



- Please complete the UNGRADED quiz on HW2 performance. It is called HW2 Review and Verilog



ECE/CS 552: Pipeline Hazards

© Prof. Mikko Lipasti

Lecture notes based in part on slides created by Mark Hill, David Wood, Guri Sohi, John Shen and Jim Smith

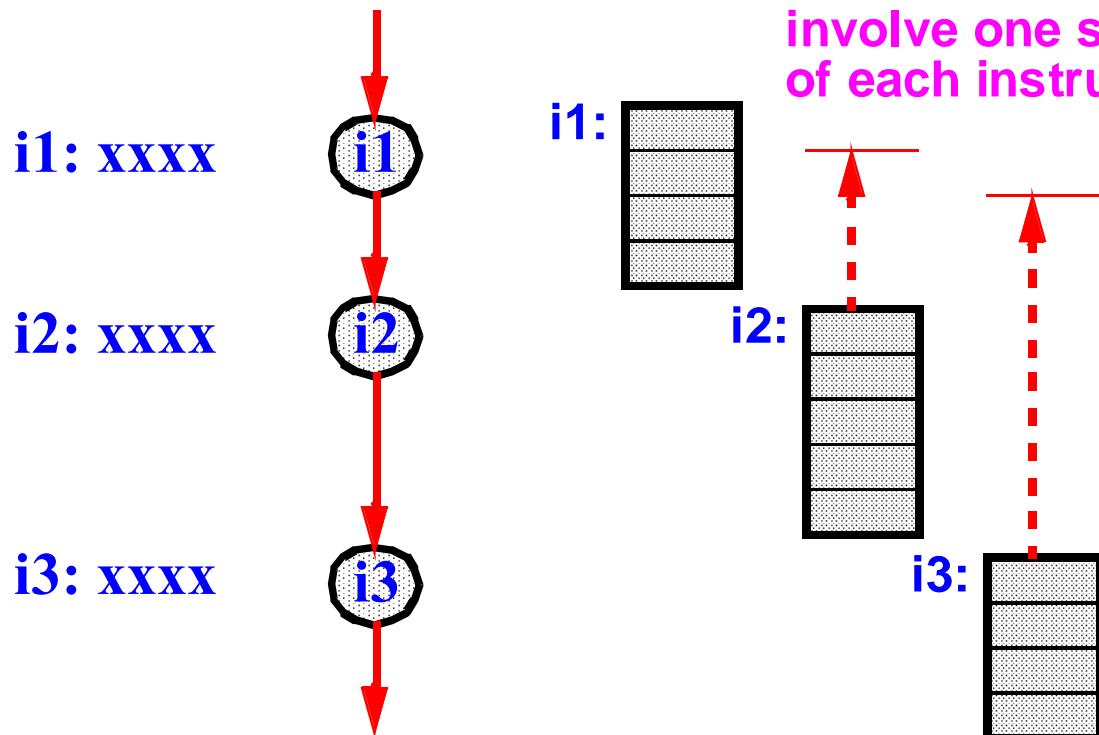
Pipeline Hazards

- Forecast
 - Program Dependences
 - Data Hazards
 - Stalls
 - Forwarding
 - Control Hazards
 - Exceptions

Sequential Execution Model

- MIPS ISA requires the appearance of *sequential execution*
 - *Precise exceptions*
 - True of most general purpose ISAs

Program Dependencies



A true dependence between two instructions may only involve one subcomputation of each instruction.

The implied sequential precedences are an overspecification. It is sufficient but not necessary to ensure program correctness.

Program Data Dependencies

- True dependence (RAW)
 - j cannot execute until i produces its result
- Anti-dependence (WAR)
 - j cannot write its result until i has read its sources
- Output dependence (WAW)
 - j cannot write its result until i has written its result

$$D(i) \cap R(j) \neq \emptyset$$

$$R(i) \cap D(j) \neq \emptyset$$

$$D(i) \cap D(j) \neq \emptyset$$

Control Dependencies

- Conditional branches
 - Branch must execute to determine which instruction to fetch next
 - Instructions following a conditional branch are control dependent on the branch instruction

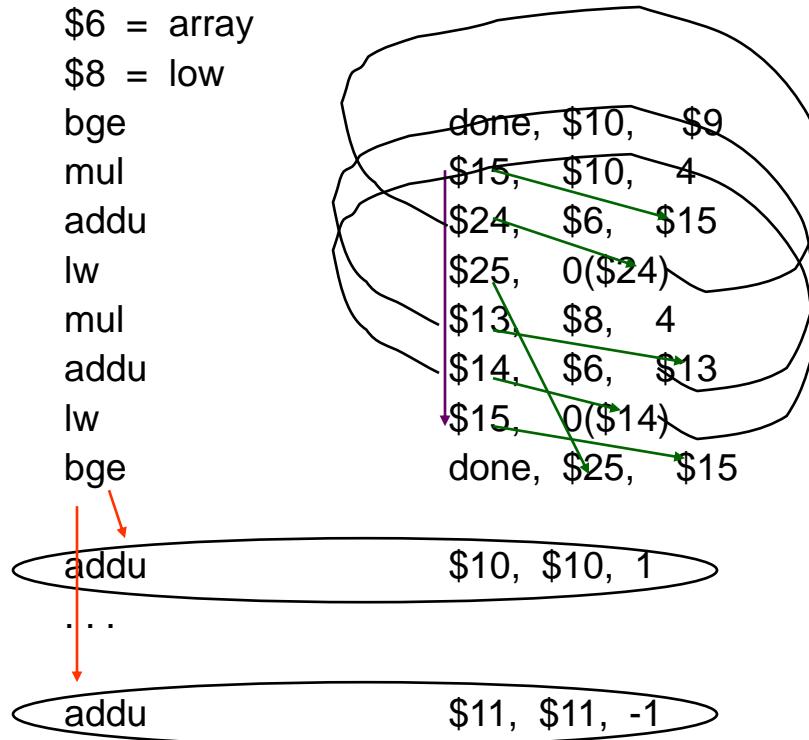
Example (quicksort/MIPS)

```
#      for ( ; (j < high) && (array[j] < array[low]) ; ++j );
#      $10 = j
#      $9 = high
#      $6 = array
#      $8 = low
```

```
bge
mul
addu
lw
mul
addu
lw
bge
```

cont:

done:



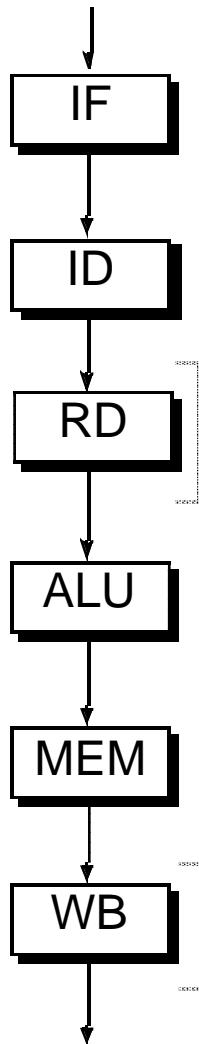
Pipeline Hazards

- Pipeline hazards
 - Potential violations of program dependences
 - Must ensure program dependences are not violated
- Hazard resolution
 - Static: compiler/programmer guarantees correctness
 - Dynamic: hardware performs checks at runtime
- Pipeline interlock
 - Hardware mechanism for dynamic hazard resolution
 - Must detect and enforce dependences at runtime

Pipeline Hazards

- Necessary conditions:
 - WAR: write stage earlier than read stage
 - Is this possible in IF-RD-EX-MEM-WB ?
 - WAW: write stage earlier than write stage
 - Is this possible in IF-RD-EX-MEM-WB ?
 - RAW: read stage earlier than write stage
 - Is this possible in IF-RD-EX-MEM-WB?
- If conditions not met, no need to resolve
- Check for both register and memory

Pipeline Hazard Analysis



- Memory hazards
 - RAW: Yes/No?
 - WAR: Yes/No?
 - WAW: Yes/No?
- Register hazards
 - RAW: Yes/No?
 - WAR: Yes/No?
 - WAW: Yes/No?

RAW Hazard

- Earlier instruction produces a value used by a later instruction:
 - add \$1, \$2, \$3
 - sub \$4, \$5, \$1

Cycle:	1	2	3	4	5	6	7	8	9	10	11	12	13
Instr:													
add	F	D	X	M	W								
sub		F	D	X	M	W							

RAW Hazard - Stall

- Detect dependence and stall:
 - add \$1, \$2, \$3
 - sub \$4, \$5, \$1

Cycle:	1	2	3	4	5	6	7	8	9	10	11	12	13
Instr:									0	1	2	3	
add	F	D	X	M	W								
sub						F	D	X	M	W			

Control Dependence

- One instruction affects which executes next
 - sw \$4, 0(\$5)
 - bne \$2, \$3, loop
 - sub \$6, \$7, \$8

Cycle:	1	2	3	4	5	6	7	8	9	1	1	1	1
Instr:									0	1	1	2	3
sw	F	D	X	M	W								
bne		F	D	X	M	W							
sub			F	D	X	M	W						

Control Dependence - Stall

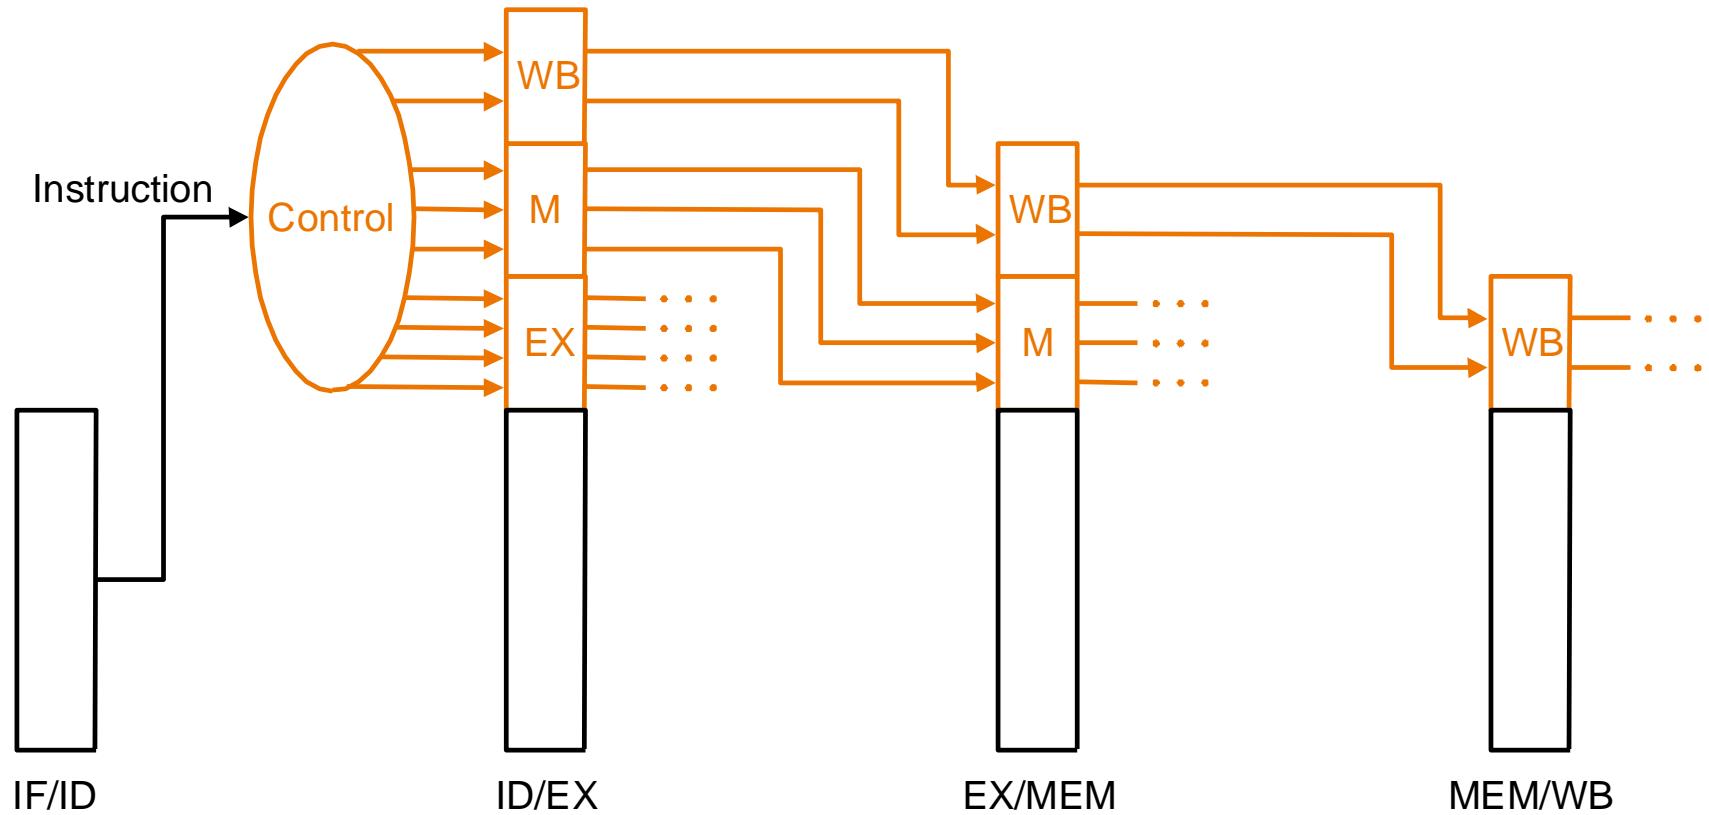
- Detect dependence and stall
 - sw \$4, 0(\$5)
 - bne \$2, \$3, loop
 - sub \$6, \$7, \$8

Cycle:	1	2	3	4	5	6	7	8	9	1	1	1	1
Instr:									0	1	2	3	
sw	F	D	X	M	W								
bne		F	D	X	M	W							
sub					F	D	X	M	W				

Pipelined Control

- Controlled by different instructions
- Decode instructions and pass the signals down the pipe
- Control sequencing is embedded in the pipeline
 - No explicit FSM
 - Instead, distributed FSM

Pipelined Control



RAW Hazards

- Must first detect RAW hazards
 - Pipeline analysis proves that WAR/WAW don't occur

ID/EX.WriteRegister = IF/ID.ReadRegister1

ID/EX.WriteRegister = IF/ID.ReadRegister2

EX/MEM.WriteRegister = IF/ID.ReadRegister1

EX/MEM.WriteRegister = IF/ID.ReadRegister2

MEM/WB.WriteRegister = IF/ID.ReadRegister1

MEM/WB.WriteRegister = IF/ID.ReadRegister2

RAW Hazards

- Not all hazards because
 - WriteRegister not used (e.g. sw)
 - ReadRegister not used (e.g. addi, jump)
 - Do something only if necessary

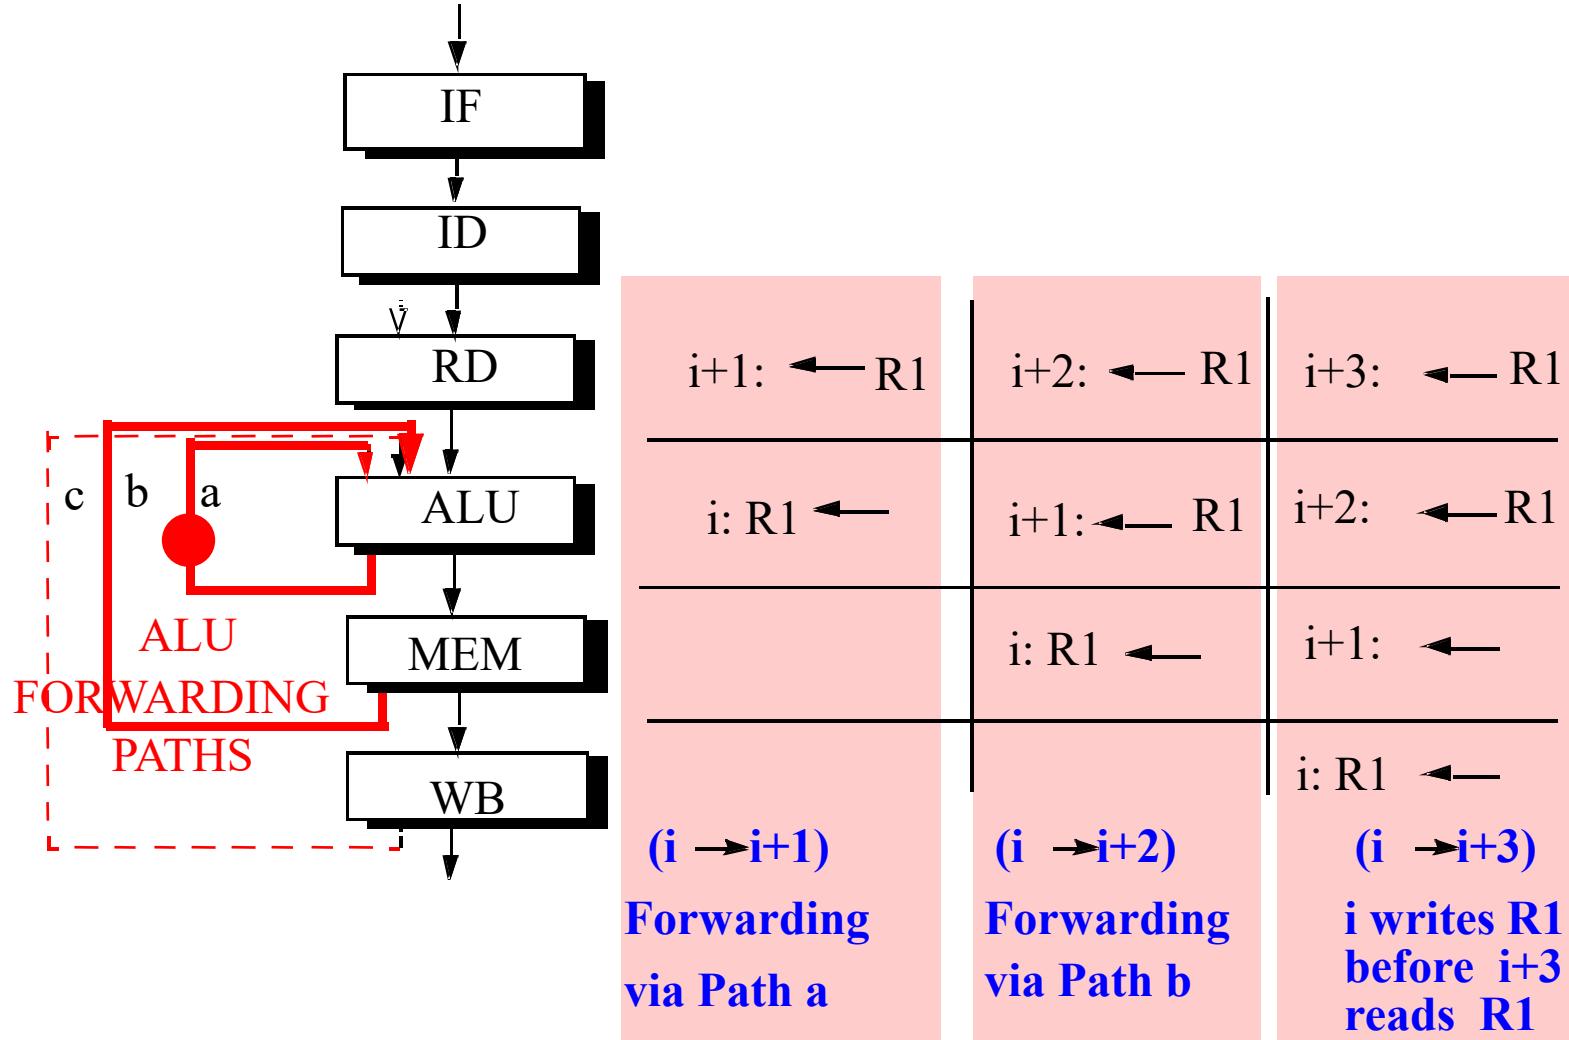
RAW Hazards

- Hazard Detection Unit
 - Several 5-bit comparators
- Response? Stall pipeline
 - Instructions in IF and ID stay
 - IF/ID pipeline latch not updated
 - Send ‘nop’ down pipeline (called a bubble)
 - PCWrite, IF/IDWrite, and nop mux

RAW Hazard Forwarding

- A better response – forwarding
 - Also called bypassing
- Comparators ensure register is read after it is written
- Instead of stalling until write occurs
 - Use mux to select forwarded value rather than register value
 - Control mux with hazard detection logic

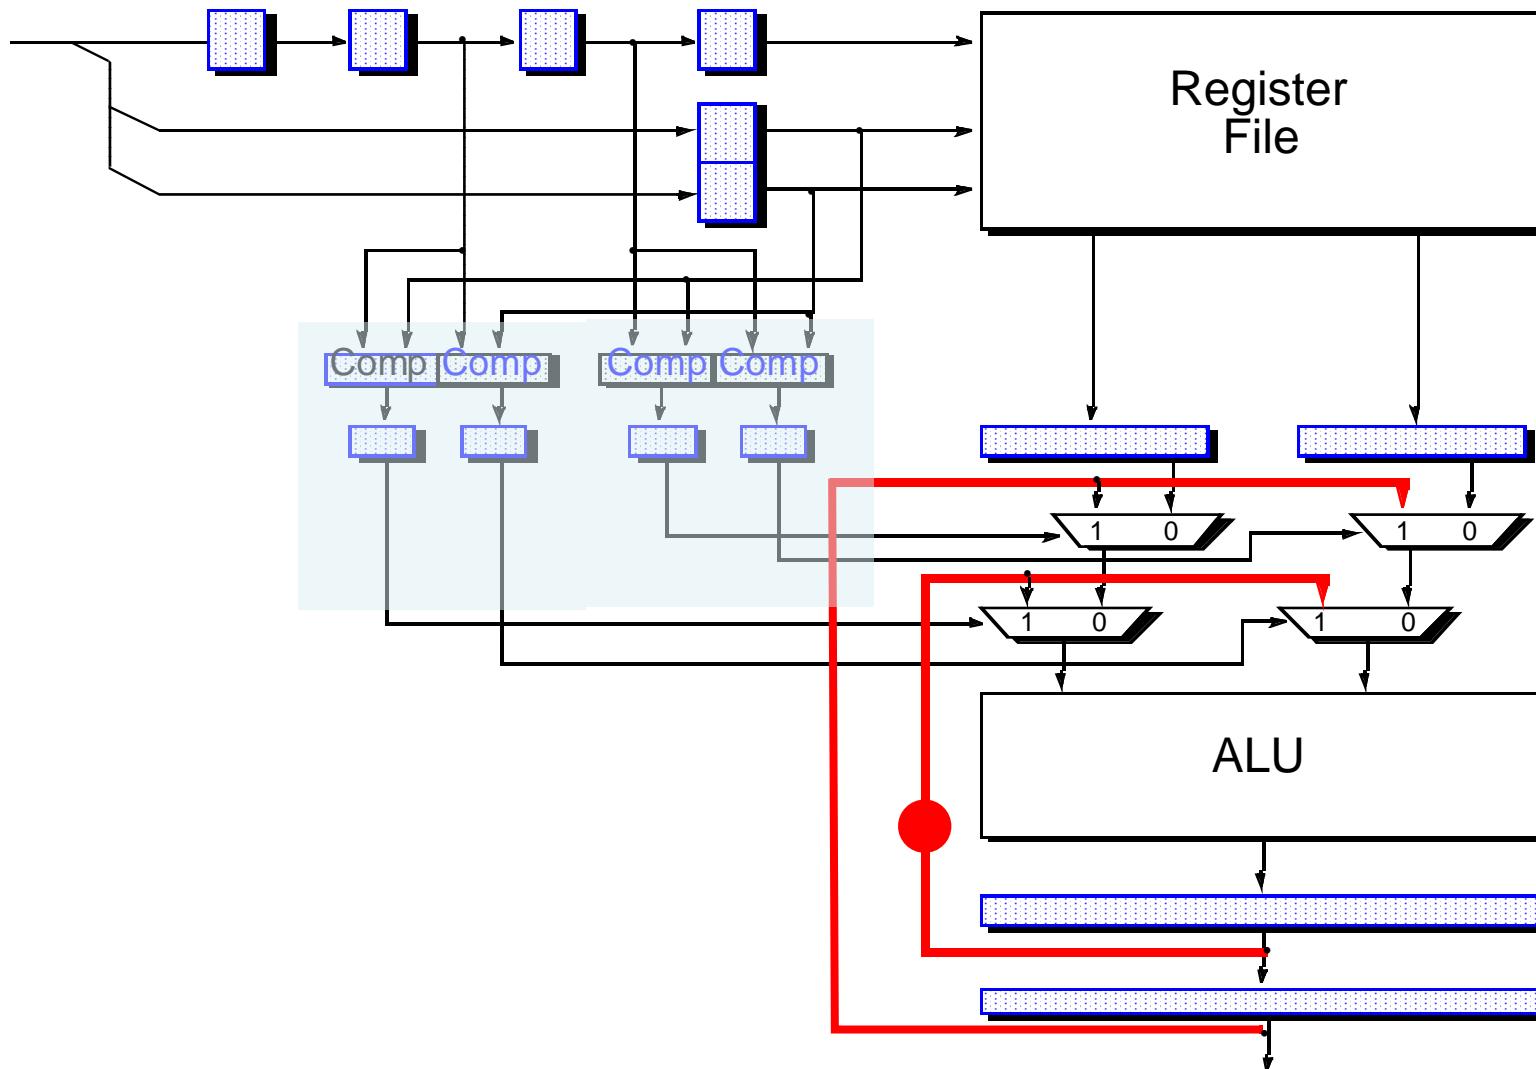
Forwarding Paths (ALU instructions)



