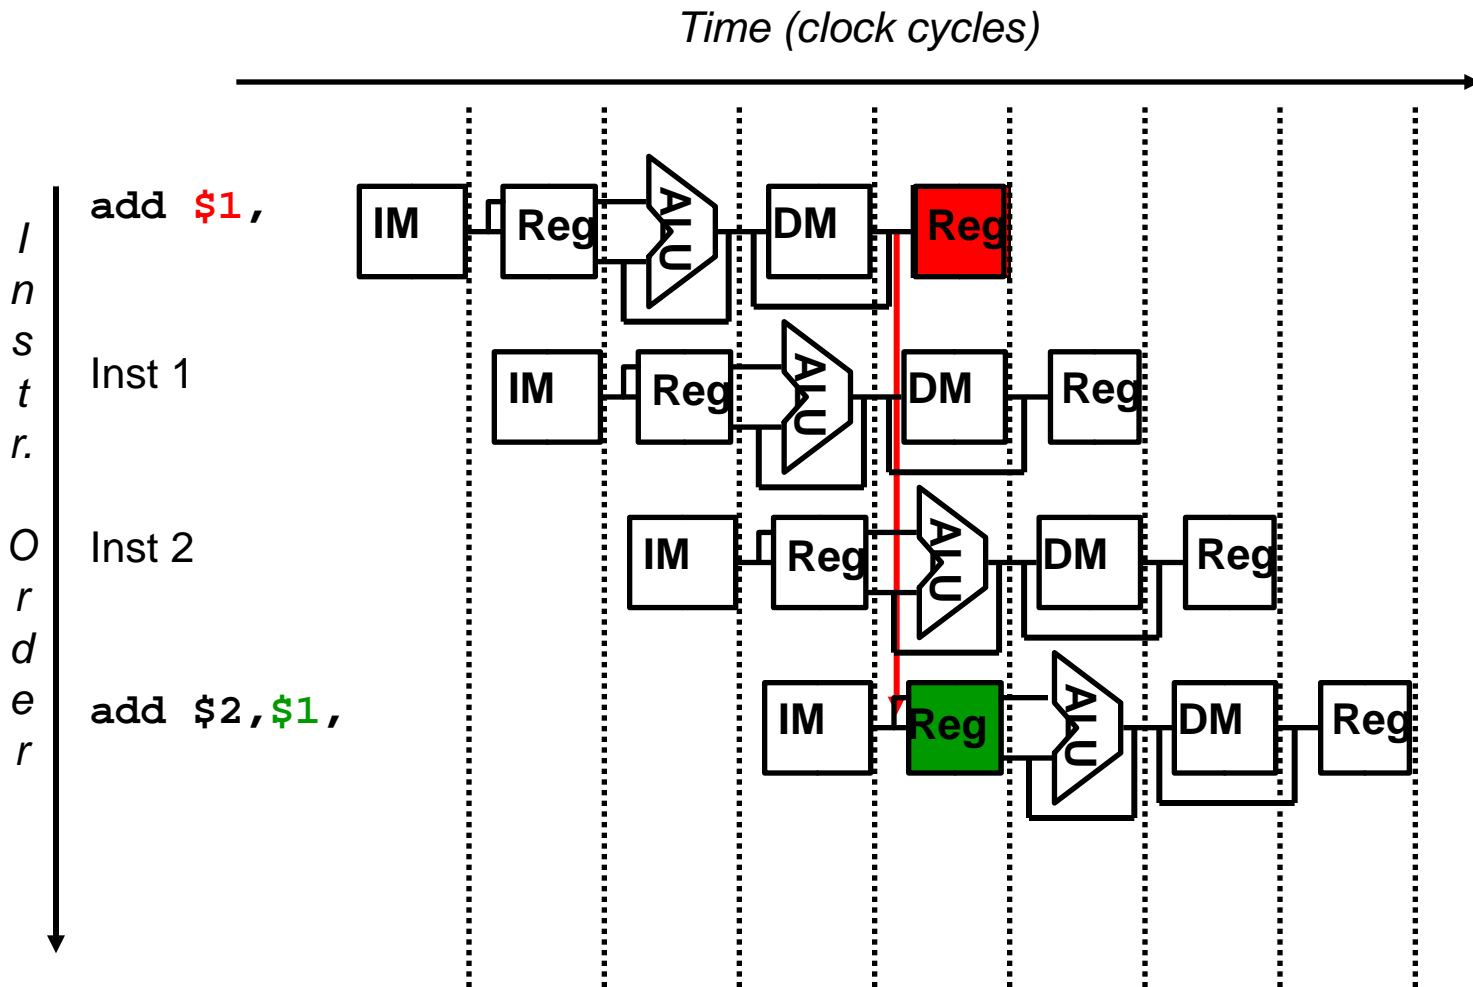
- A data hazard? Must read after write (RAW)