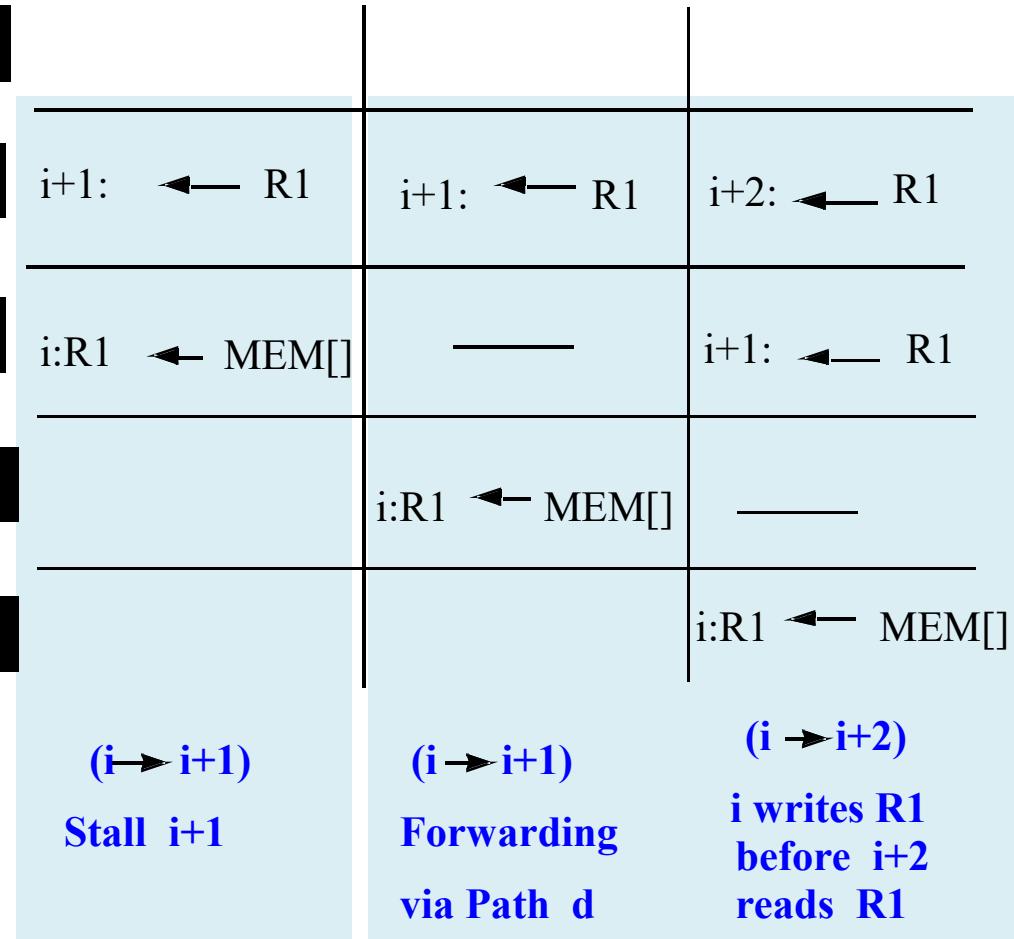
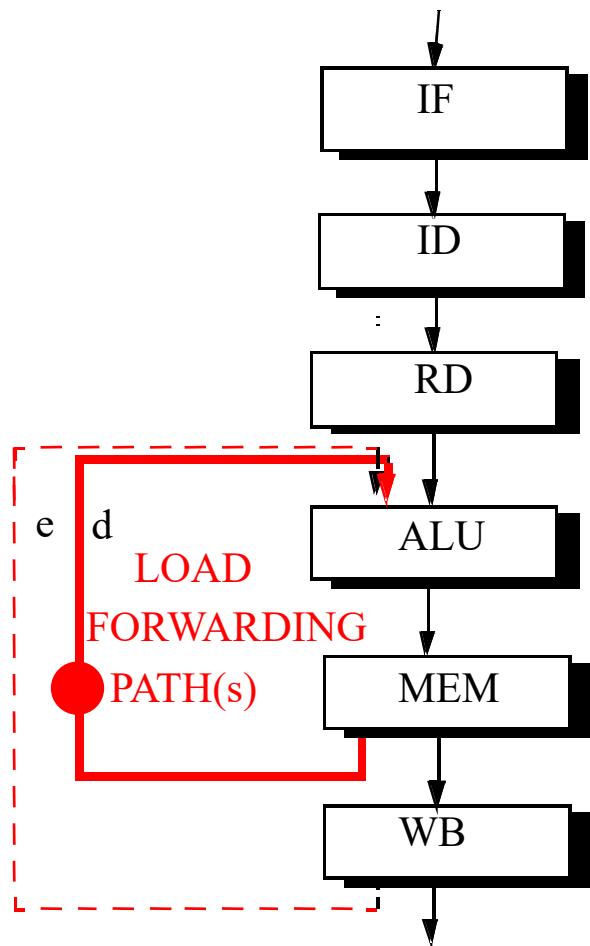
Write before Read RF

- Register file design
 - 2-phase clocks common
 - Write RF on first phase
 - Read RF on second phase
- Hence, same cycle:
 - Write \$1
 - Read \$1
- No bypass needed
 - If read before write or DFF-based, need bypass

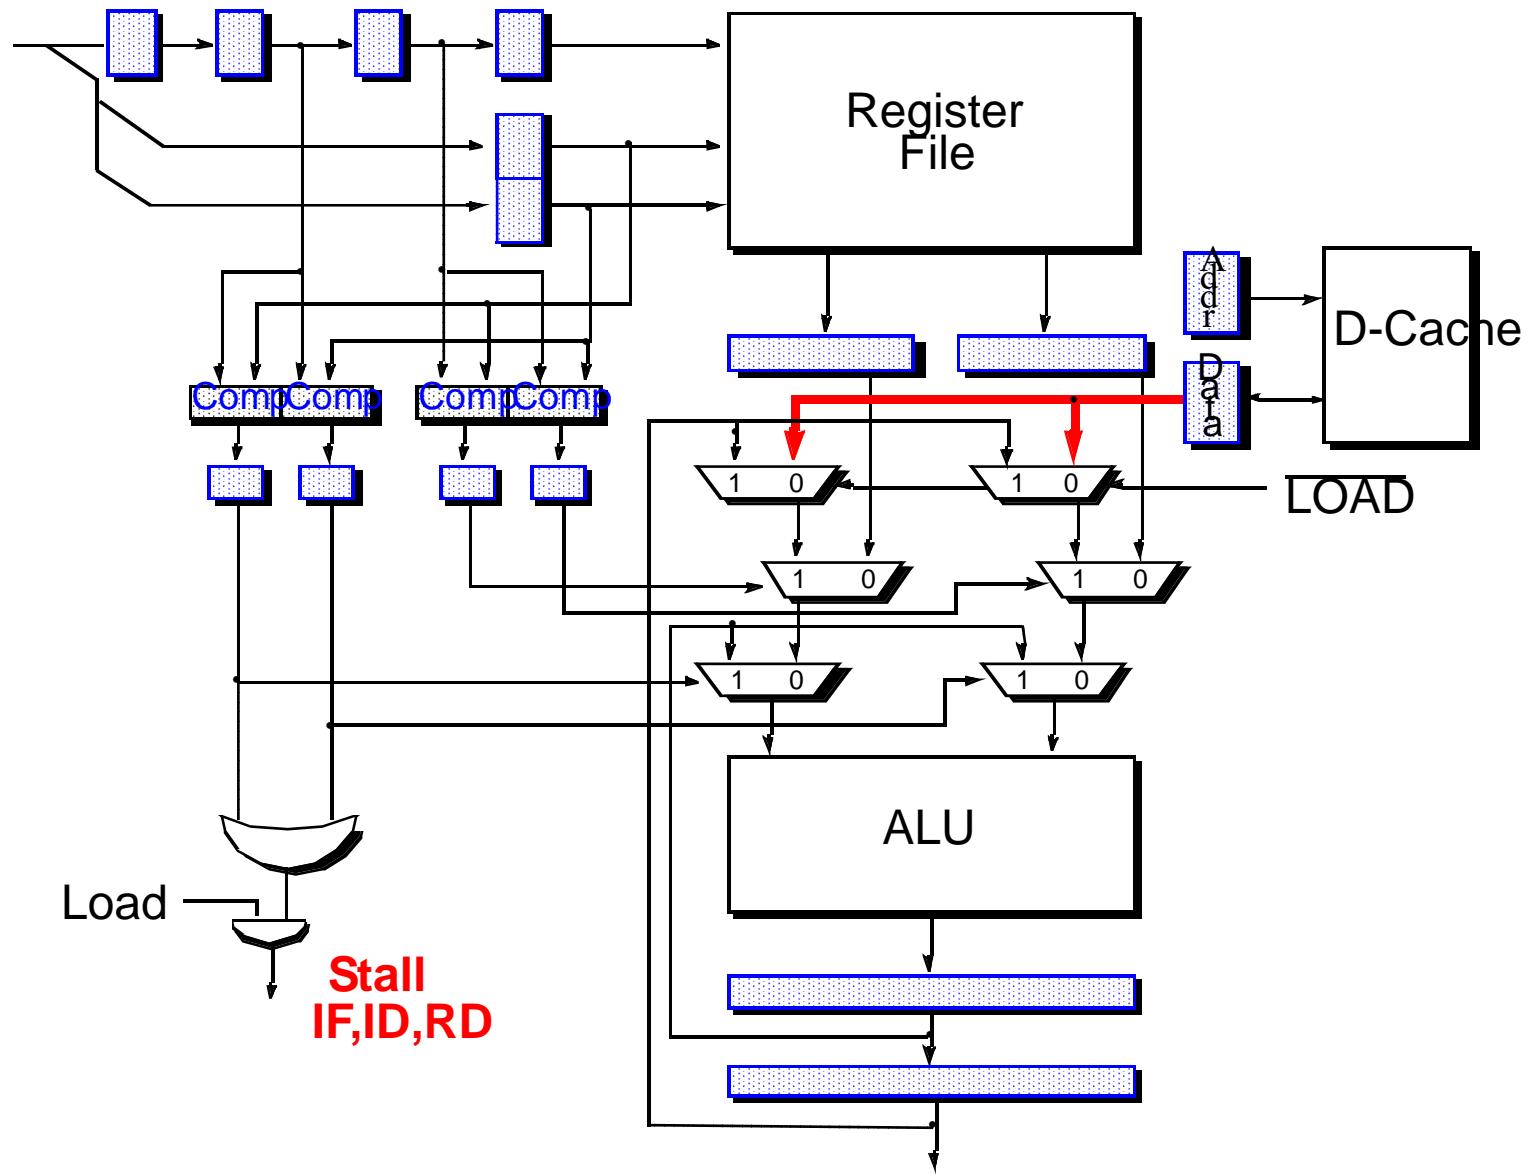
ALU Forwarding



Forwarding Paths (Load instructions)



Implementation of Load Forwarding



Pipeline with Forwarding

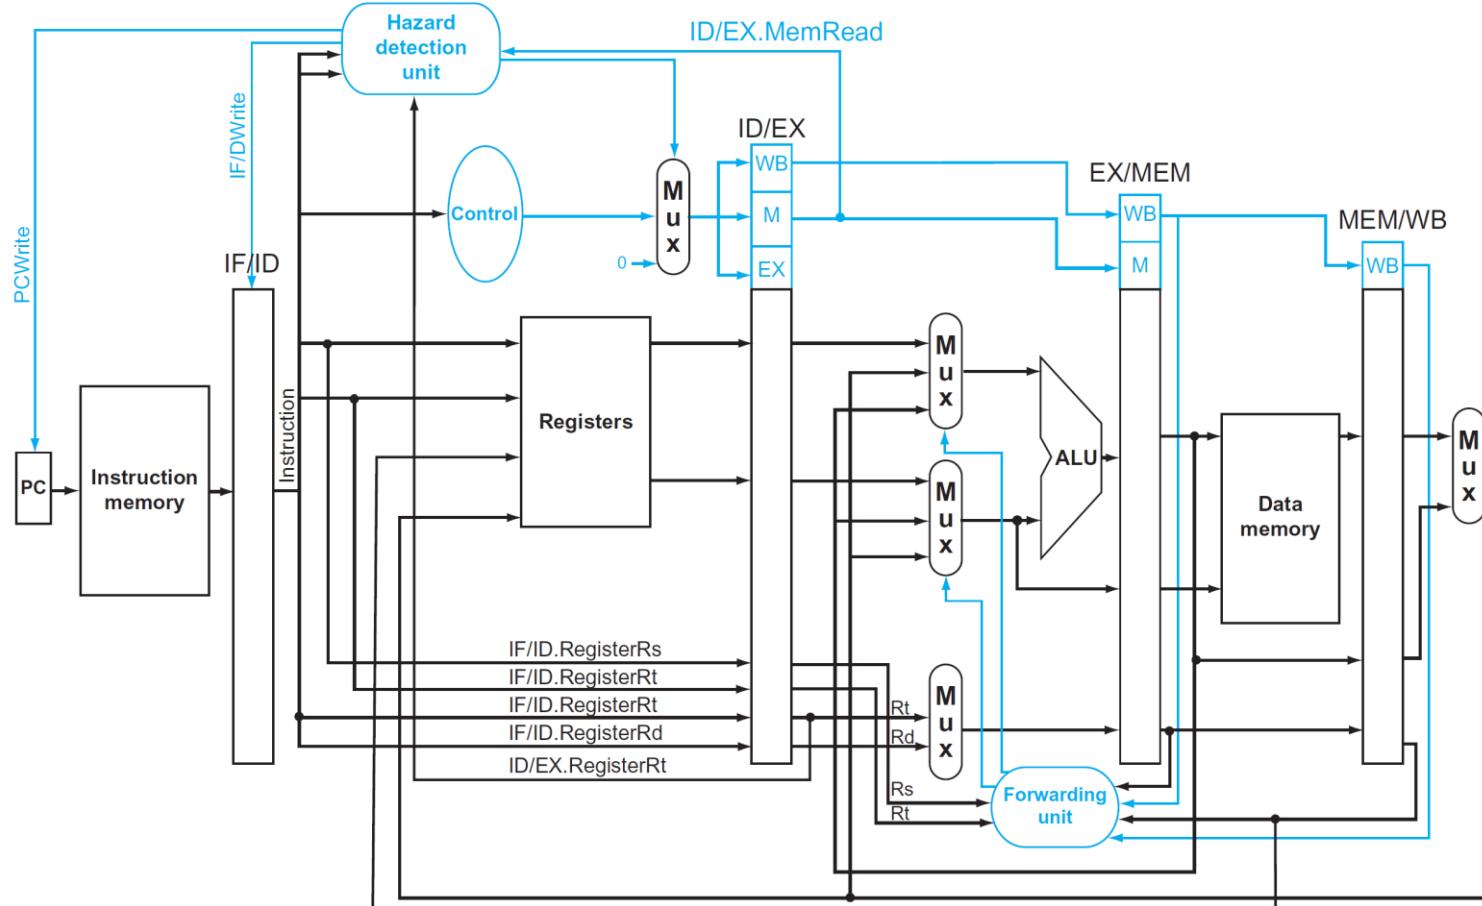


FIGURE 4.60 Pipelined control overview, showing the two multiplexors for forwarding, the hazard detection unit, and the forwarding unit. Although the ID and EX stages have been simplified—the sign-extended immediate and branch logic are missing—this drawing gives the essence of the forwarding hardware requirements.

Control Flow Hazards

- Control flow instructions
 - branches, jumps, jals, returns
 - Can't fetch until branch outcome known
 - Too late for next IF

Control Flow Hazards

- What to do?
 - Always stall
 - Easy to implement
 - Performs poorly
 - $1/6^{\text{th}}$ instructions are branches
 - each branch takes 3 cycles
 - $\text{CPI} = 1 + 3 \times 1/6 = 1.5$ (lower bound)

Control Flow Hazards

- Predict branch not taken
- Send sequential instructions down pipeline
- Kill instructions later if incorrect
- Must stop memory accesses and RF writes
- Late flush of instructions on misprediction
 - Complex
 - Global signal (wire delay)

Control Flow Hazards

- Even better but more complex
 - Predict taken
 - Predict both (eager execution)
 - Predict one or the other dynamically
 - Adapt to program branch patterns
 - Lots of chip real estate these days
 - Core i7, ARM A15
 - Current research topic
 - More later, covered in detail in ECE752

Control Flow Hazards

- Another option: delayed branches
 - Always execute following instruction
 - “delay slot” (later example on MIPS pipeline)
 - Put useful instruction there, otherwise ‘nop’
- A mistake to cement this into ISA
 - Just a stopgap (one cycle, one instruction)
 - Superscalar processors (later)
 - Delay slot just gets in the way

Pipelined Datapath and Control path

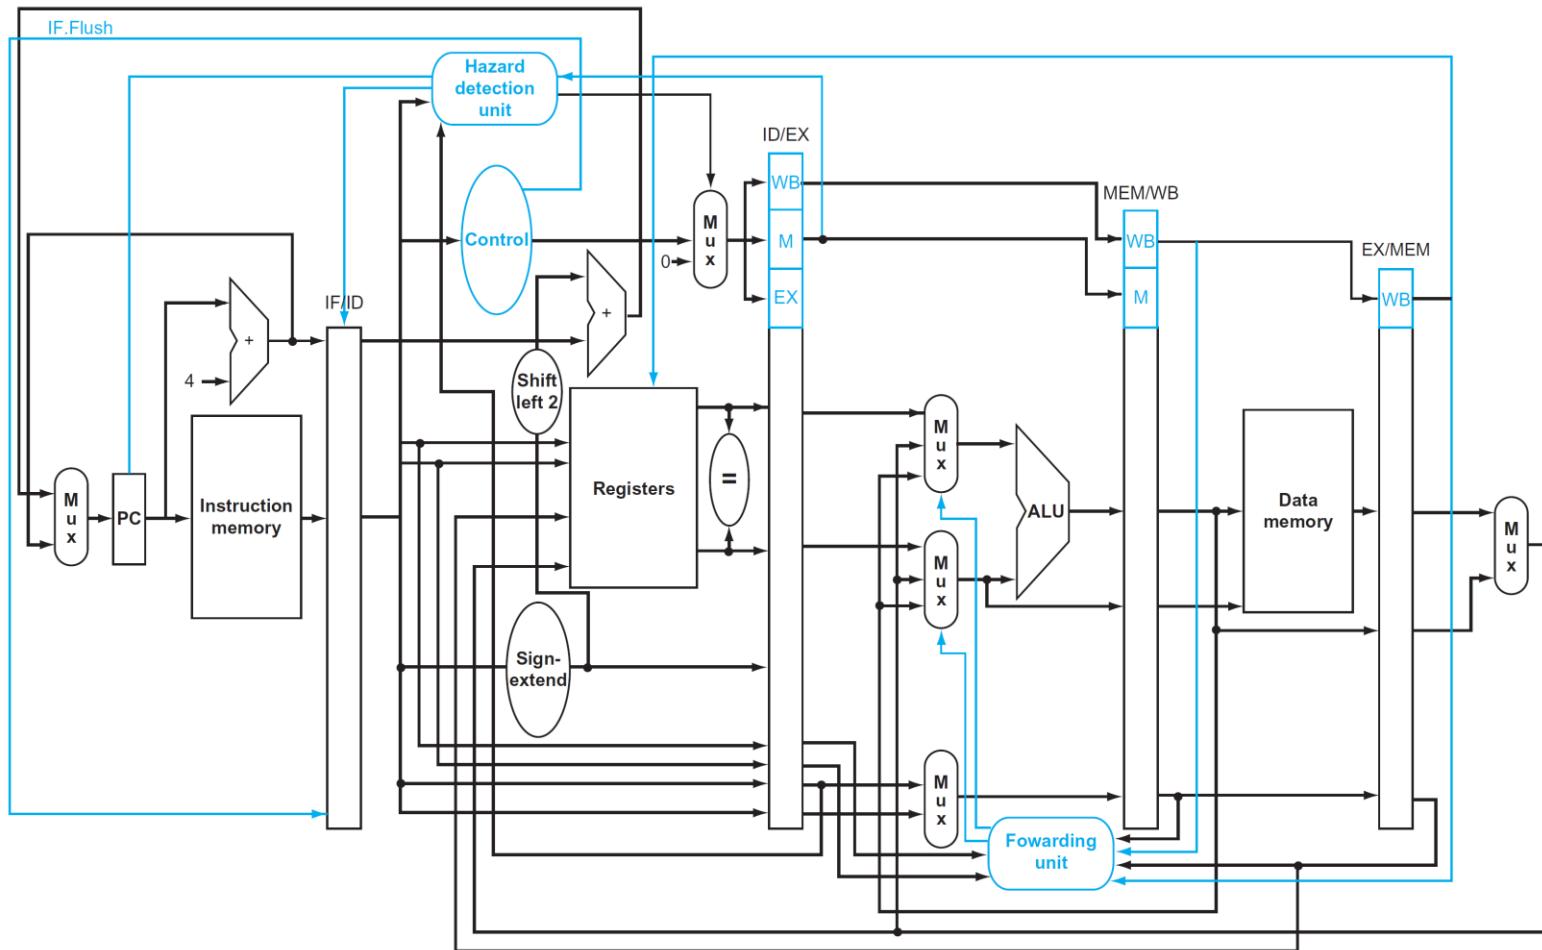


FIGURE 4.65 The final datapath and control for this chapter. Note that this is a stylized figure rather than a detailed datapath, so it's missing the ALUsrc Mux from Figure 4.57 and the multiplexer controls from Figure 4.51.

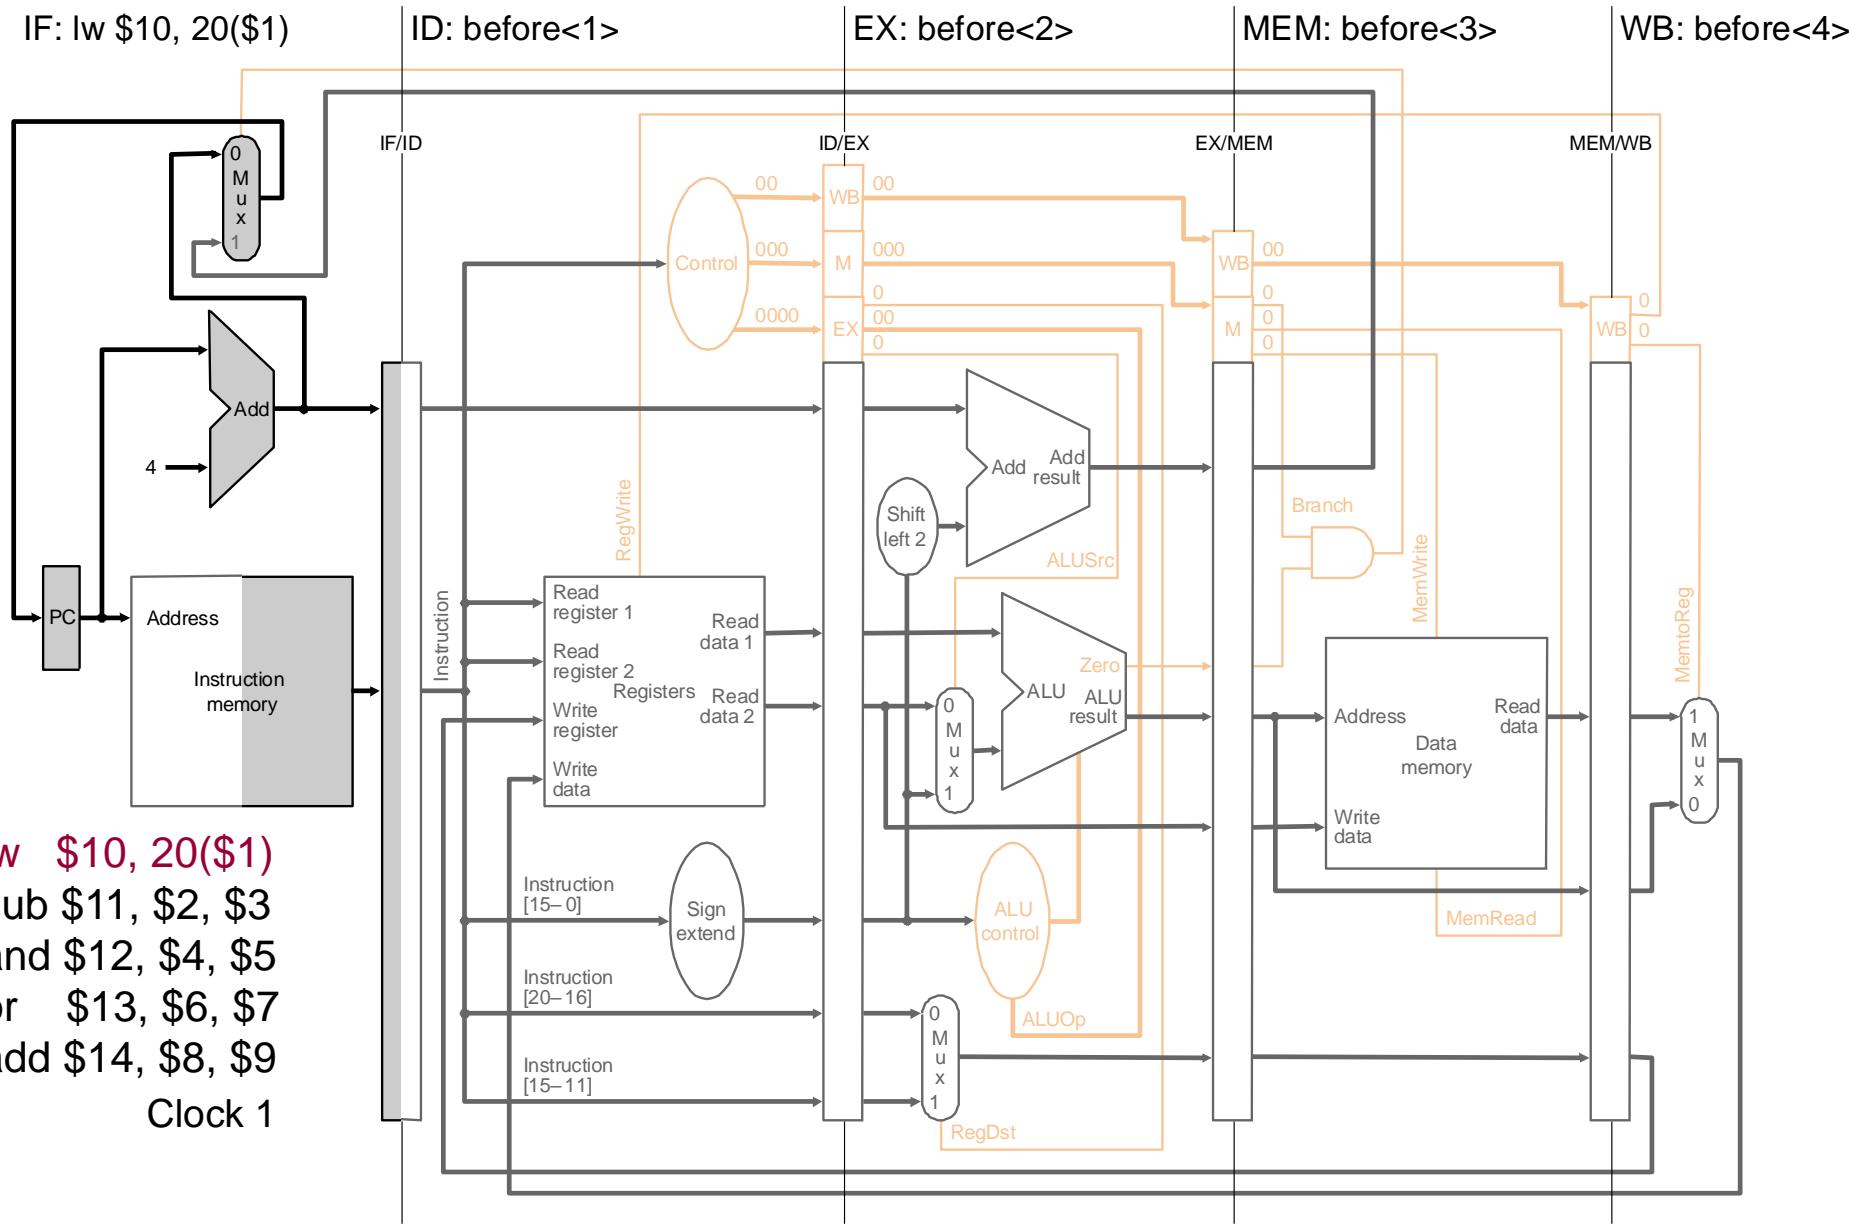
Pipeline Walkthrough Examples

- Example complex sequence
- Detailed walkthrough
 - No dependences
 - Load dependence
 - Branch instruction

Pipeline Walkthrough with controls

- Use walkthrough worksheets
 - Use code segment shown
 - Fill in controls
 - Interesting stages
 - Controls **generated** in Decode stage
 - Controls **consumed** in subsequent stages
 - NOTE: THERE ARE INTENTIONALLY NO DATA DEPENDENCES OR BRANCH INSTRUCTIONS IN SNIPPET

lw \$10, 20(\$1)
sub \$11, \$2, \$3
and \$12, \$4, \$5
or \$13, \$6, \$7
add \$14, \$8, \$9



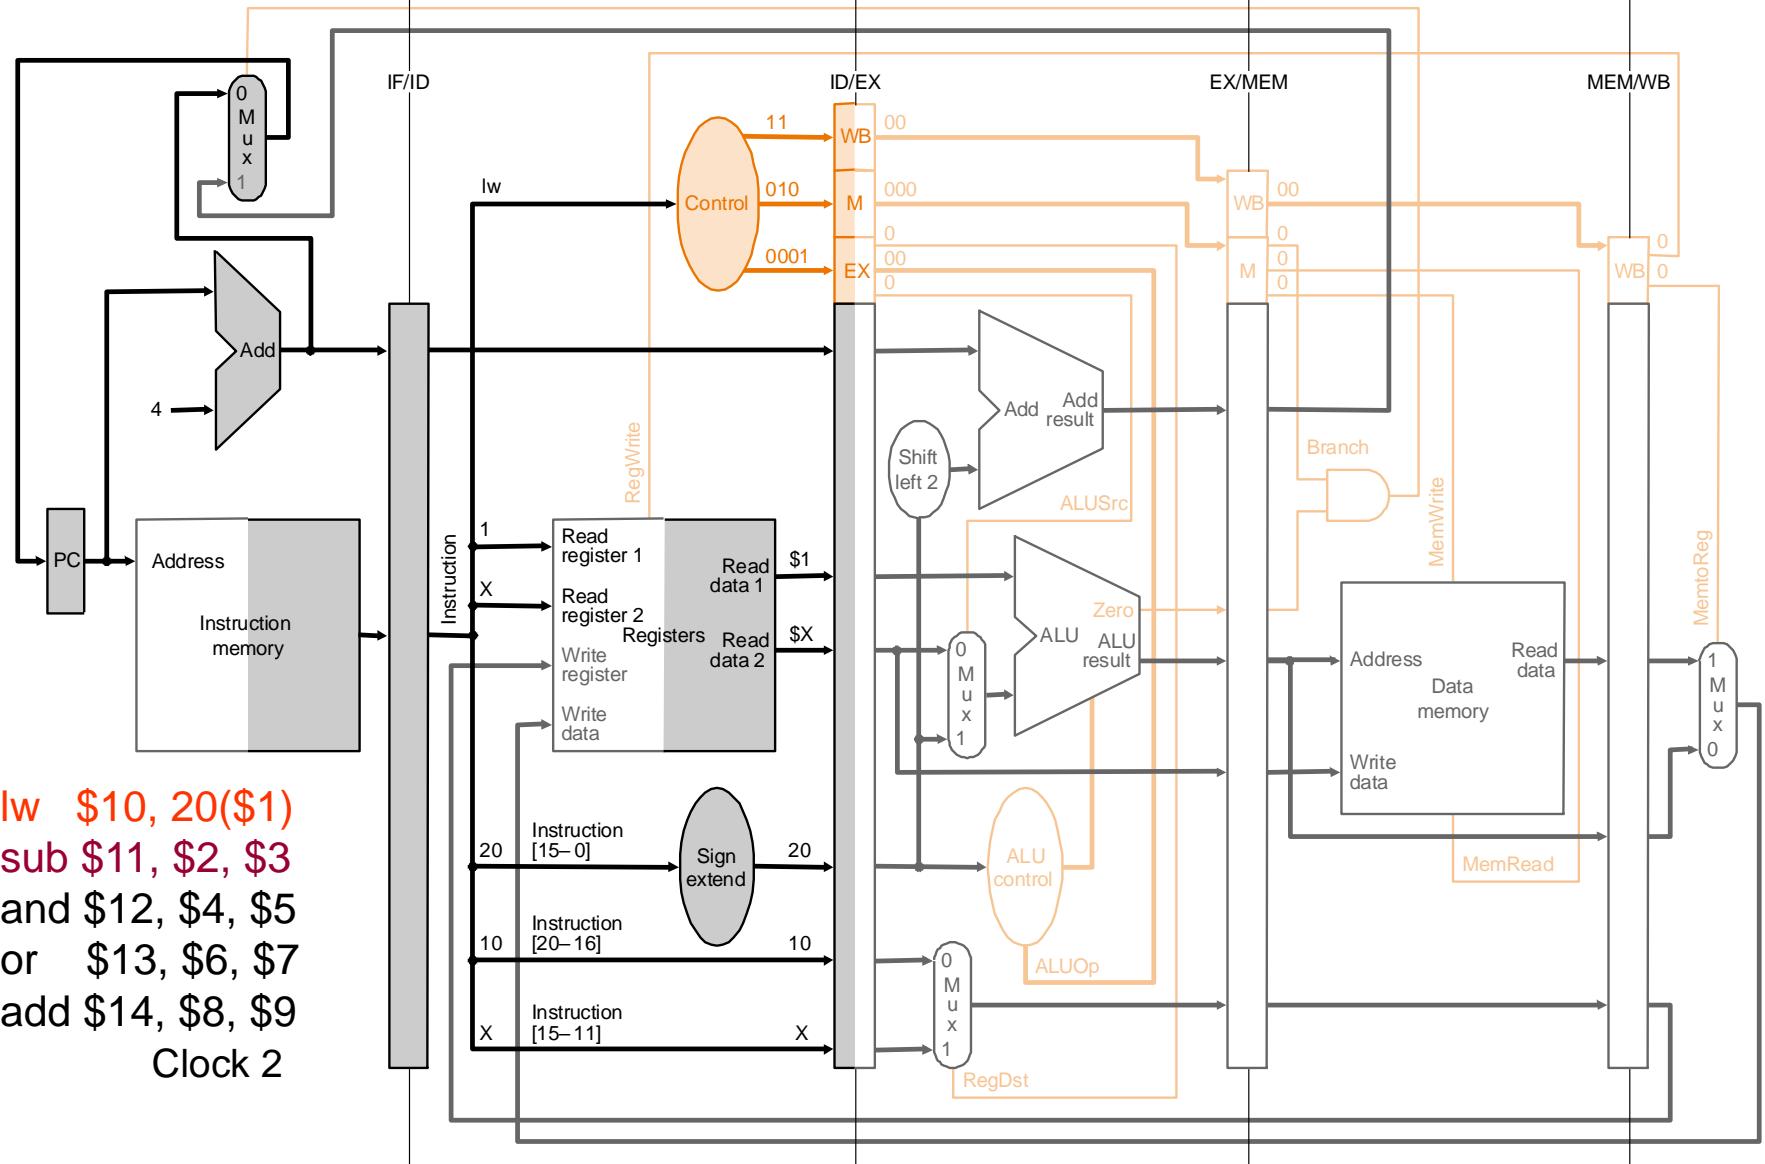
IF: sub \$11, \$2, \$3

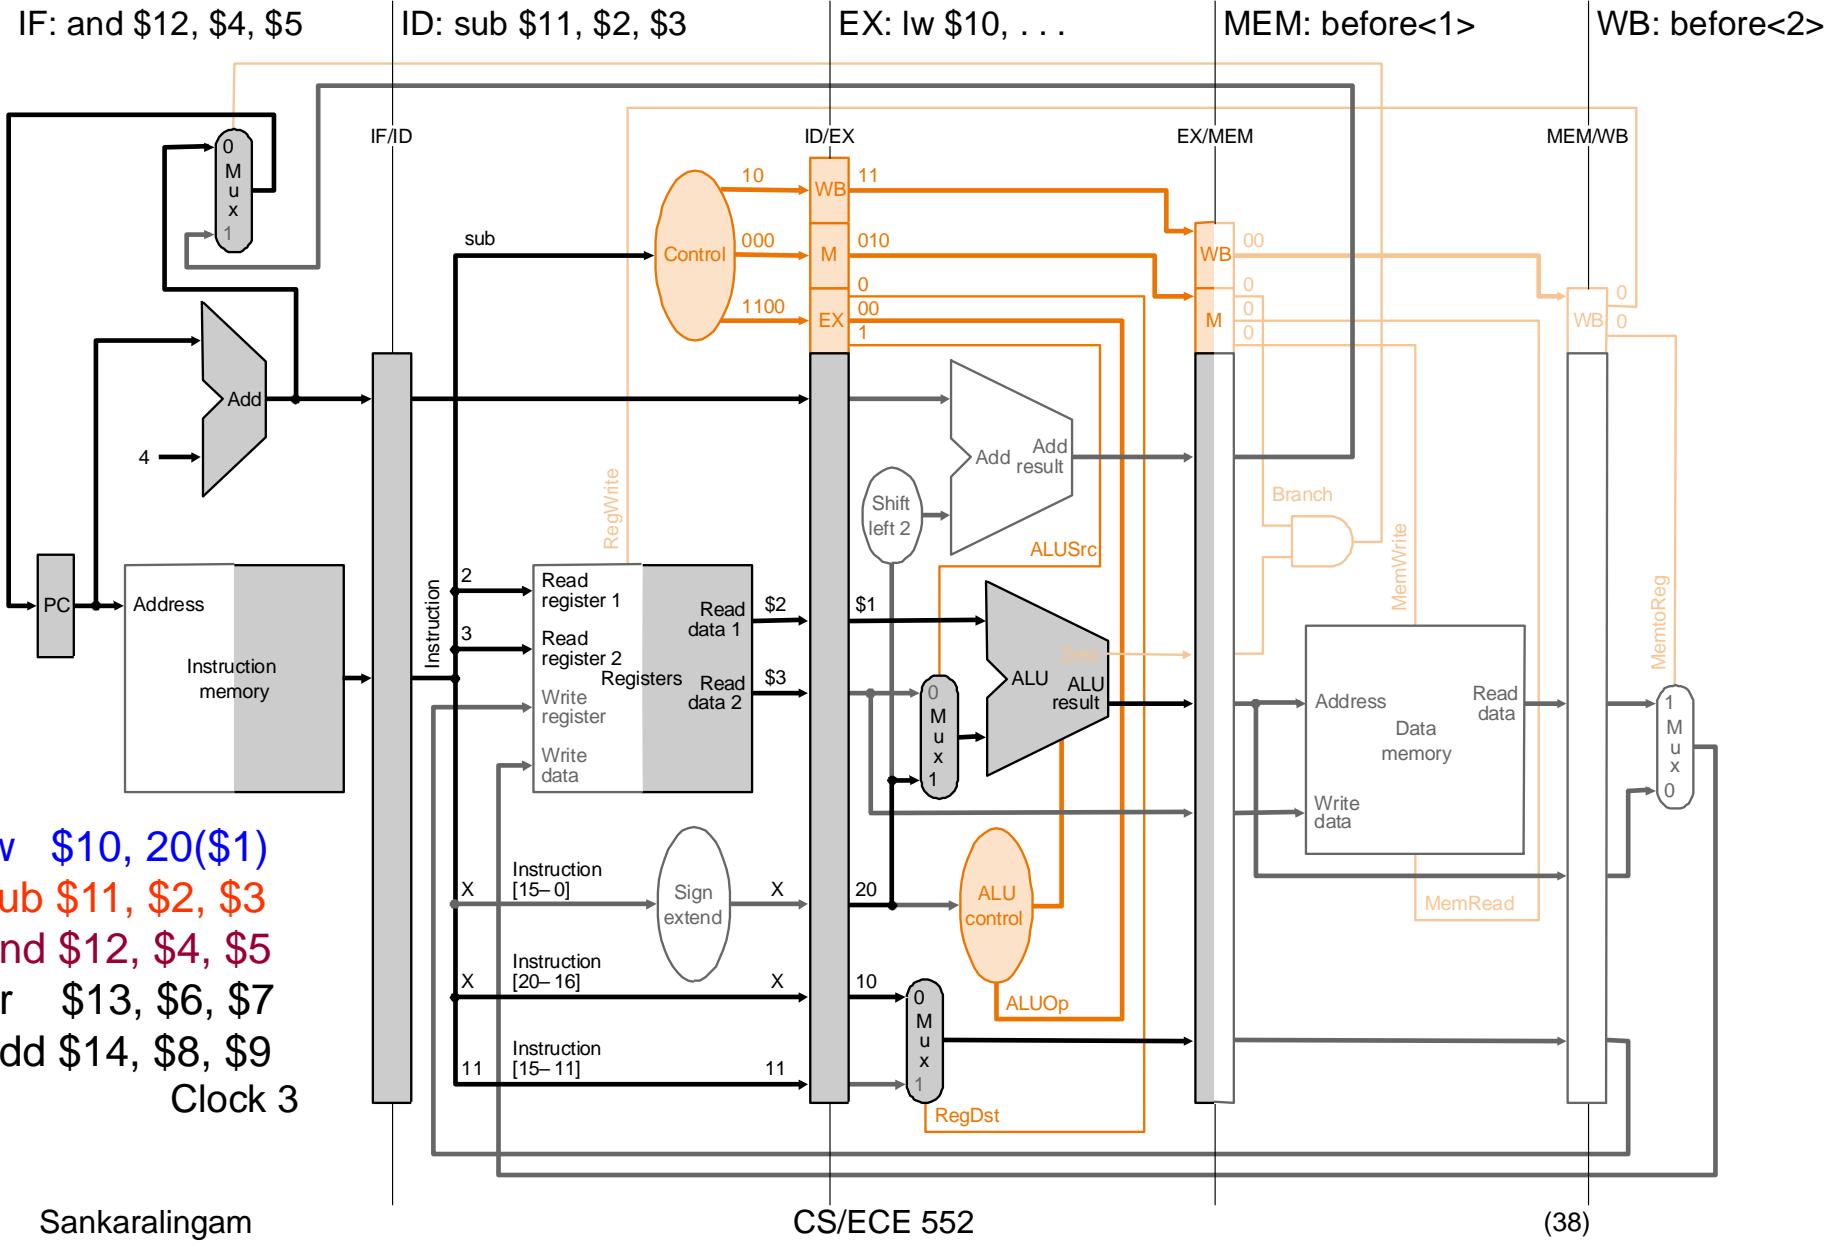
ID: lw \$10, 20(\$1)