What if both instructions are accessing the same register?



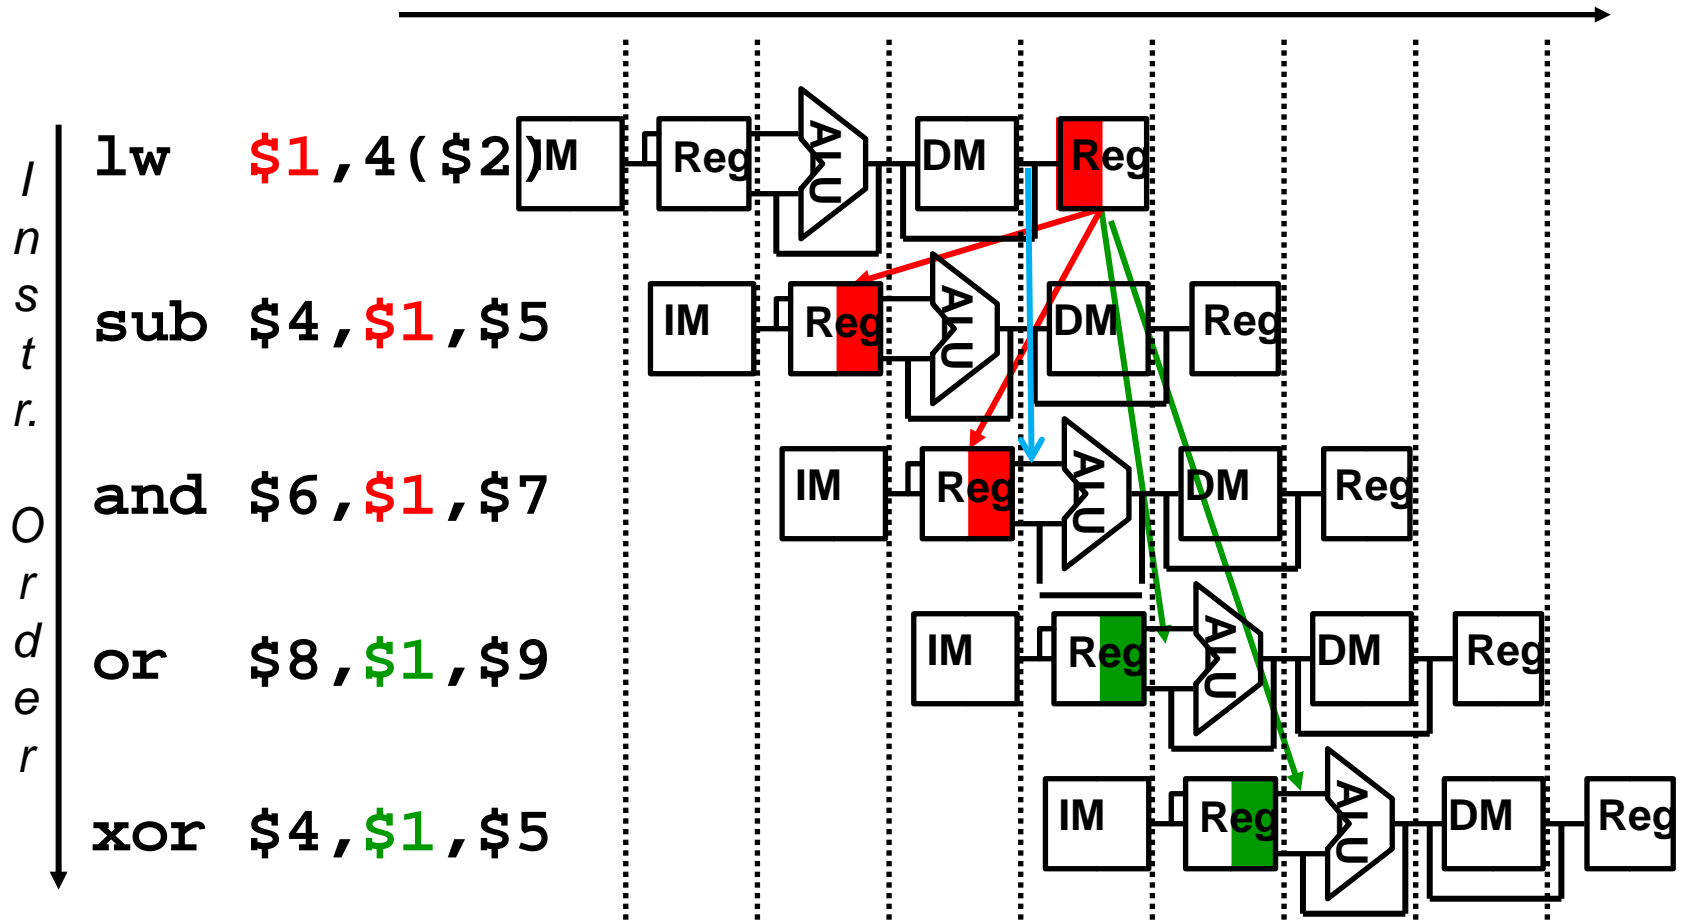
- A data hazard? Must read after write (RAW)