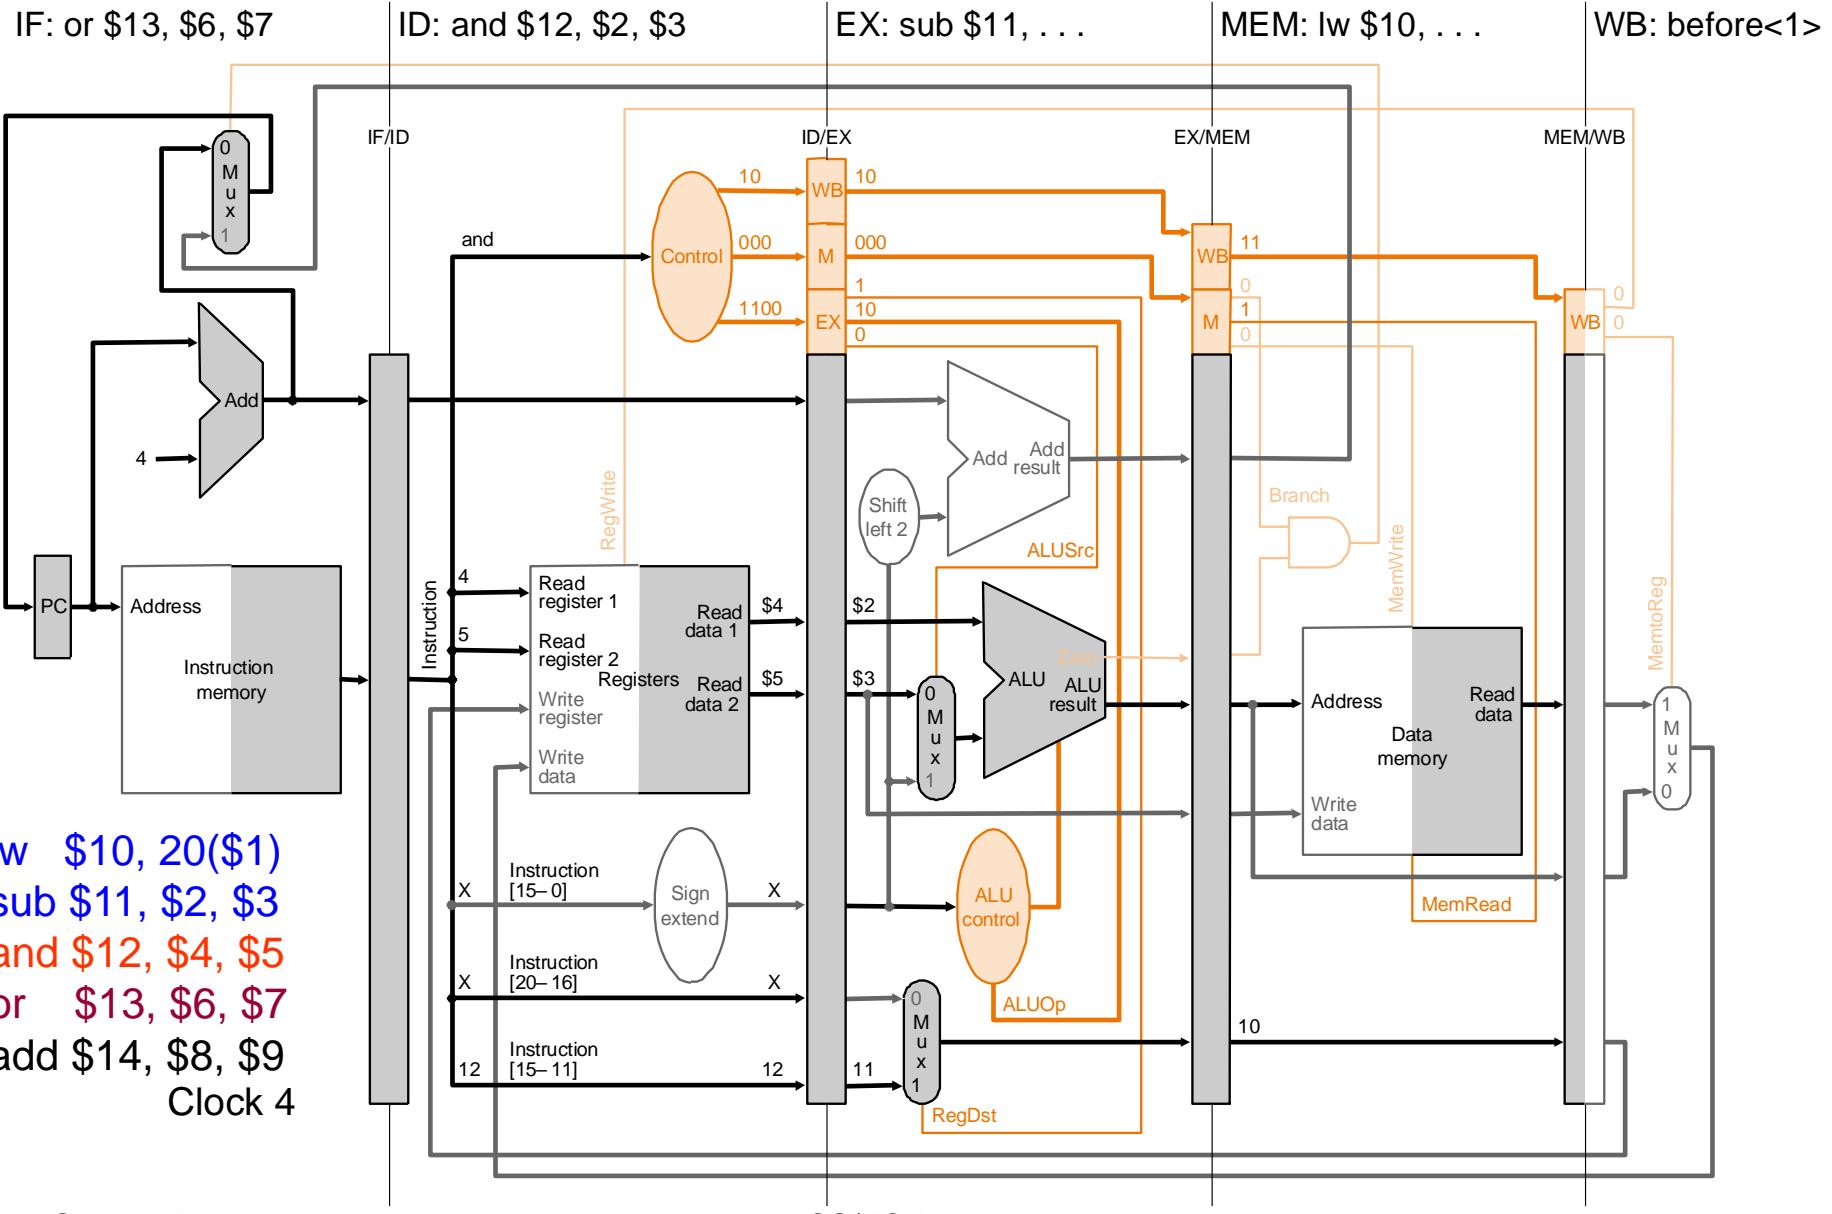
EX: before<1>

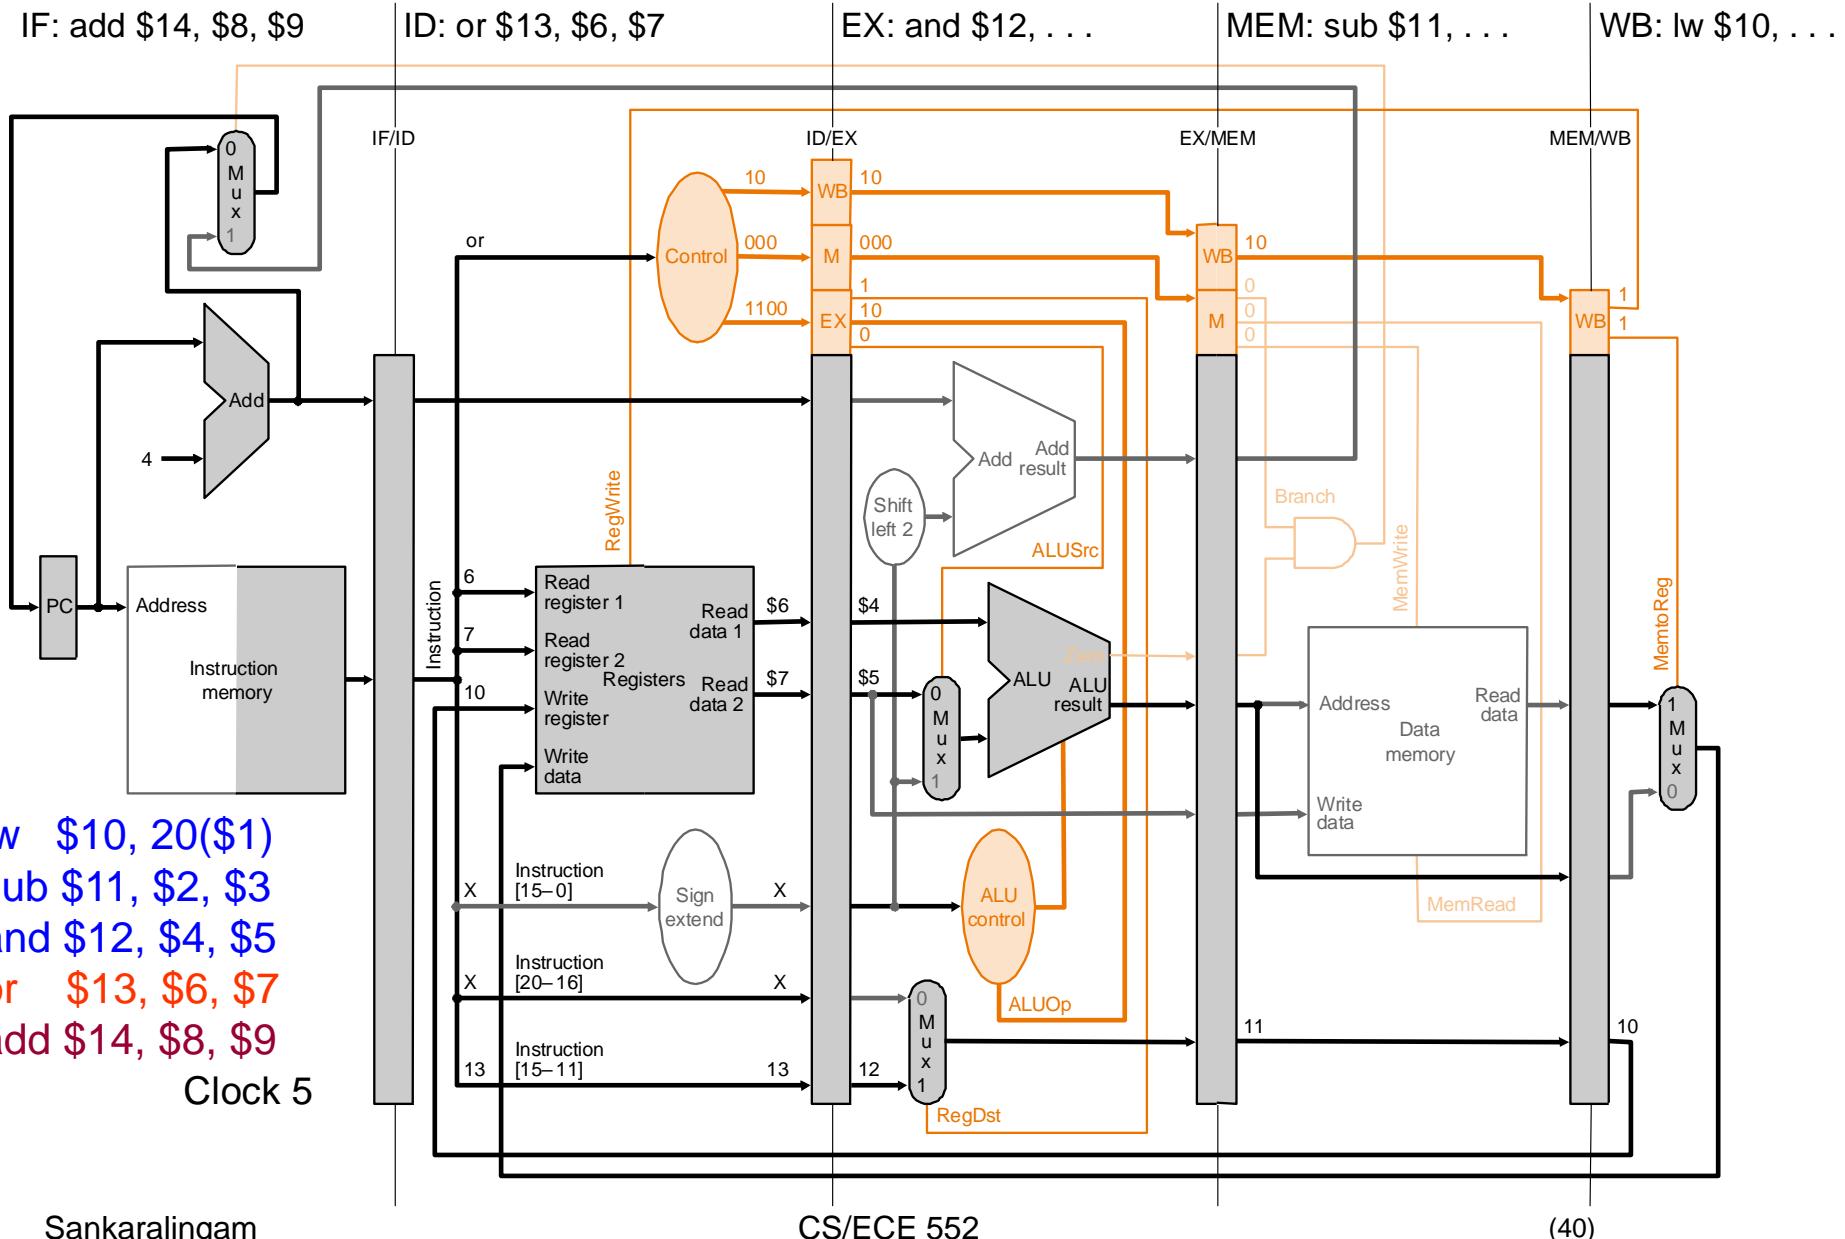
MEM: before<2>

WB: before<3>

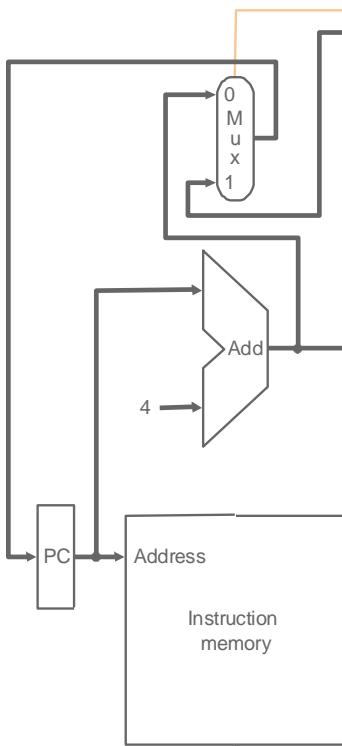




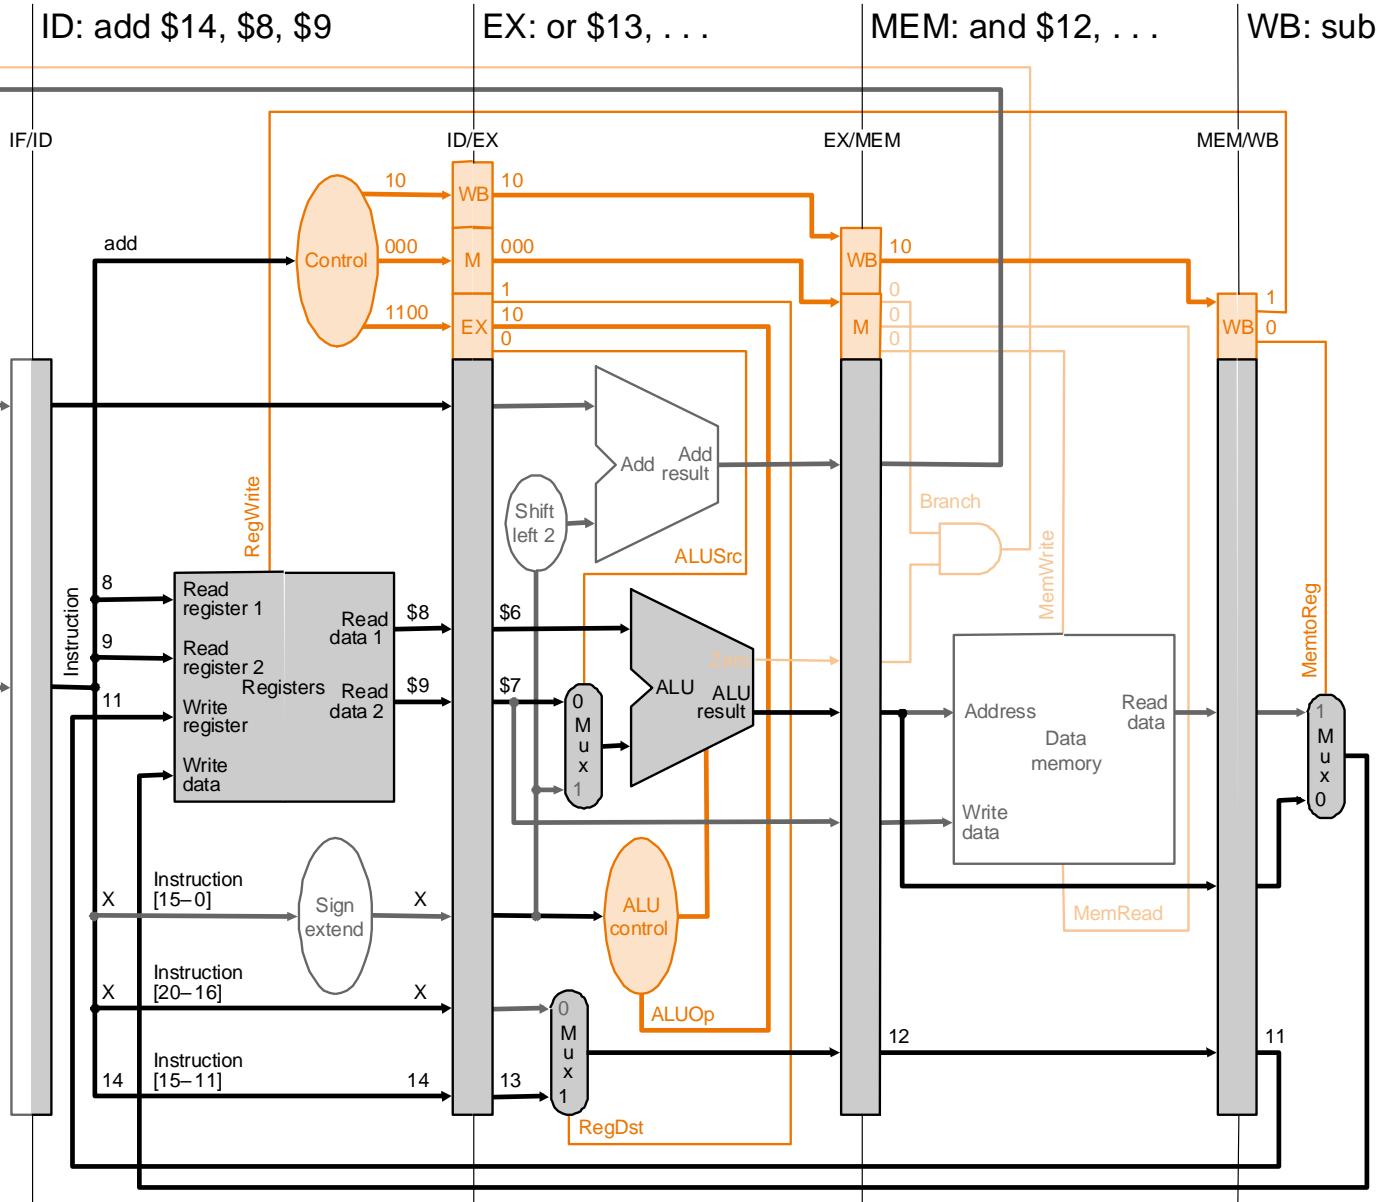




IF: after<1>



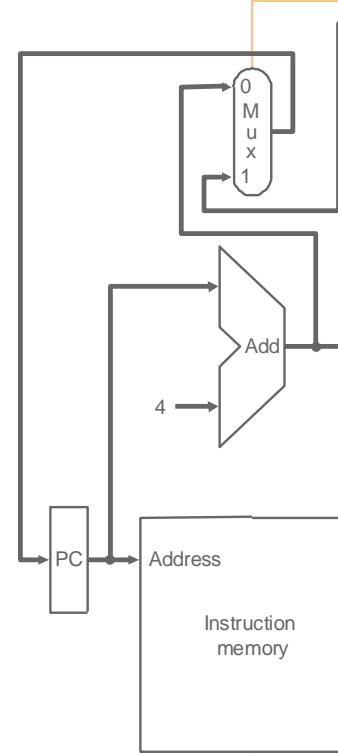
ID: add \$14, \$8, \$9



Iw \$10, 20(\$1)
sub \$11, \$2, \$3
and \$12, \$4, \$5
or \$13, \$6, \$7
add \$14, \$8, \$9

Clock 6

IF: after<2>



Iw \$10, 20(\$1)
sub \$11, \$2, \$3
and \$12, \$4, \$5
or \$13, \$6, \$7
add \$14, \$8, \$9

Clock 7

ID: after<1>

IF/ID

ID

RegWrite

Instruction

12

Sign extend

Instruction [20-16]

Instruction [15-11]

CS/ECE 552

EX: add \$14, ...

ID/EX

WB

00

000

M

10

0

1

0

Mux

1

0

RegDst

14

(42)

MEM: or \$13, ...

EX/MEM

WB

0

0

M

0

0

Branch

MemWrite

13

12

WB: and \$12,

MEM/WB

WB

0

1

WB

0

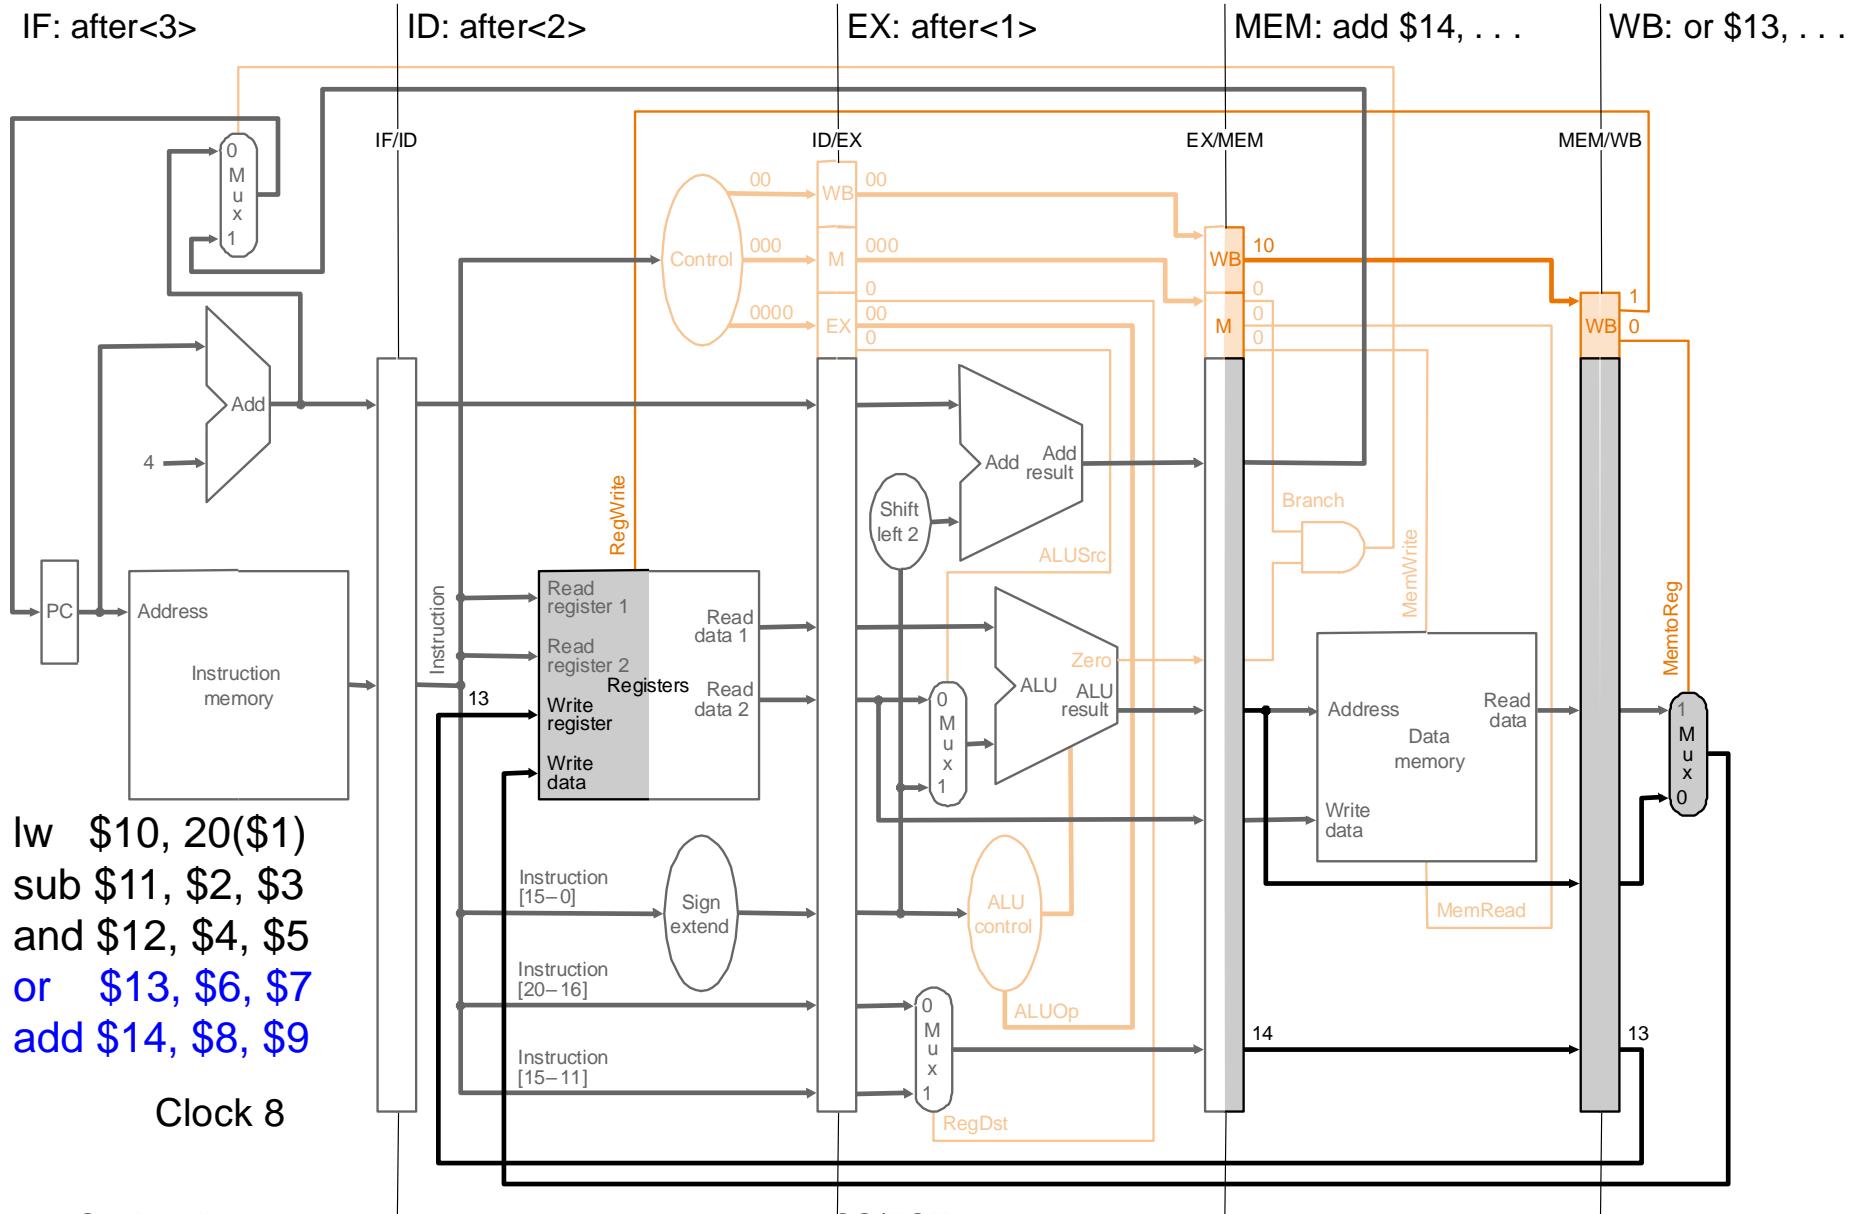
1

Mux

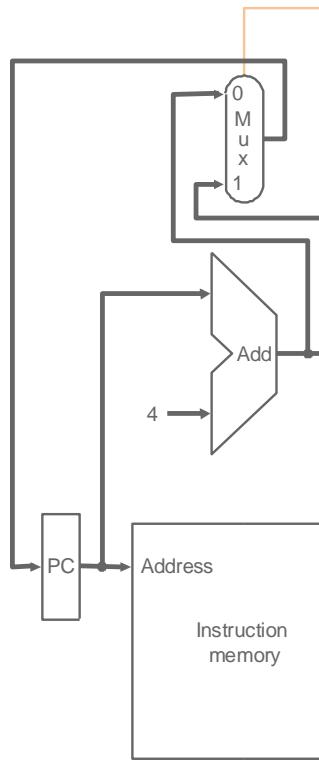
0

1

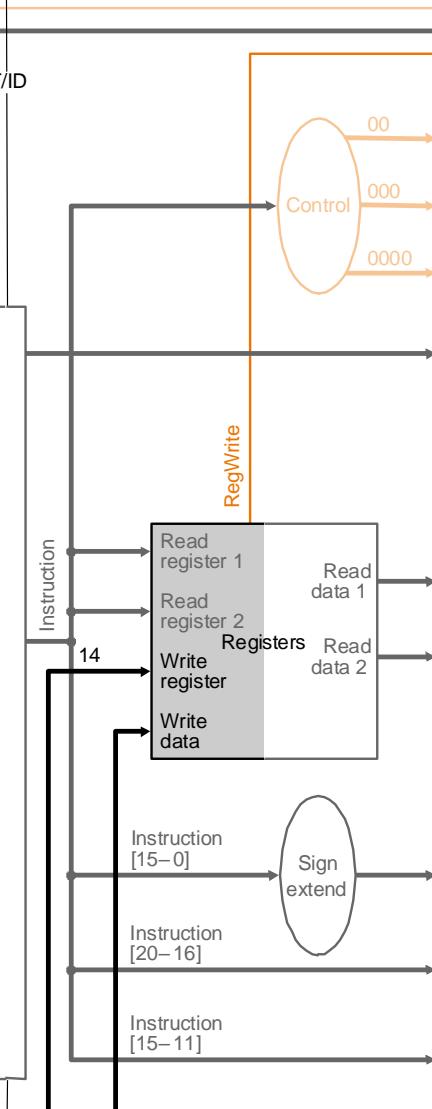
MemReg



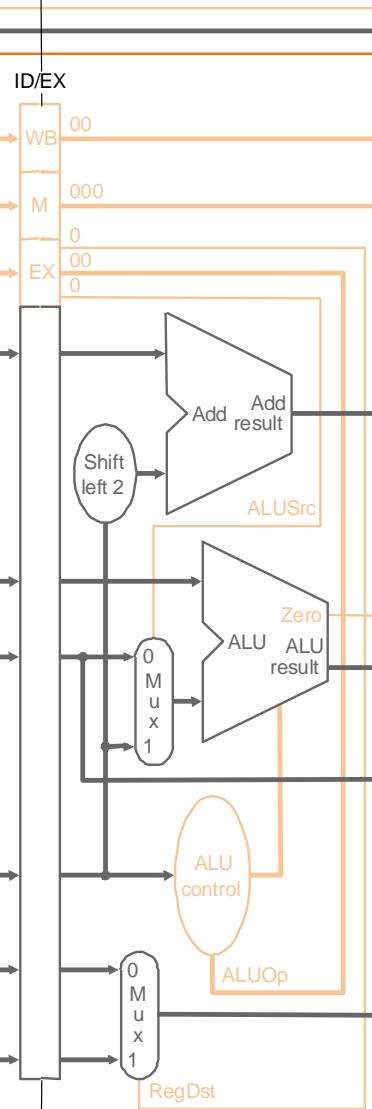
IF: after<4>



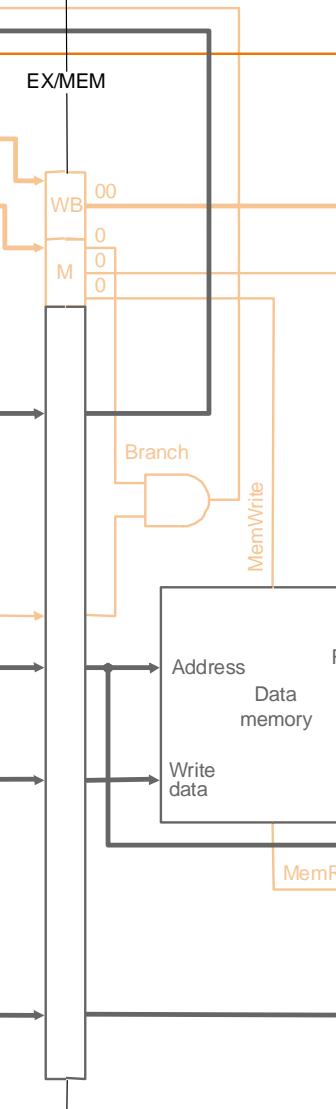
ID: after<3>



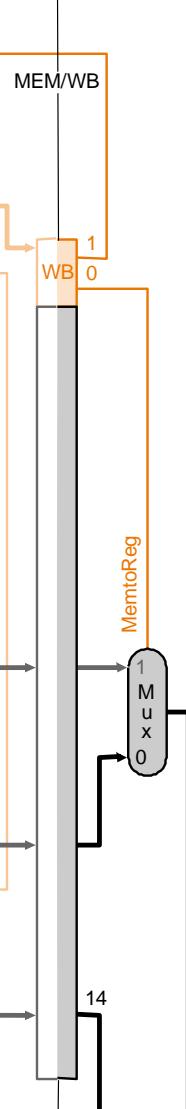
EX: after<2>



MEM: after<1>



WB: add \$14, .



Iw \$10, 20(\$1)
sub \$11, \$2, \$3
and \$12, \$4, \$5
or \$13, \$6, \$7
add \$14, \$8, \$9

Clock 9

Dependence : Backward in time



THE UNIVERSITY
of
WISCONSIN
MADISON

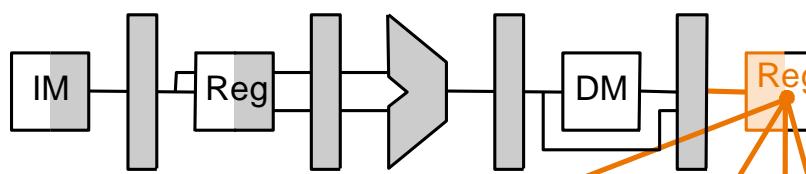
Time (in clock cycles)

Value of register \$2:	CC 1	CC 2	CC 3	CC 4	CC 5	CC 6	CC 7	CC 8	CC 9
	10	10	10	10	10/-20	-20	-20	-20	-20

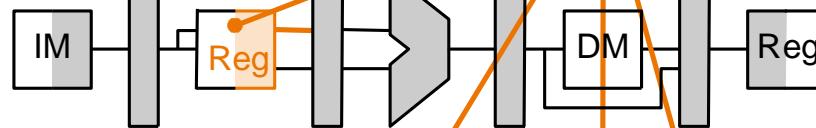
Program execution order

(in instructions)

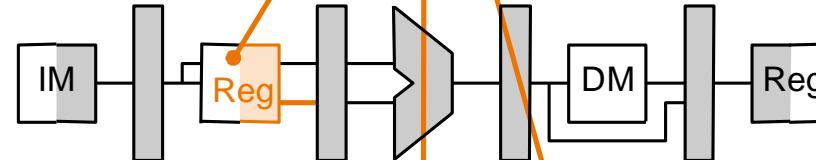
sub \$2, \$1, \$3



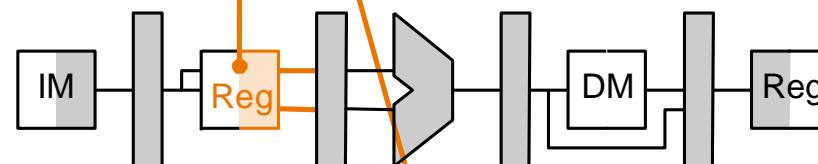
and \$12, \$2, \$5



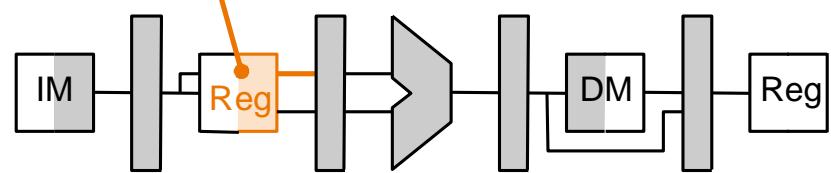
or \$13, \$6, \$2



add \$14, \$2, \$2



sw \$15, 100(\$2)



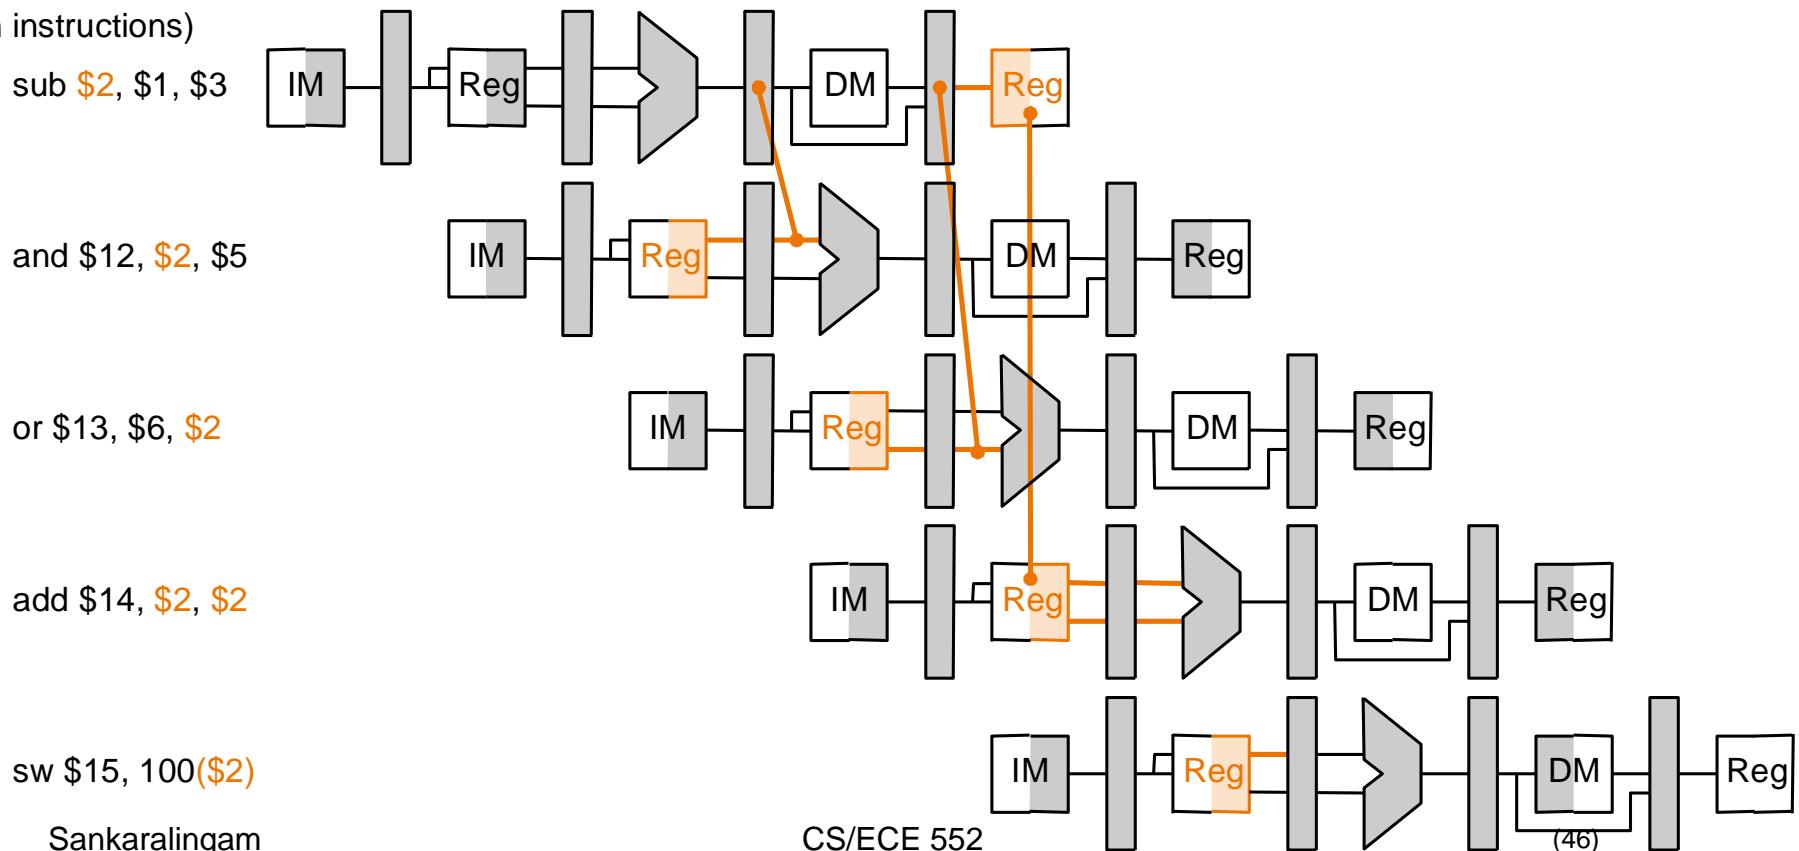
True dependence : Forward in time



THE UNIVERSITY
of
WISCONSIN
MADISON

	Time (in clock cycles)								
	CC 1	CC 2	CC 3	CC 4	CC 5	CC 6	CC 7	CC 8	CC 9
Value of register \$2 :	10	10	10	10	10/-20	-20	-20	-20	-20
Value of EX/MEM :	X	X	X	-20	X	X	X	X	X
Value of MEM/WB :	X	X	X	X	-20	X	X	X	X

Program
execution order
(in instructions)



Data Hazards and Forwarding: Walkthrough



- Code snippet
 - identify hazards
 - identify forwarding paths
- NOTE:
 - DATA DEPENDENCES ON REG \$2, REG \$4
 - ALSO TWO INSTRUCTIONS WRITE TO REG \$4

sub \$2, \$1, \$3
and \$4, \$2, \$5
or \$4, \$4, \$2
add \$9, \$4, \$2

Walkthrough

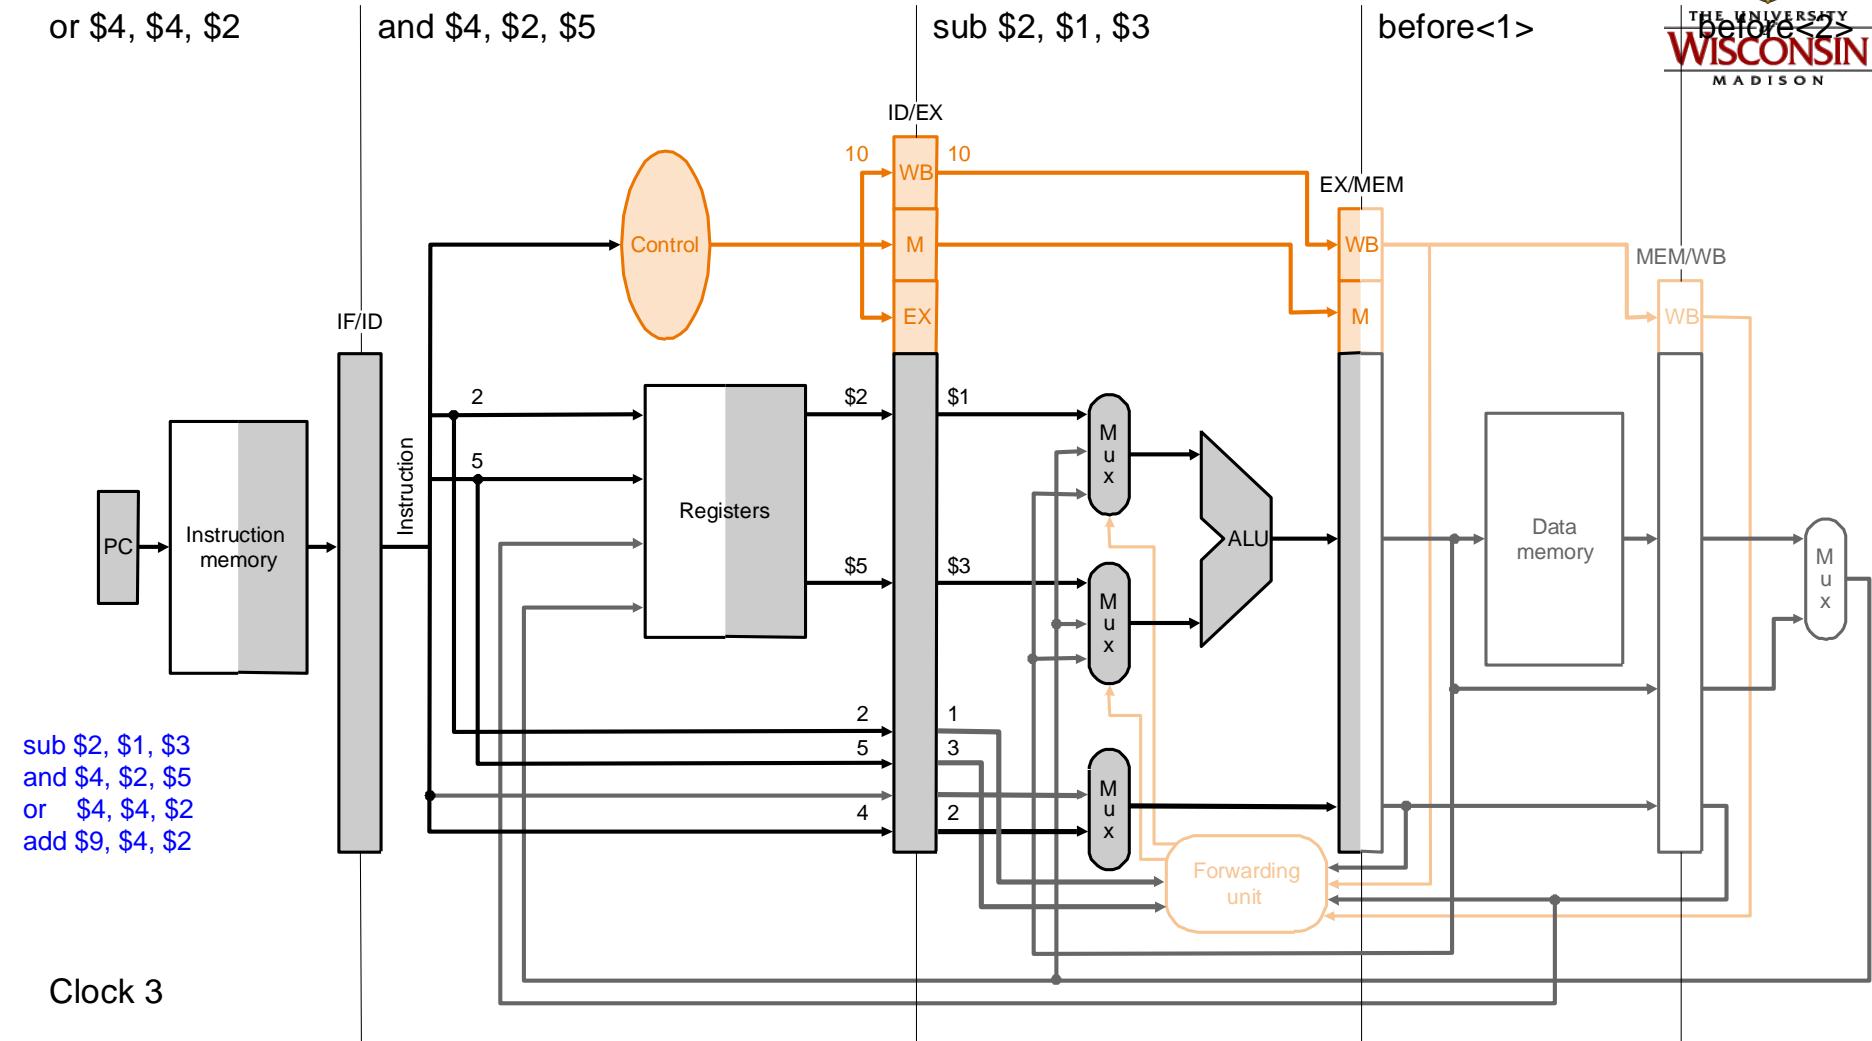


or \$4, \$4, \$2

and \$4, \$2, \$5

sub \$2, \$1, \$3

| before<1>



- Skip the boring stuff, jump to cycle 3



add \$9, \$4, \$2

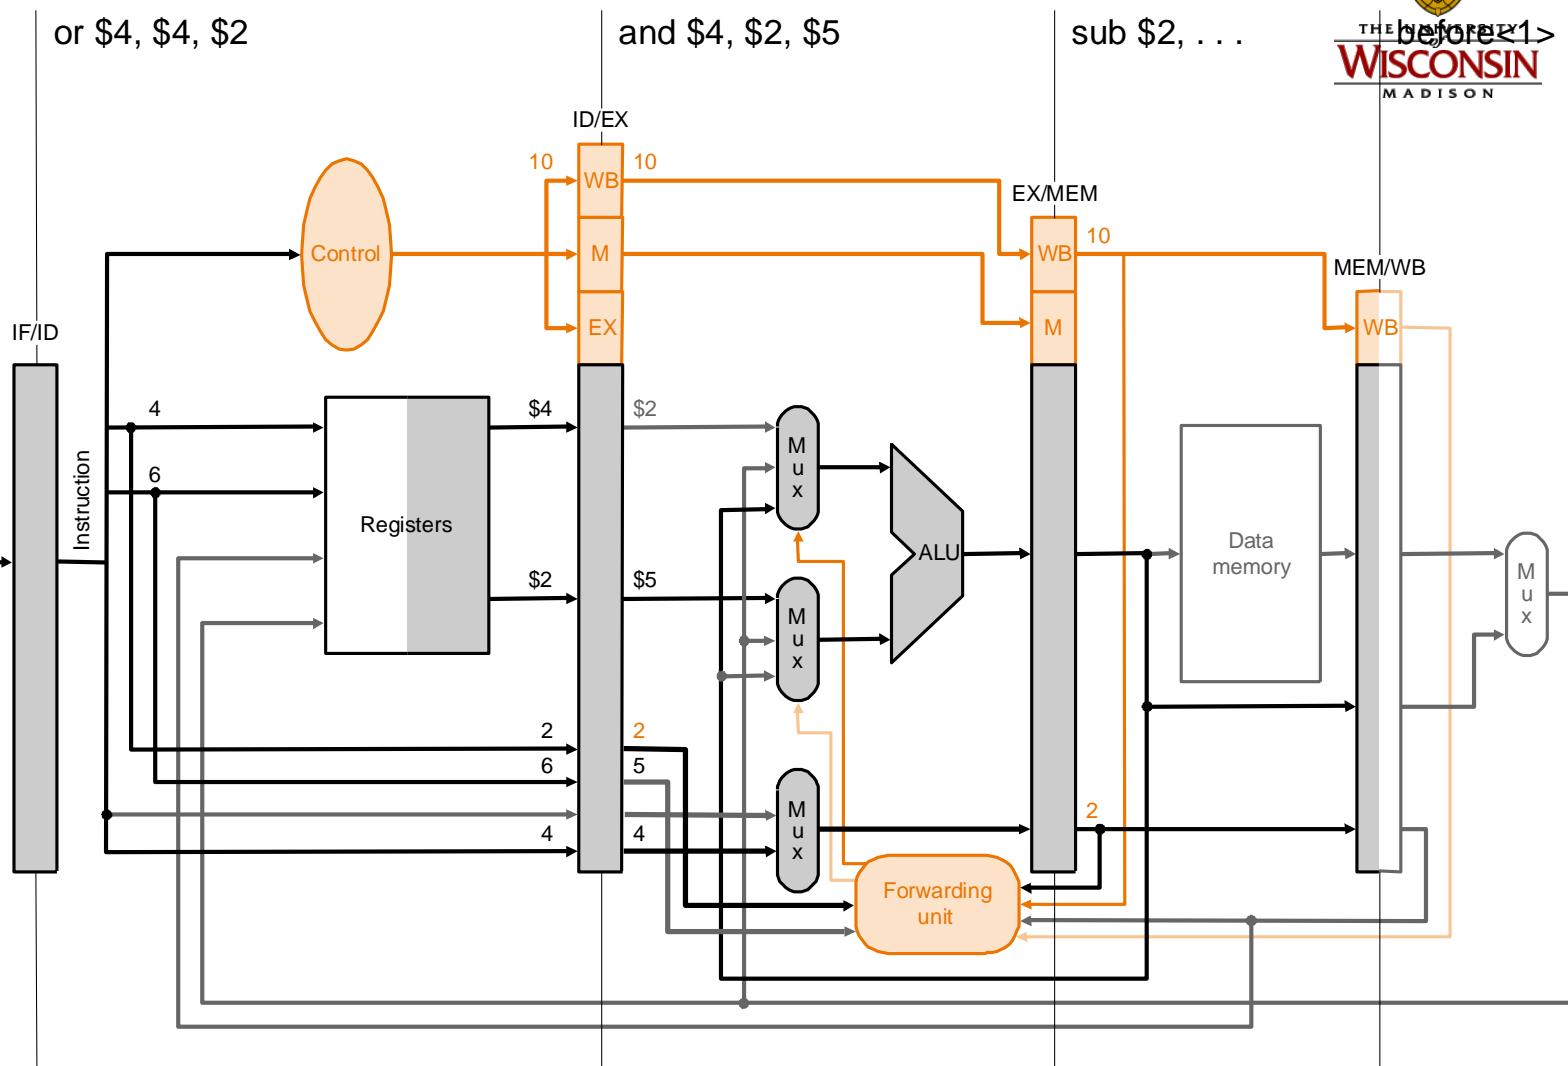
or \$4, \$4, \$2

and \$4, \$2, \$5

sub \$2, ...