How About Register File Access?



Today: Register reads/writes take a whole cycle!
So in this case we need to “bypass” the data
from the write to the read in the same cycle.

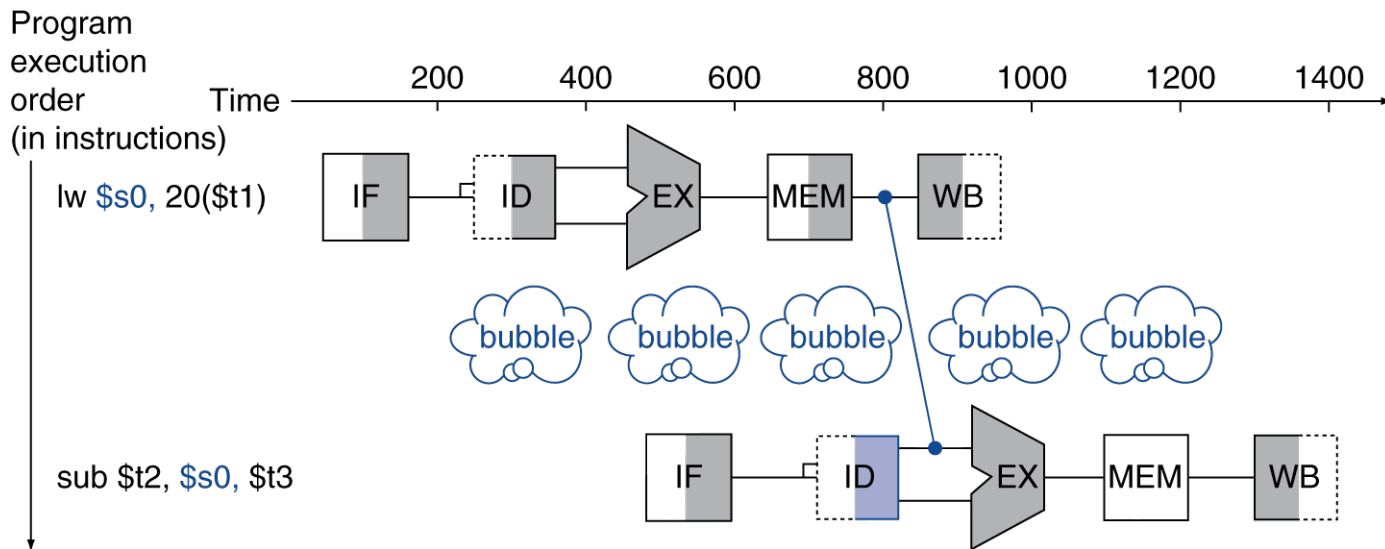
Loads Can Cause Data Hazards



- Load-use data hazard

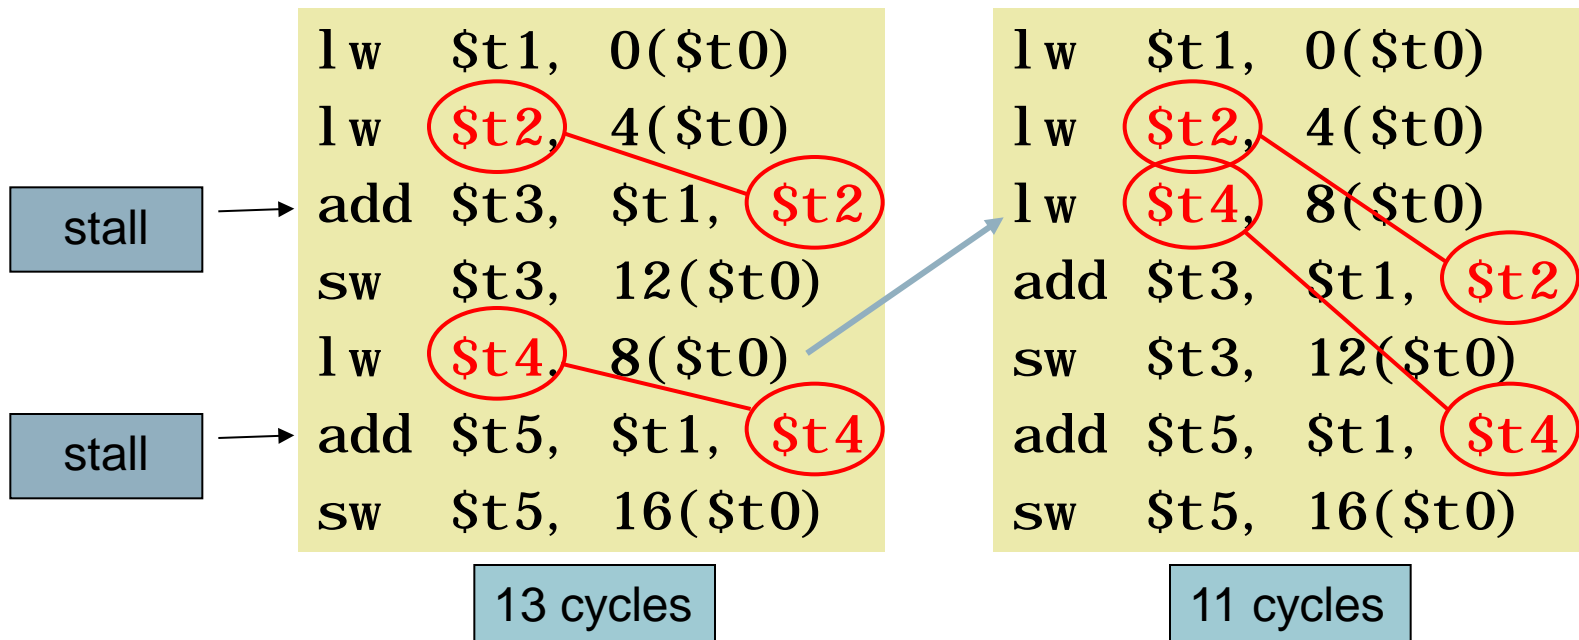
Load-Use Data Hazard

- Can't always **avoid stalls by forwarding**
 - If value not computed when needed
 - **Can't forward backward in time!**



Code Scheduling to Avoid Stalls

- Reorder code to avoid use of load result in the next instruction
- C code for $A = B + E$; $C = B + F$;



Types of Data Hazards

- **RAW** (read after write): **true dependency**
 - **only hazard for** 'fixed' pipelines
 - later instruction must *read* after earlier instruction *writes*
- **WAW** (write after write): **output dependency**
 - **only hazard for** variable-length pipeline
 - later instruction must *write* after earlier instruction *writes*
- **WAR** (write after read): **anti dependency or name dependency**
 - **only hazard for** pipelines with late read
 - later instruction must *write* after earlier instruction *reads*
- **RAR** (read after read): **no hazard**