sub \$2, \$1, \$3
and \$4, \$2, \$5
or \$4, \$4, \$2
add \$9, \$4, \$2

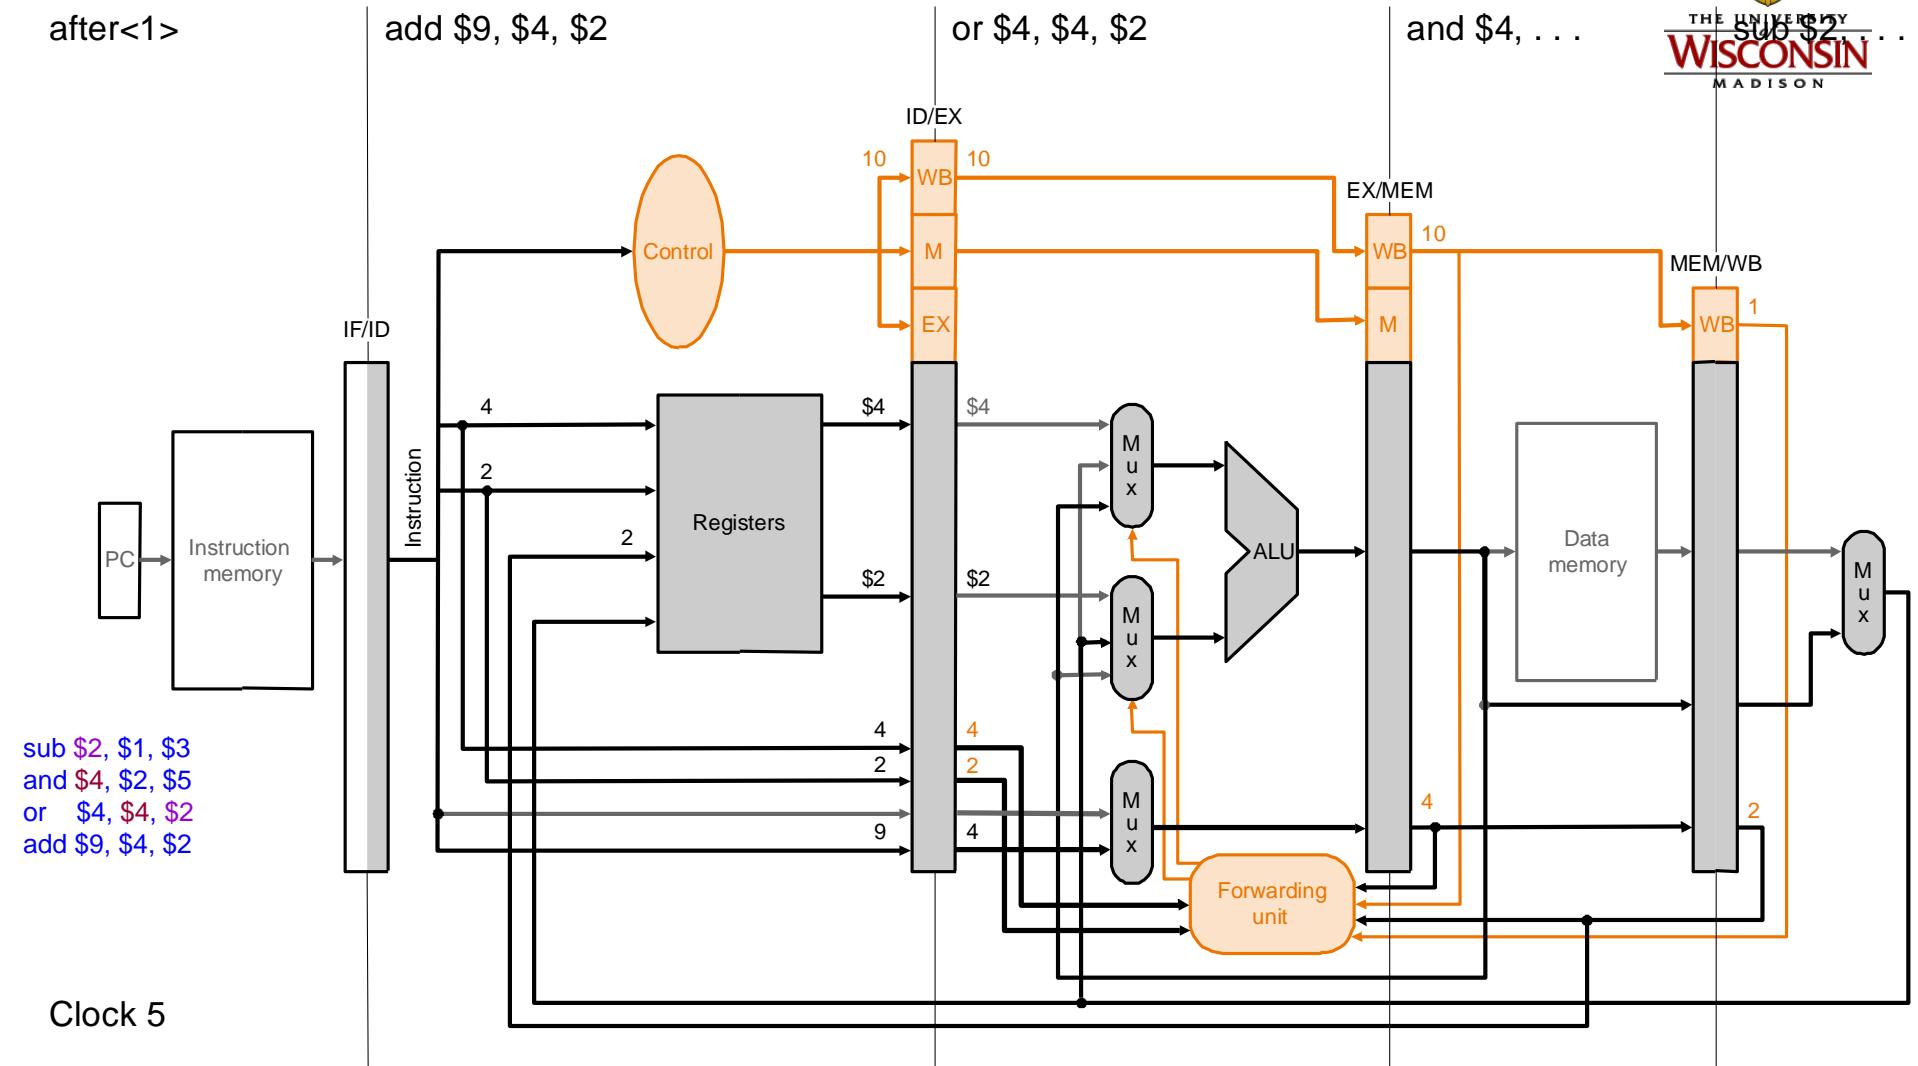
Clock 4



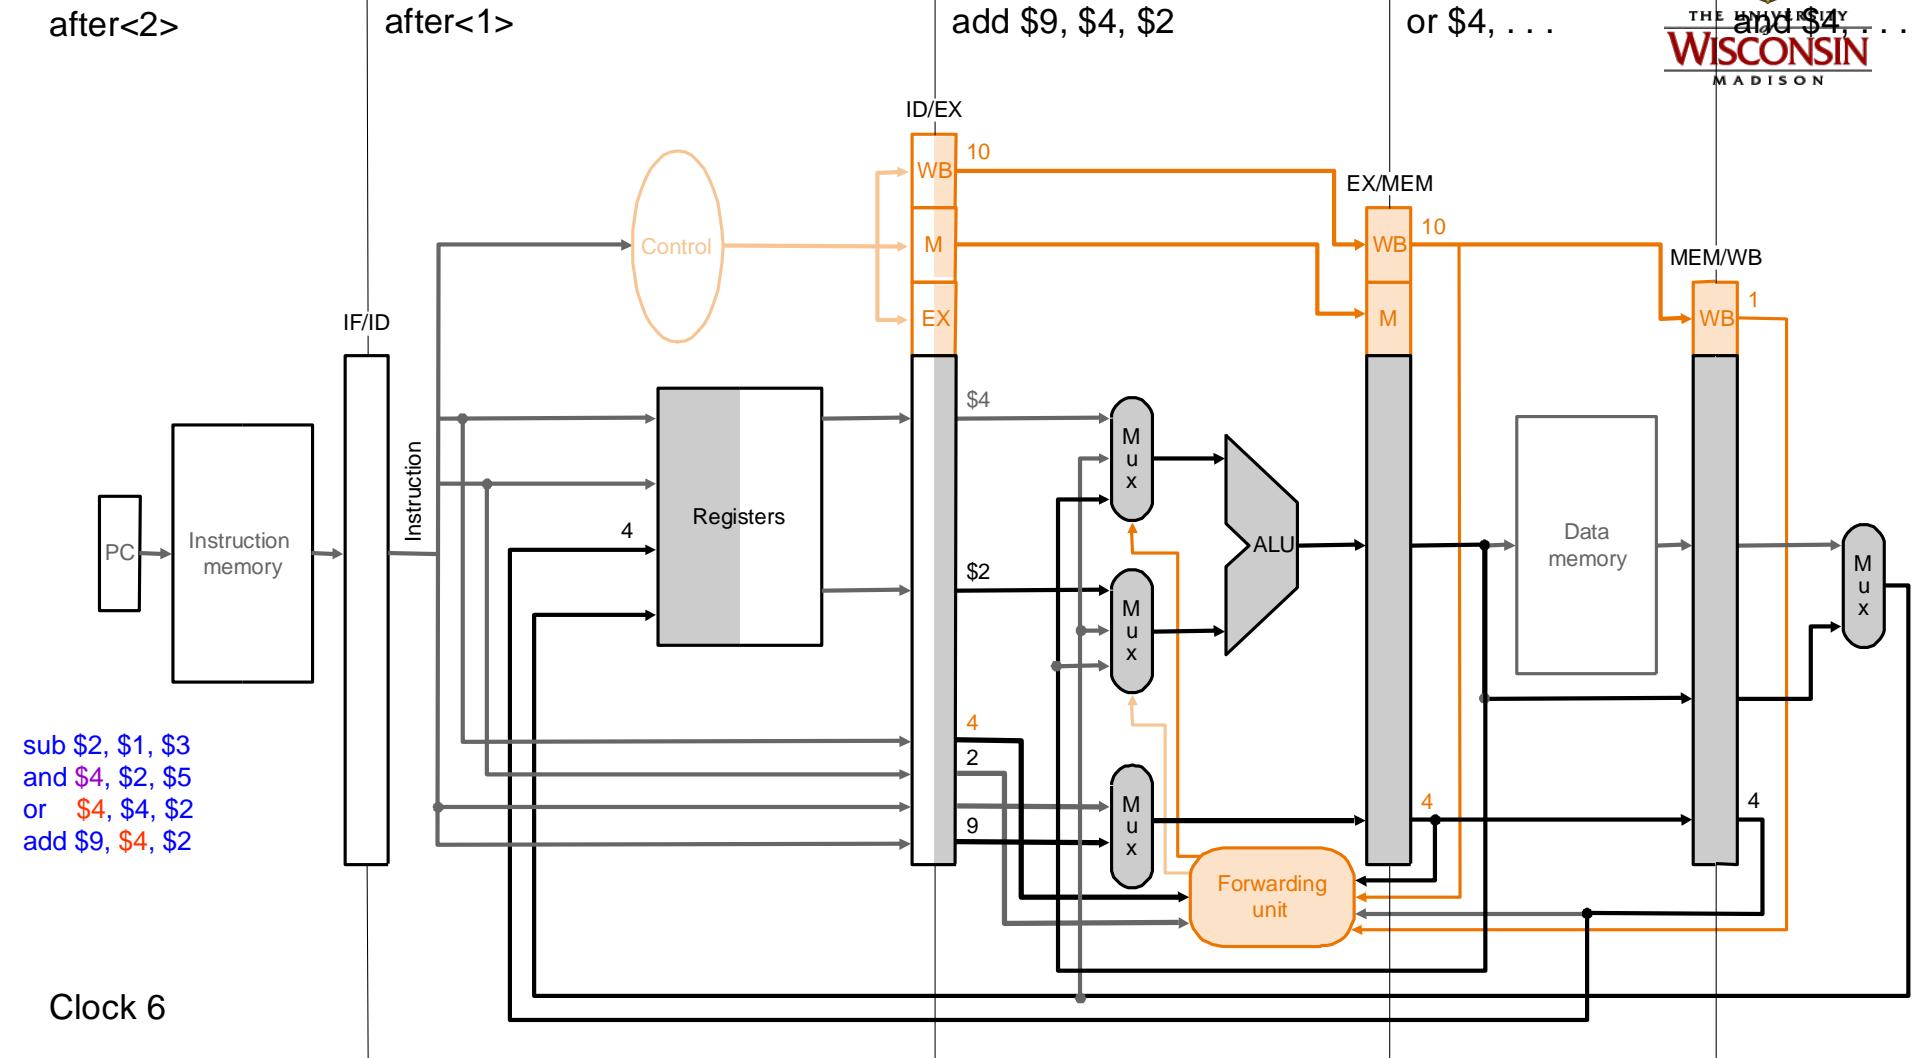
- Forward ALUOut to Operand 1



after<1>

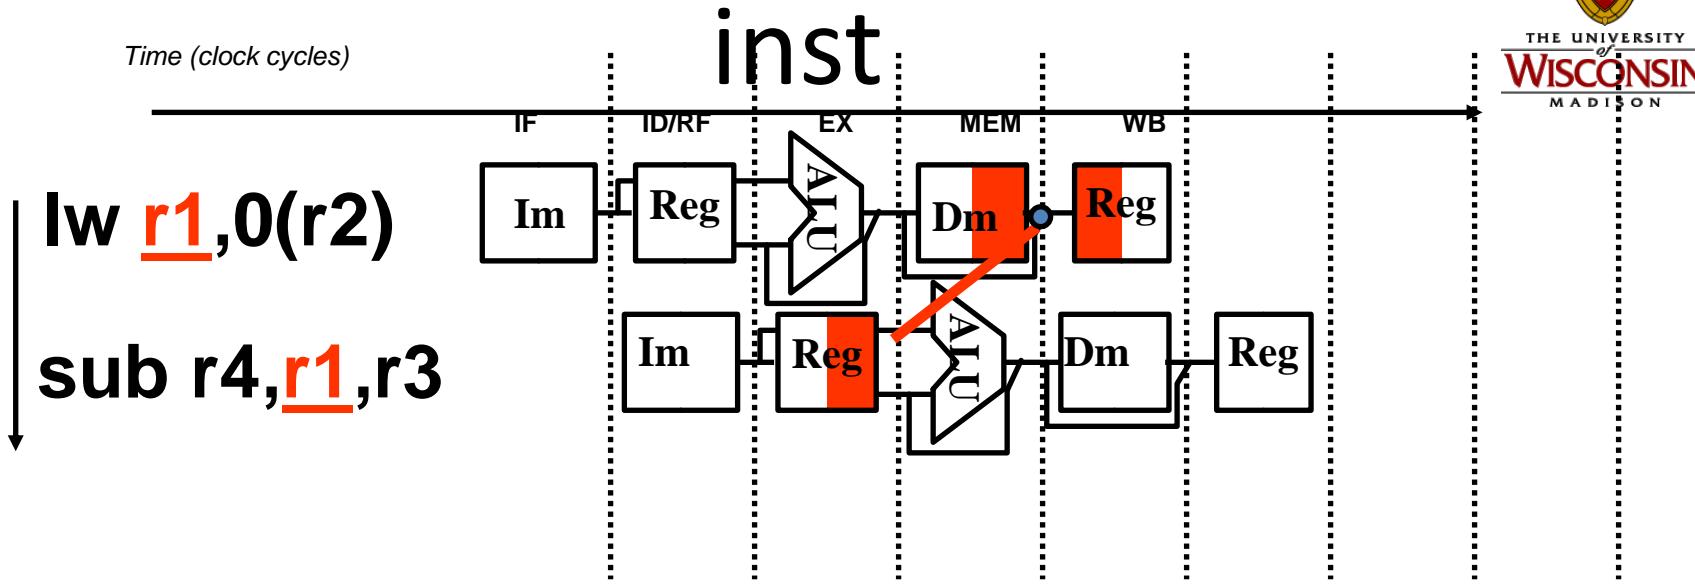


- Forward ALUout to Op1, Mem to Op2



- Two candidates match, forward the latest

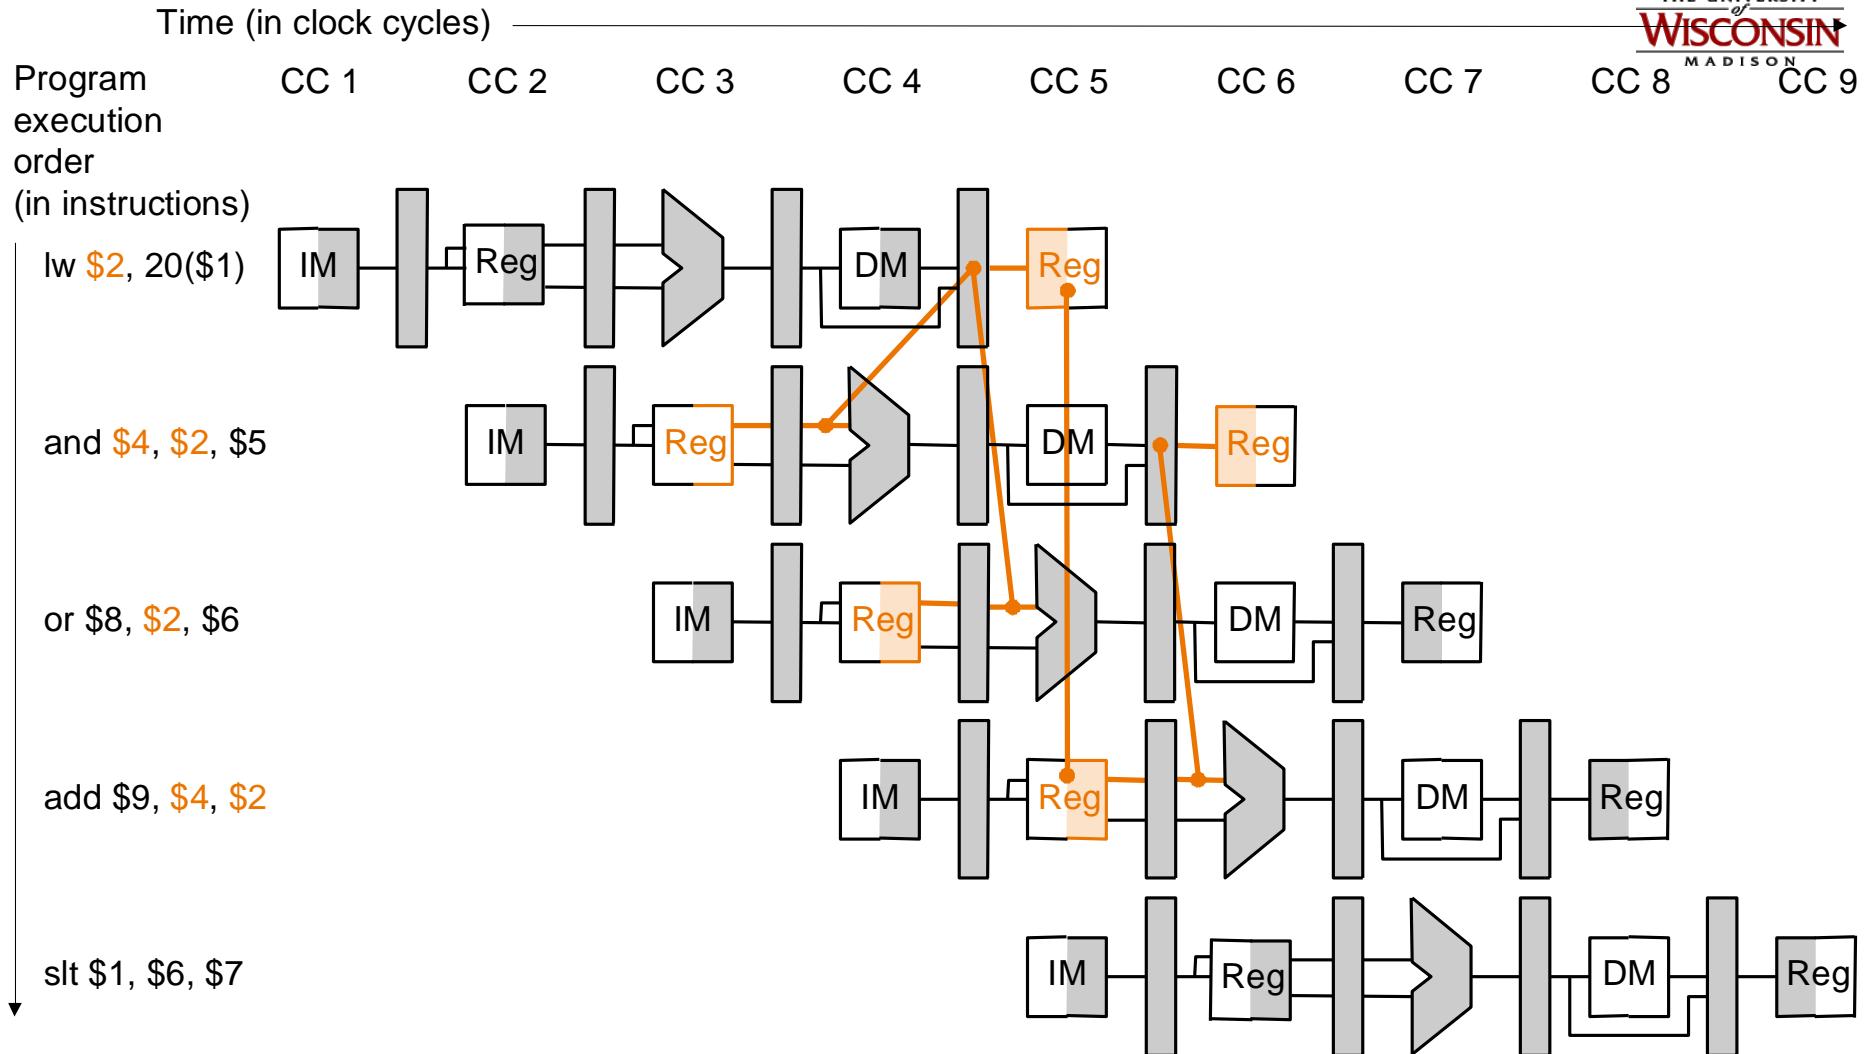
Lookahead: RAW hazard with load



- Forwarding as solution to RAW hazard
 - possible if no (true) dependence going backwards in time
 - True for R-type instructions
 - Data available after EX stage (i.e., at ALUOut)
 - Not true for load instruction



Load instruction



- Replaced “sub” with “lw” in previous code-example

Hazards with load instruction

- True dependencies: backward in time
- Stall the pipeline
- Minor change in terminology
 - If **forwarding** can solve it, it is not a hazard!
 - “**Hazard**” refers only to true backward dependencies in time.

Detection

- Conditions
 - Preceding instruction must read memory
 - MemRead must be asserted
 - Destination of preceding instruction (*rt*) must be one of operands of current instruction
- Logic equations— restate above conditions formally
 - If(ID/EX.MemRead AND

$$((ID/EX.RegRt = IF/ID.RegRs) \text{ OR } (ID/EX.RegRt = IF/ID.RegRt))$$

STALL

Iw \$2, 20(\$1)
and \$4, \$2, \$5
or \$4, \$4, \$2
add \$9, \$4, \$2

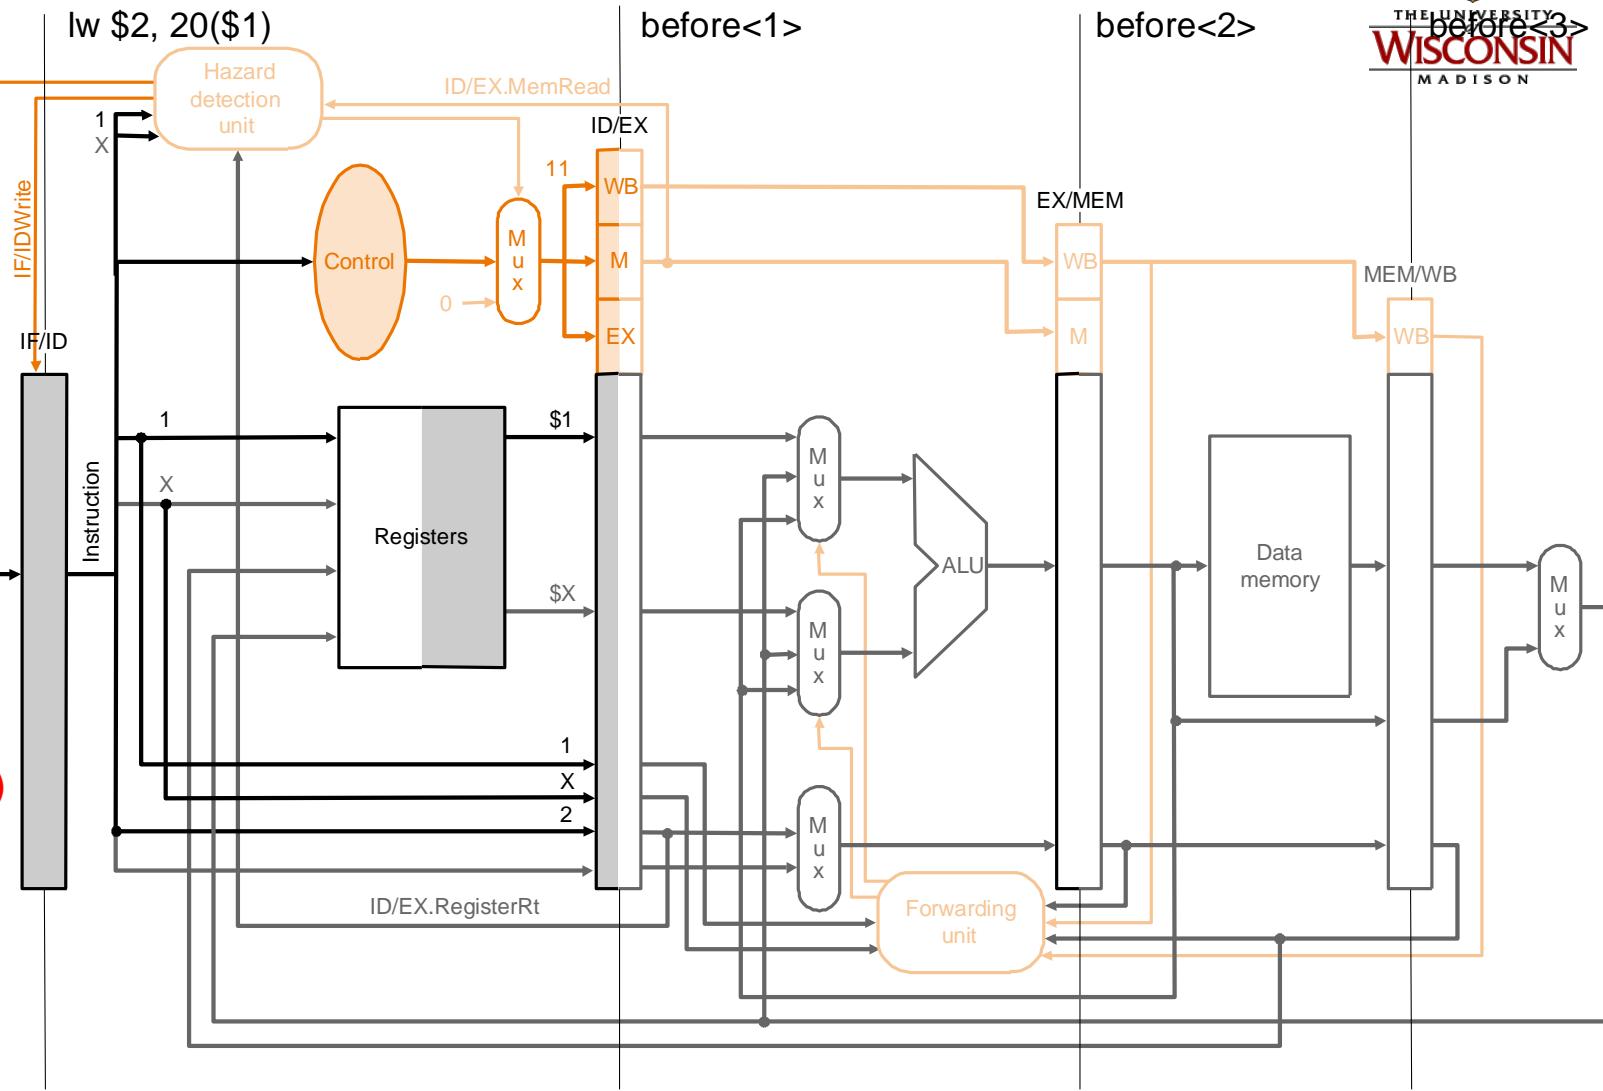
Stalling the pipeline

- Instruction cannot proceed
 - Following instruction must be stalled too.
 - Otherwise state in pipeline registers is overwritten
- Preceding instructions may proceed as usual
- Solution
 - inject NOP into EX/Mem pipeline
 - Prevent writes to PC to IF/ID register



Walk-through (1 of 6)

and \$4, \$2, \$5



Iw \$2, 20(\$1)
and \$4, \$2, \$5
or \$4, \$4, \$2
add \$9, \$4, \$2

Clock 2

- Skip to cycle 2

Walk-through (2 of 6)

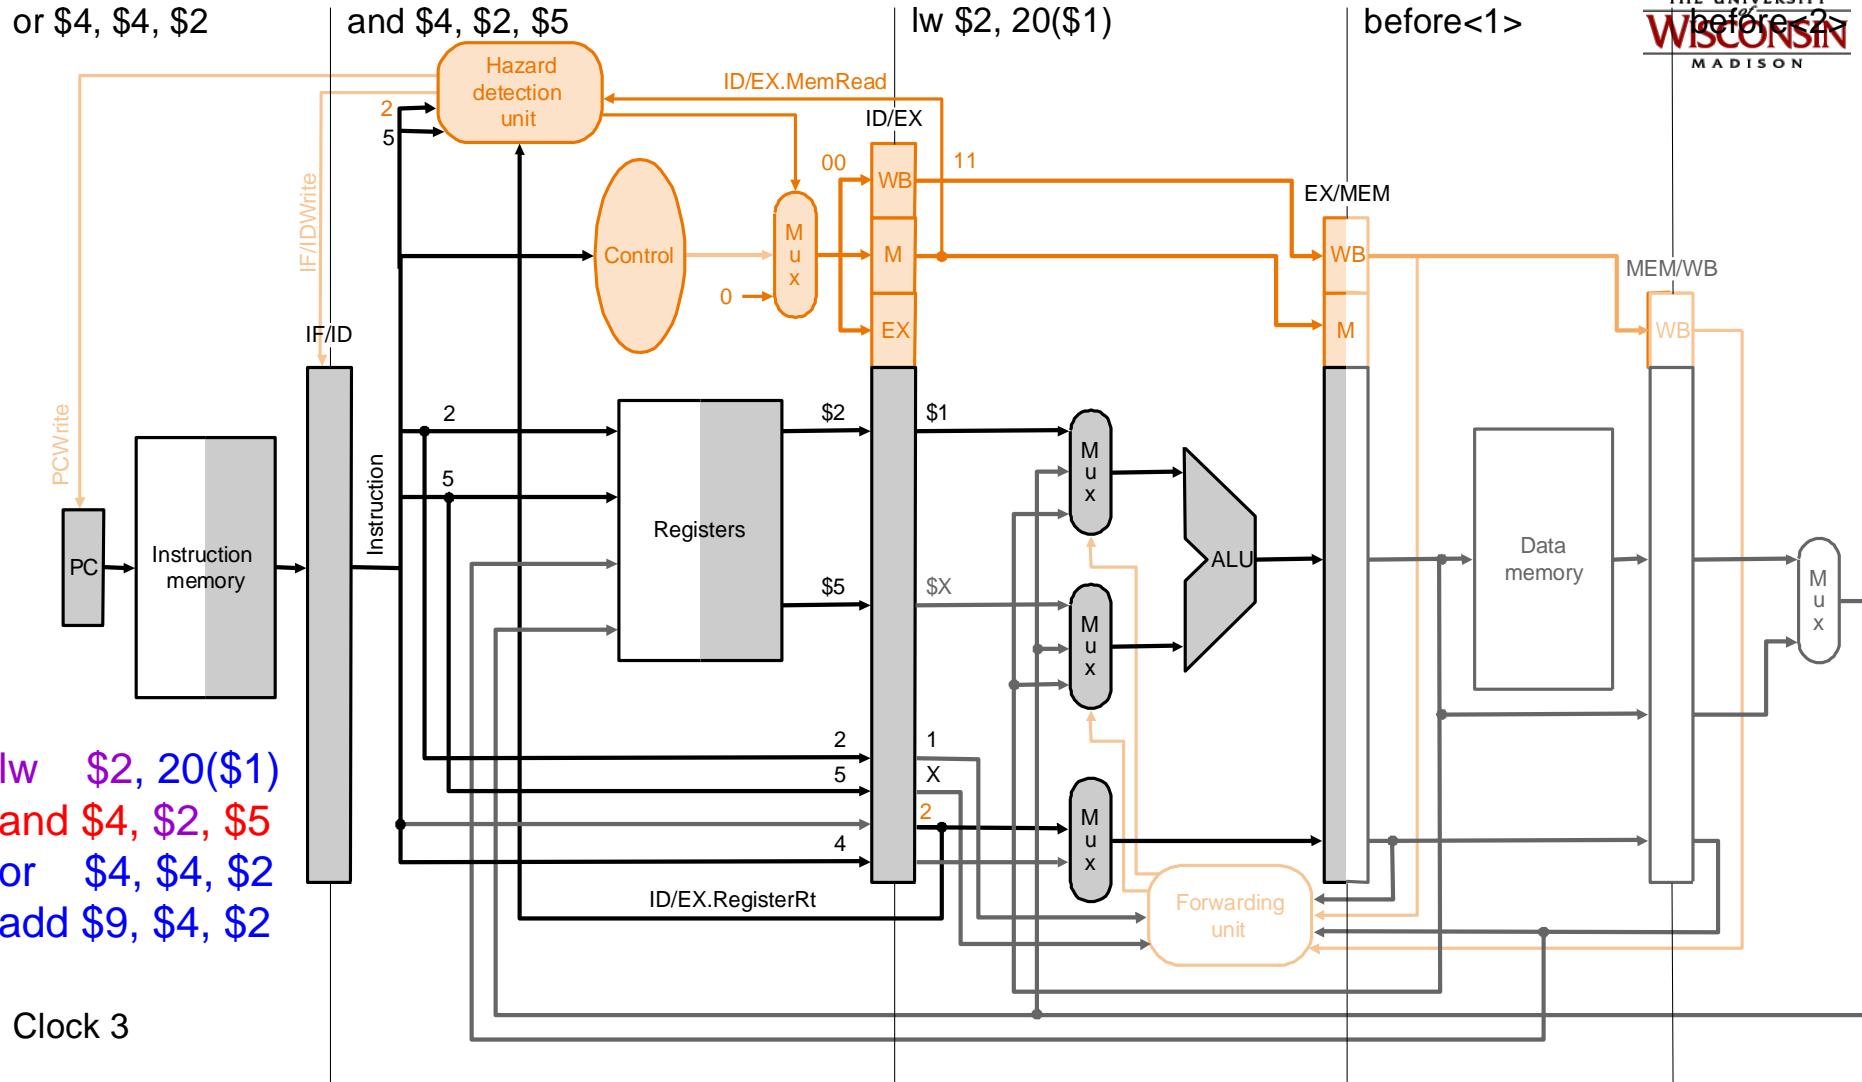


or \$4, \$4, \$2

and \$4, \$2, \$5

| Iw \$2, 20(\$1)

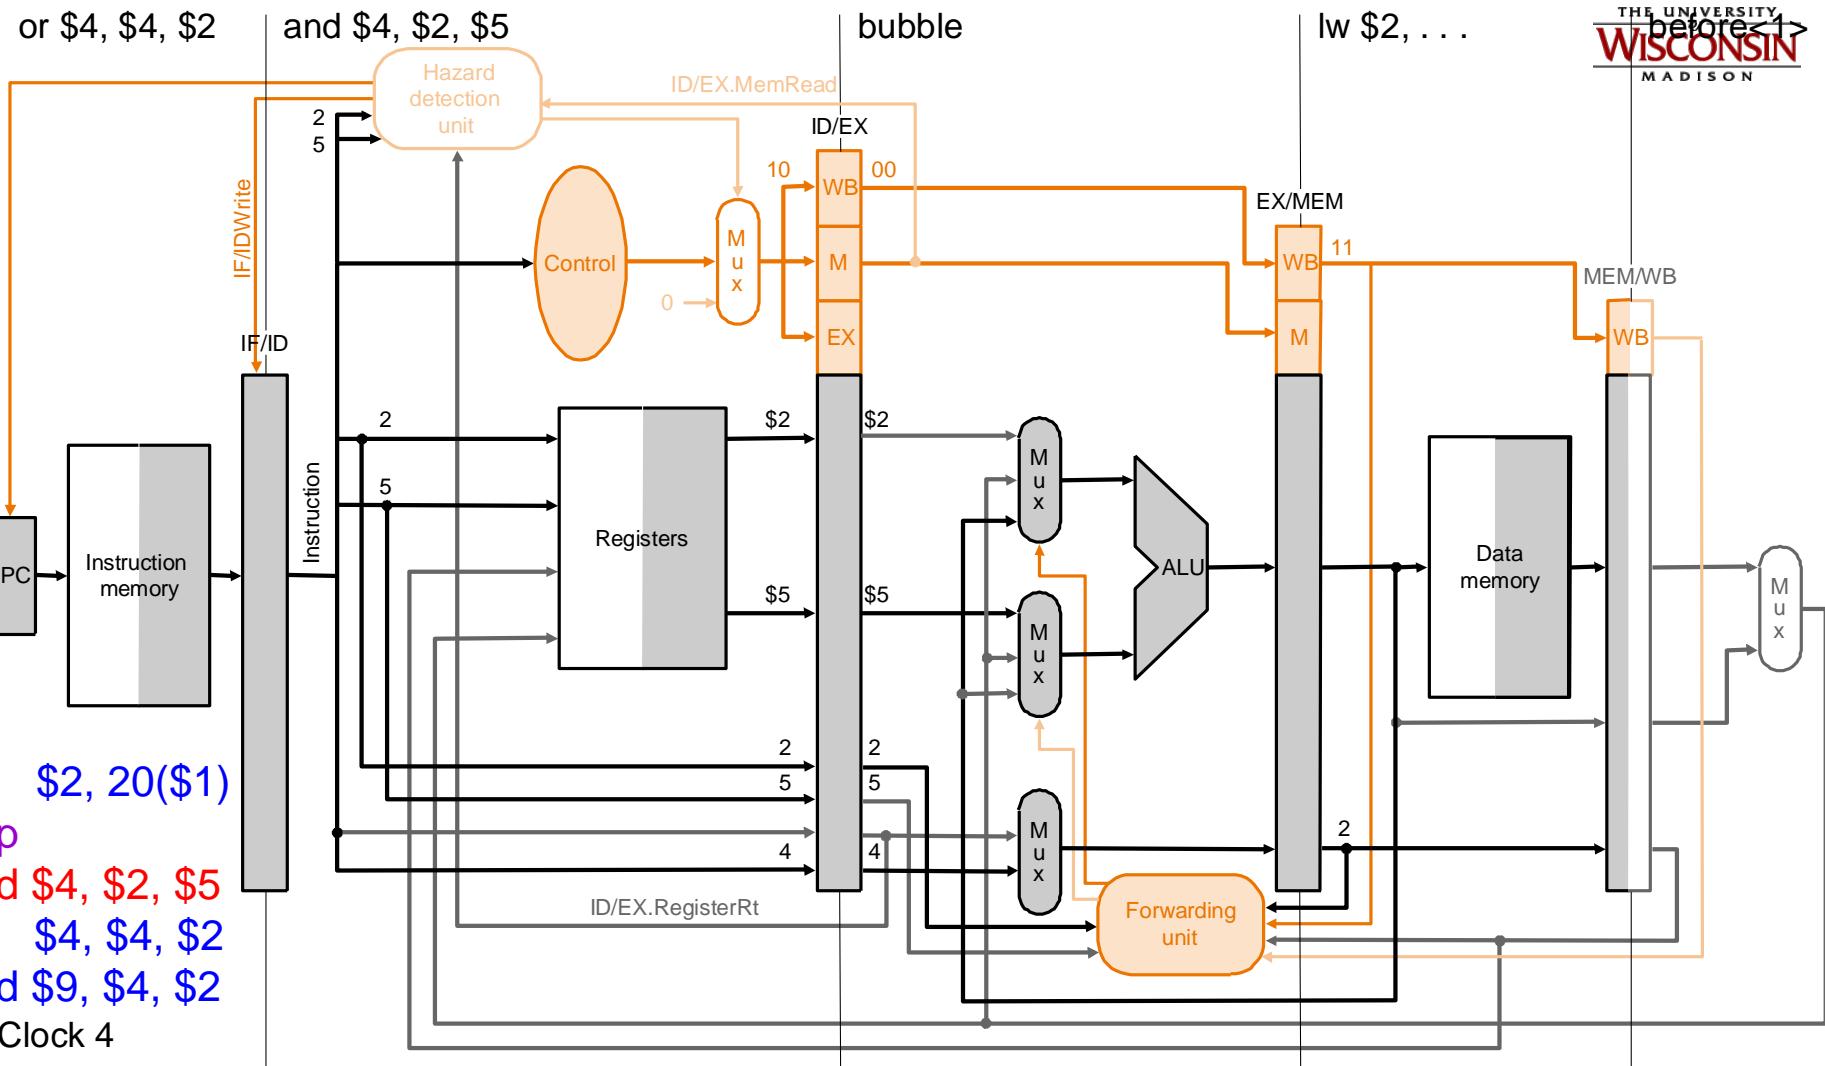
| before<1>



- All '0's => NOP (MemWr, RegWr, deasserted)



Walk-through (3 of 6)





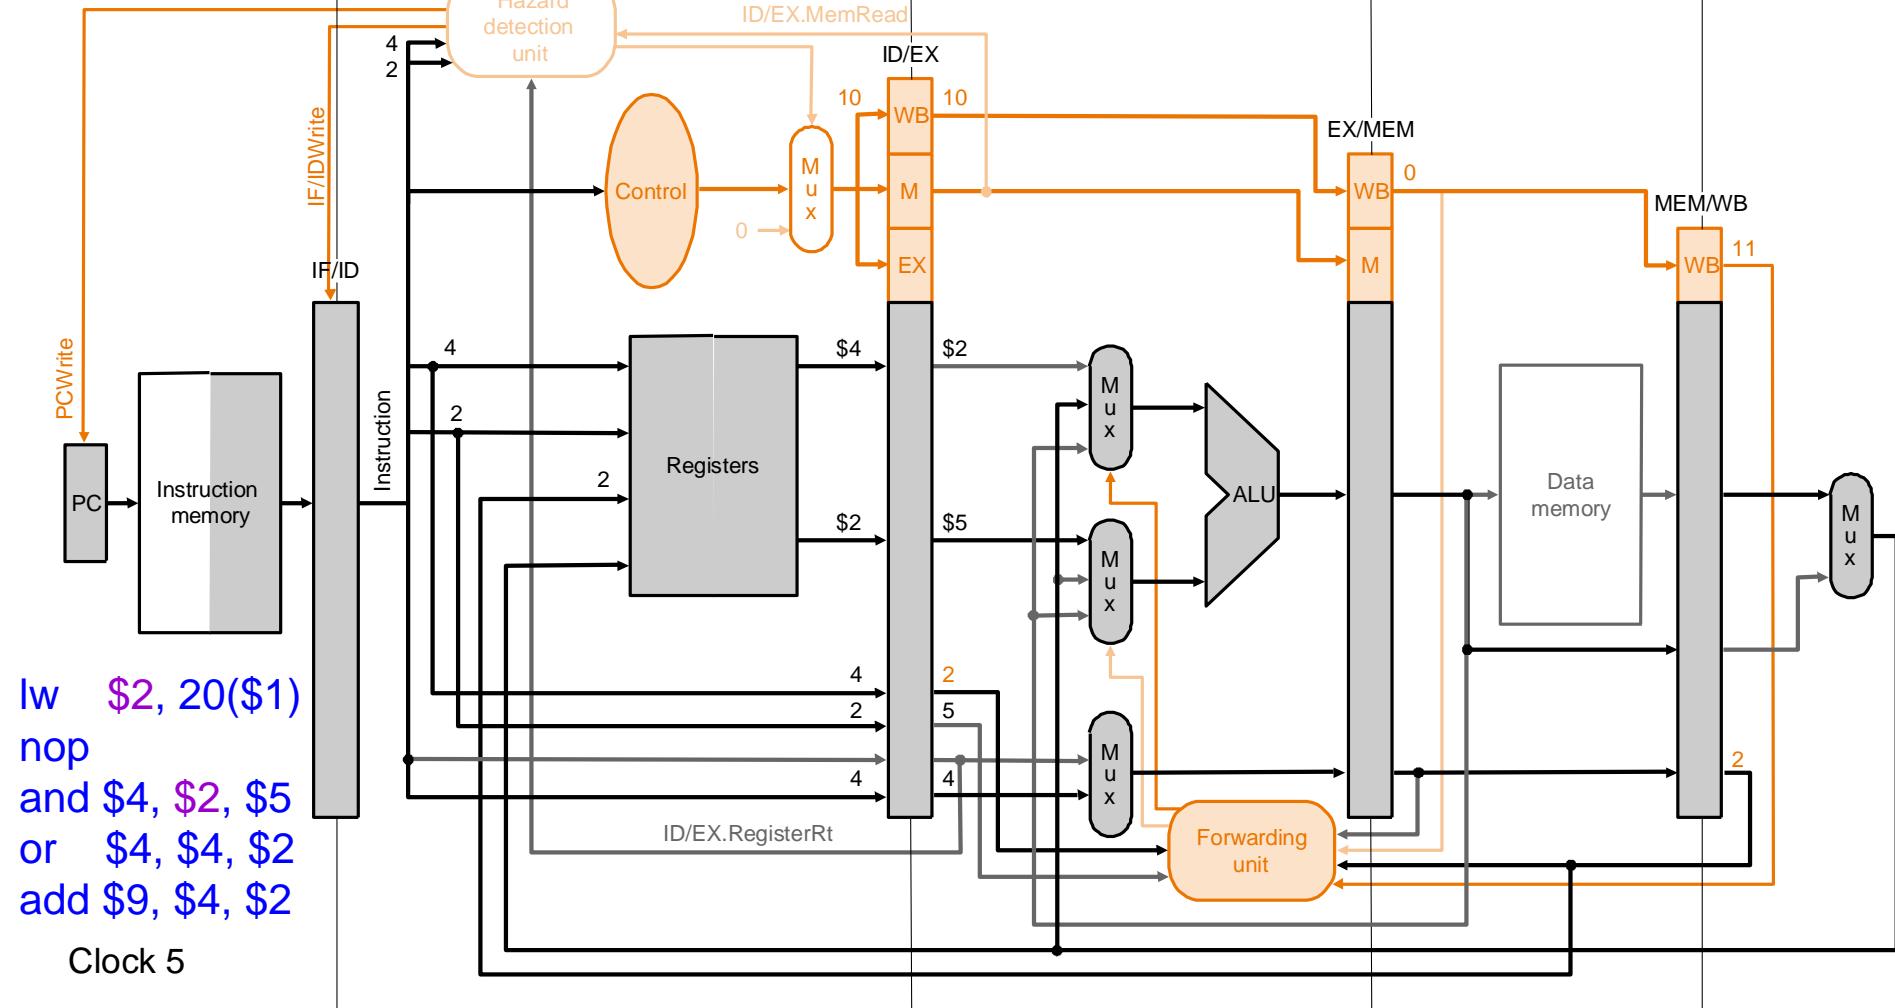
Walk-through (4 of 6)

add \$9, \$4, \$2

or \$4, \$4, \$2

and \$4, \$2, \$5

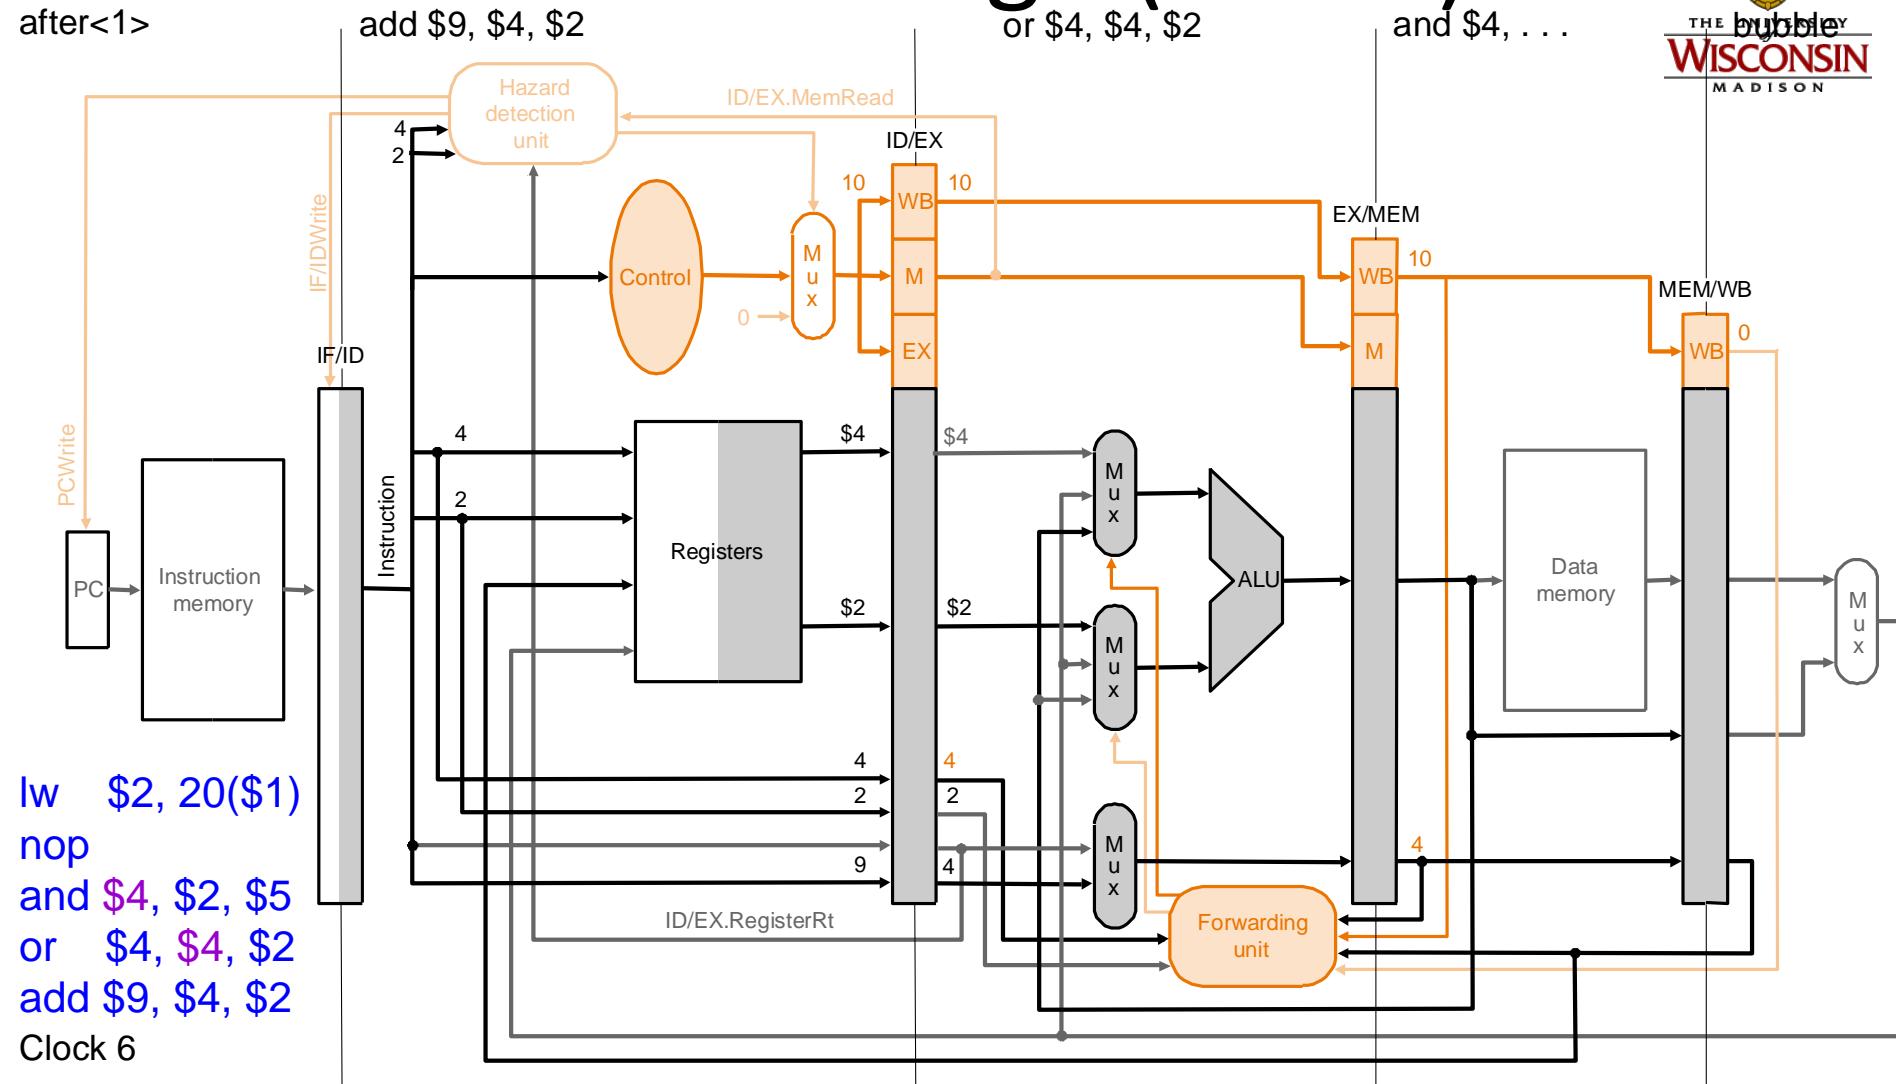
bubble



- Load value forwarded from MEM/WB register

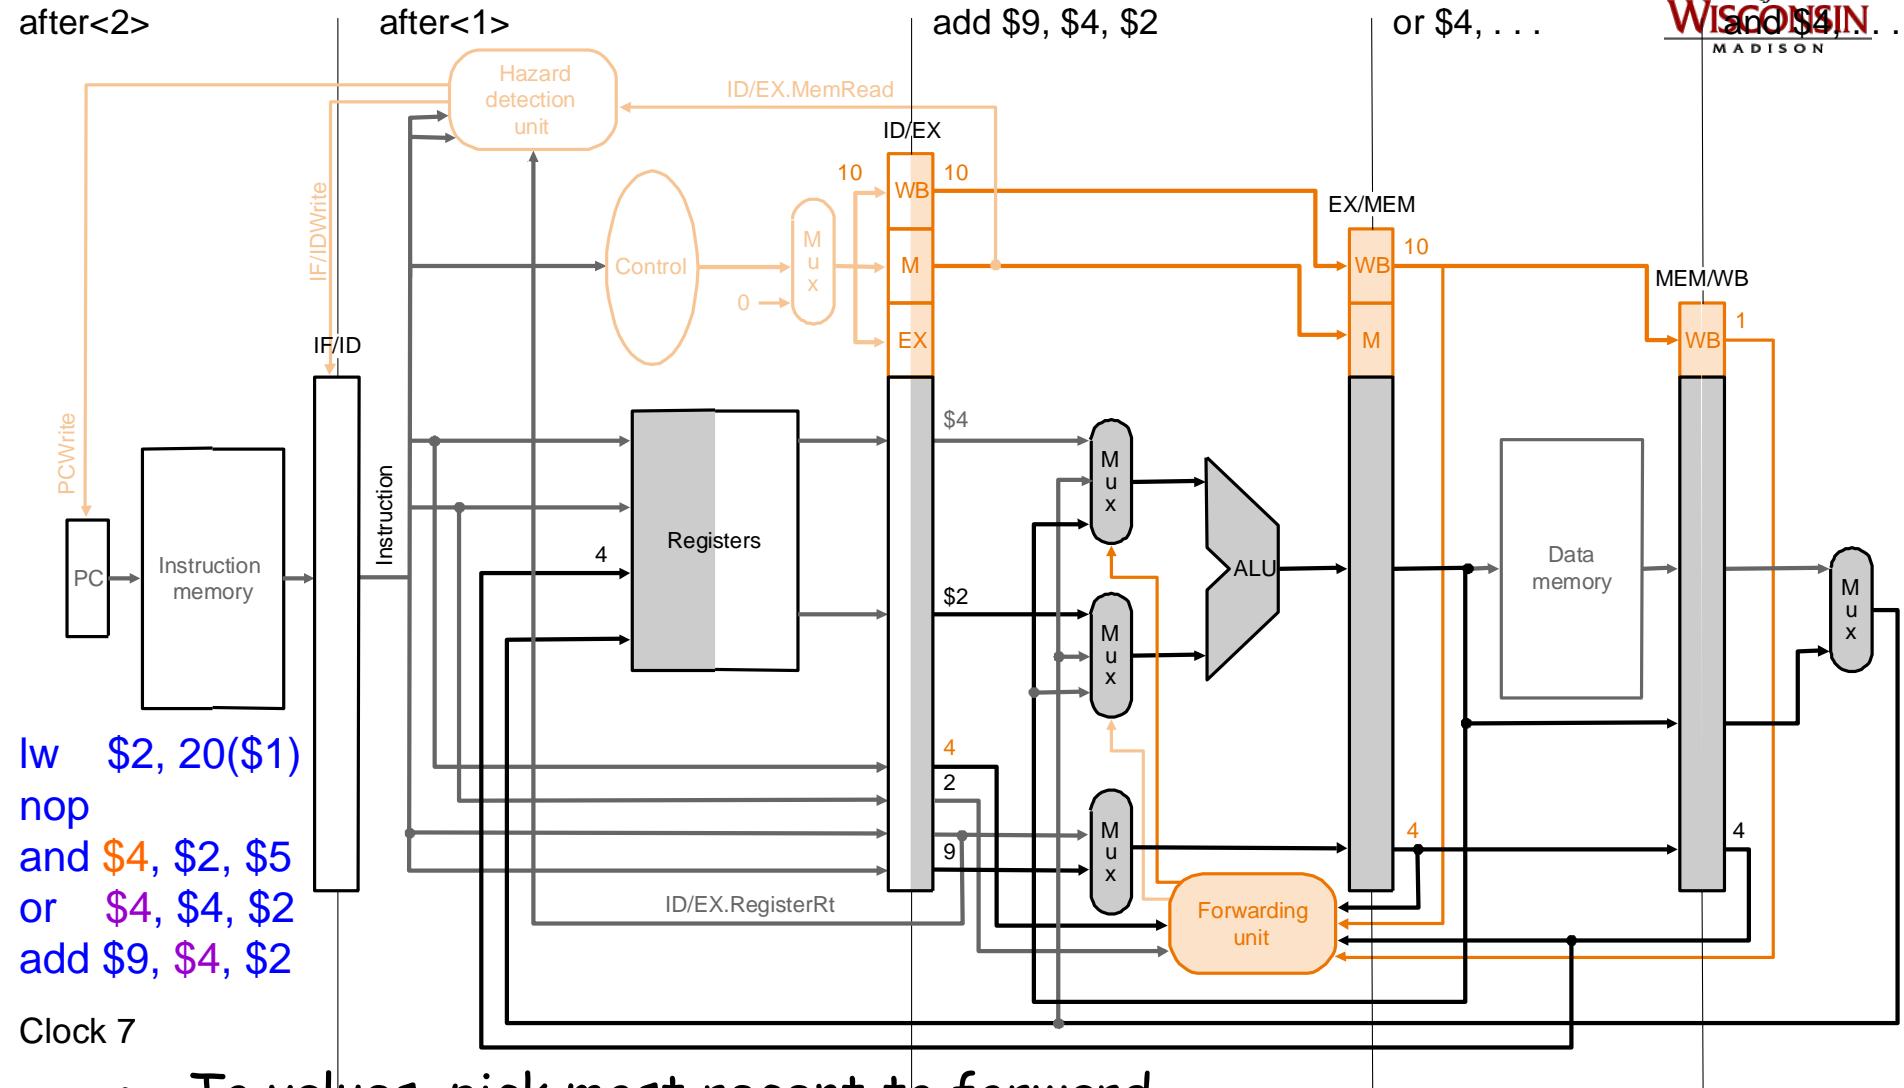


Walk-through (5 of 6)



- \$4 value forwarded from EX/MEM register

Walk-through (6 of 6)

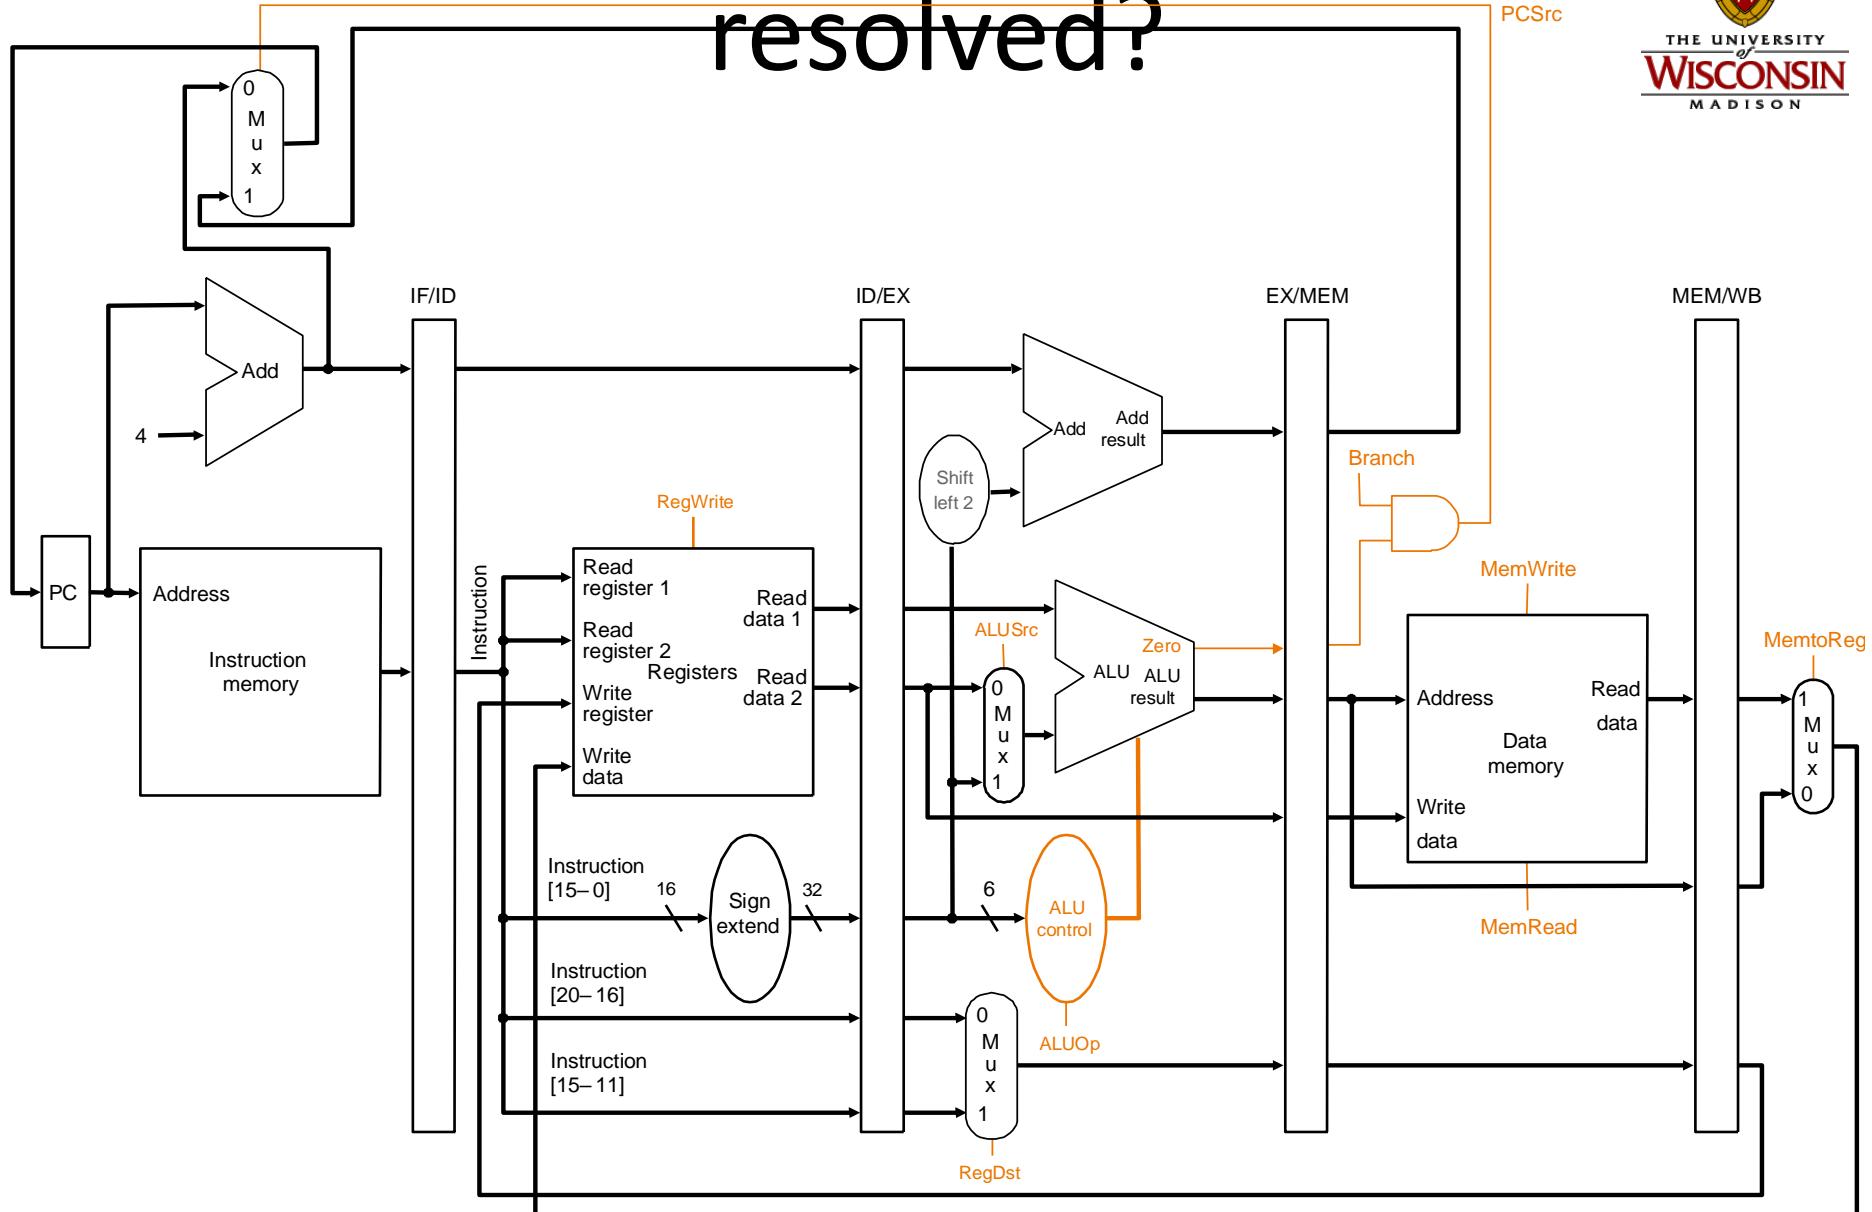


- To values, pick most recent to forward

RAW Hazard with Loads: Summary

- True backward dependencies in time
 - Need to stall
- Stall achieved by
 - Detecting hazard (remember logic equation)
 - Inserting NOP (all EX/MEM/WB controls set to 0)
 - Preventing IF/ID register and PC from being overwritten
- Next Branch/Control Hazards

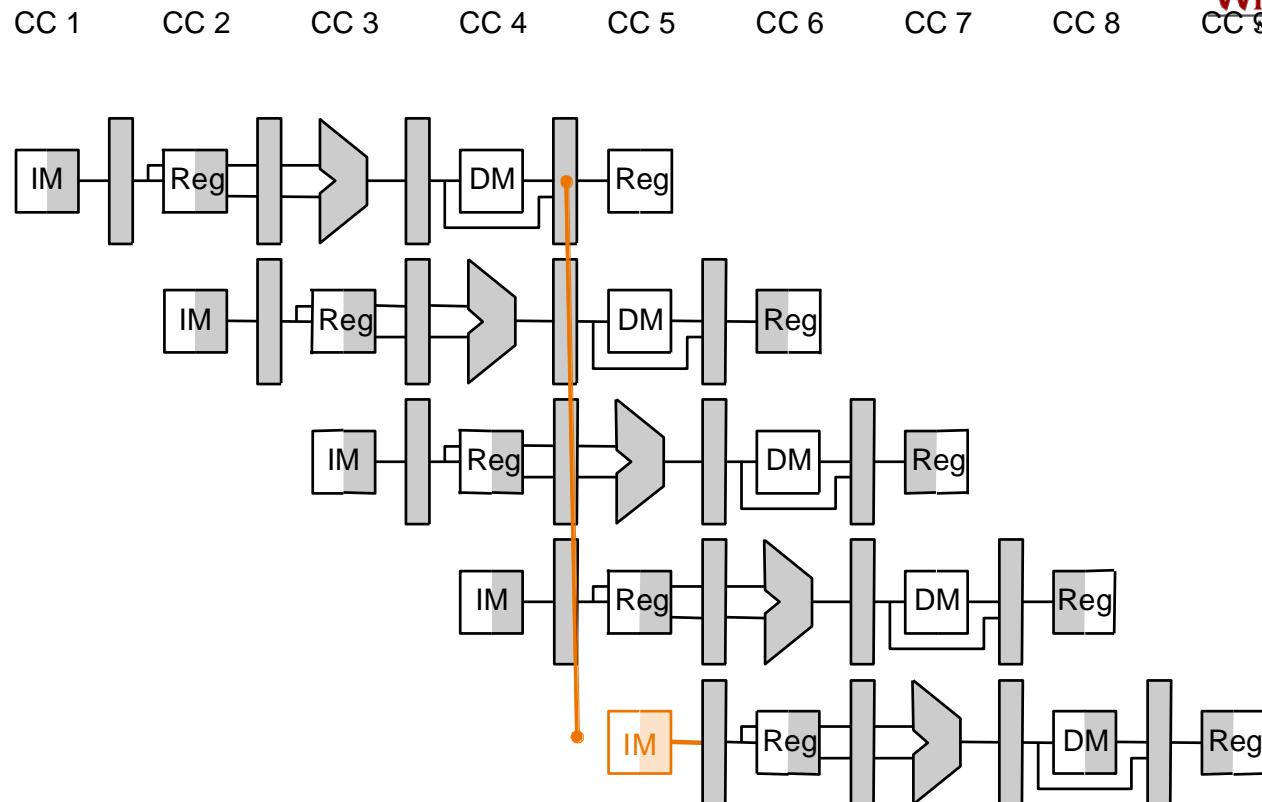
When conditional branches resolved?





Program
execution
order
(in instructions)

Time (in clock cycles)

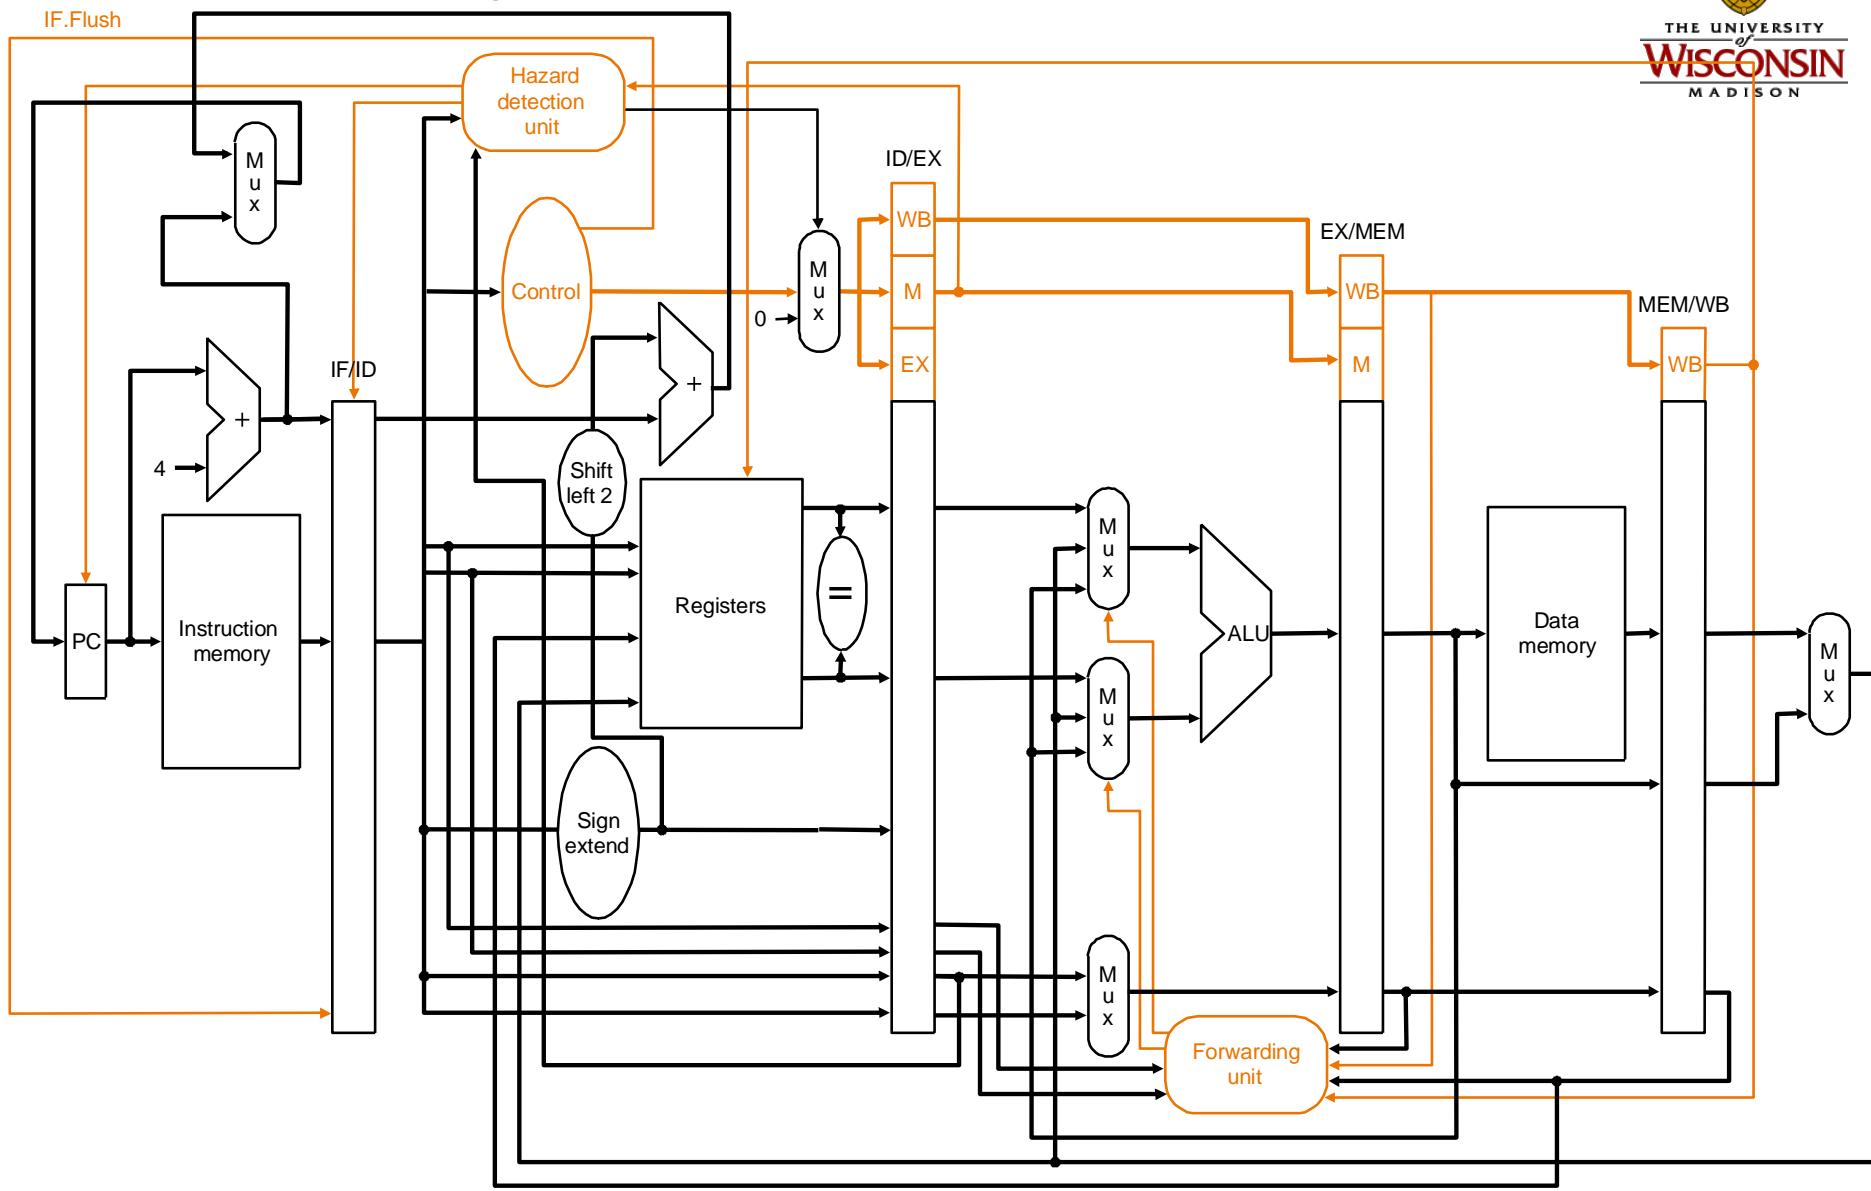


- Branch resolved in the MEM stage
- If taken,
 - $PC \leftarrow PC + 4 + SX(Imm * 4)$
 - $40 + 4 + 7 * 4 = 72$

Control/Branch Hazards

- Branch resolved in the MEM stage
 - But next instruction has to fetched in the next cycle
 - Reduce the penalty by moving decision earlier in pipeline
 - Need additional **comparator** ($r1=r2?$) and **adder** ($PC+4+SX(IMM)*4$)
 - Reduced penalty from 3 cycles to 1 cycle

Datapath for branch hazards



Eliminate 1-cycle stall?

- Two solutions
 - Predict branch is always not taken
 - More sophisticated prediction schemes
 - Delay slots
 - Compiler's problem
- Walkthrough example for solution #1
 - Predict not taken

Walkthrough (1 of 2)



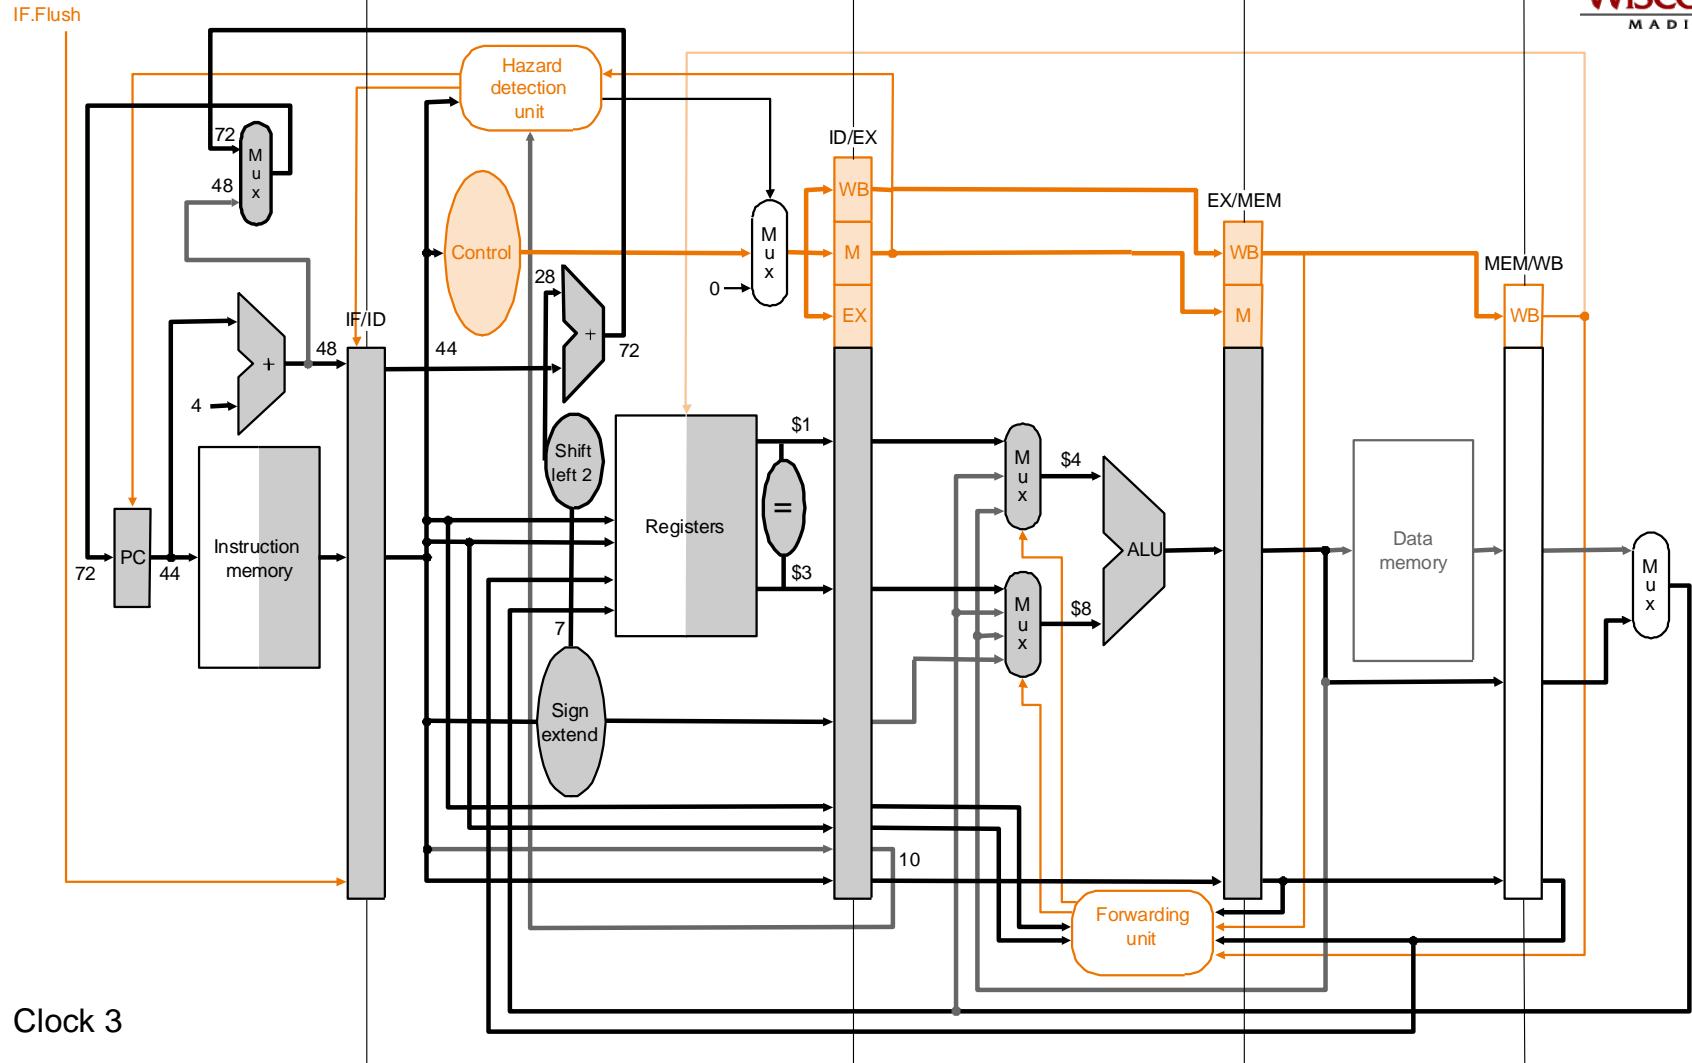
before<
THE UNIVERSITY
of
WISCONSIN
MADISON

and \$12, \$2, \$5

beq \$1, \$3, 7

sub \$10, \$4, \$8

before<1>



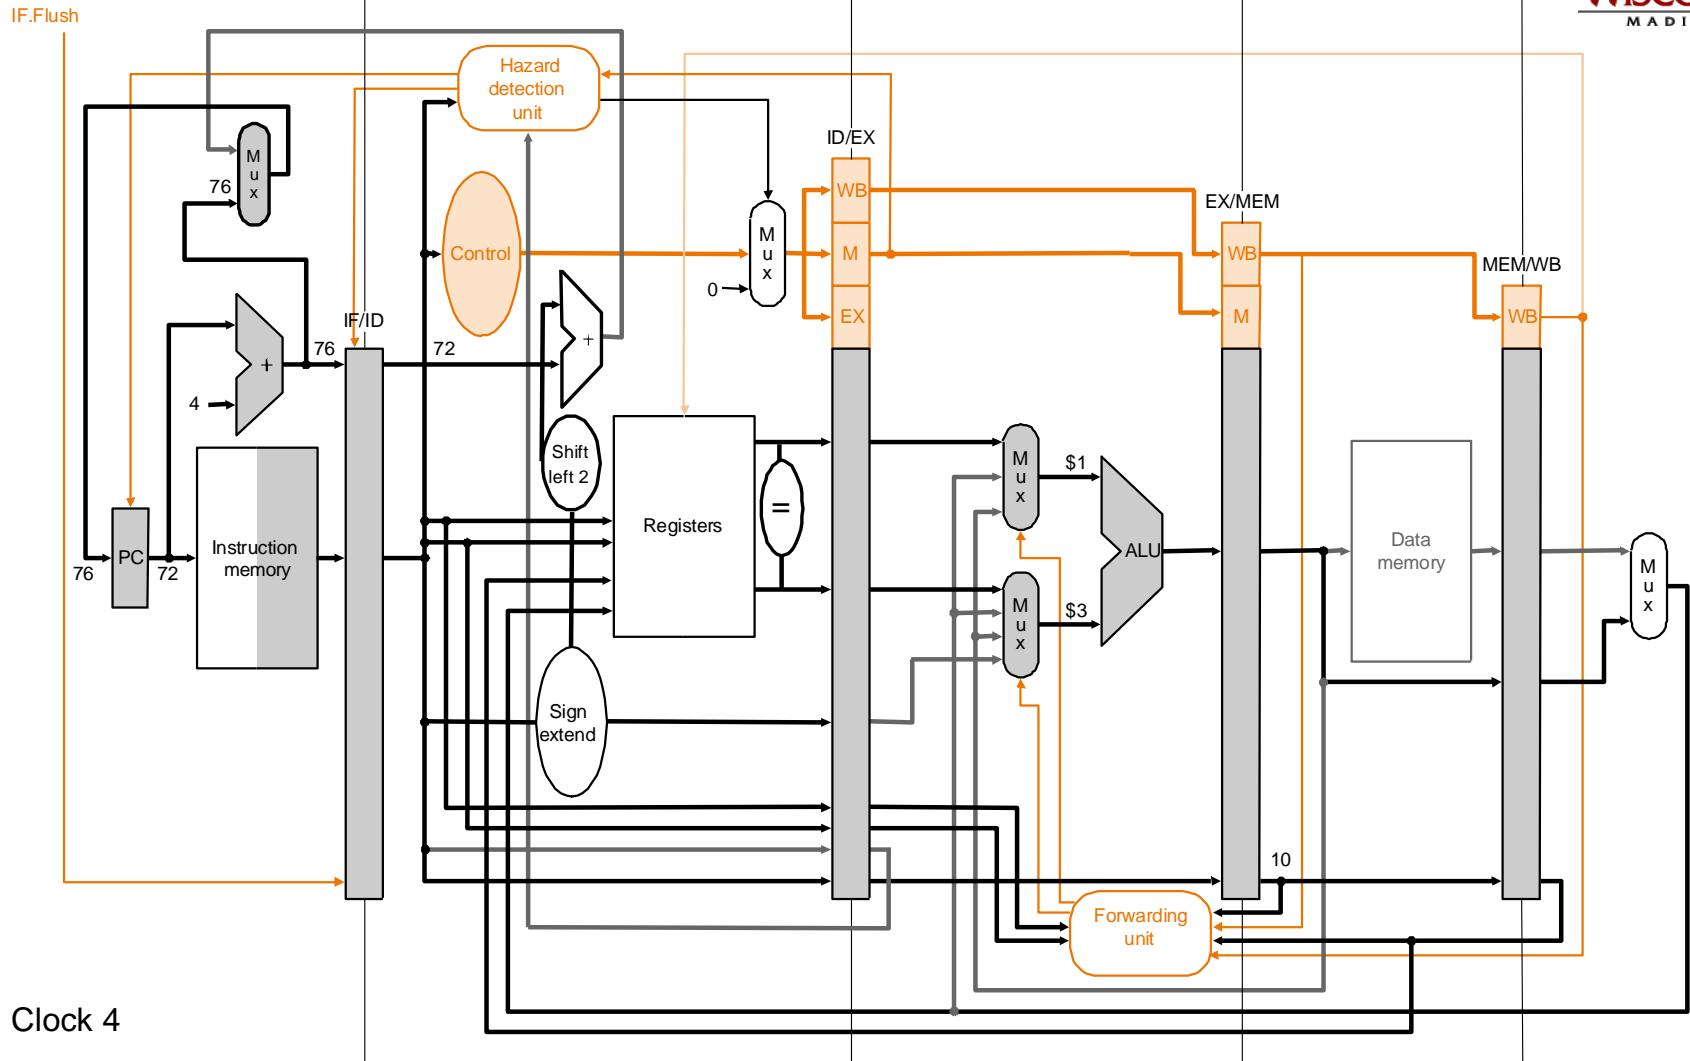
Walkthrough (2 of 2)

lw \$4, 50(\$7)

bubble (nop)

beq \$1, \$3, 7

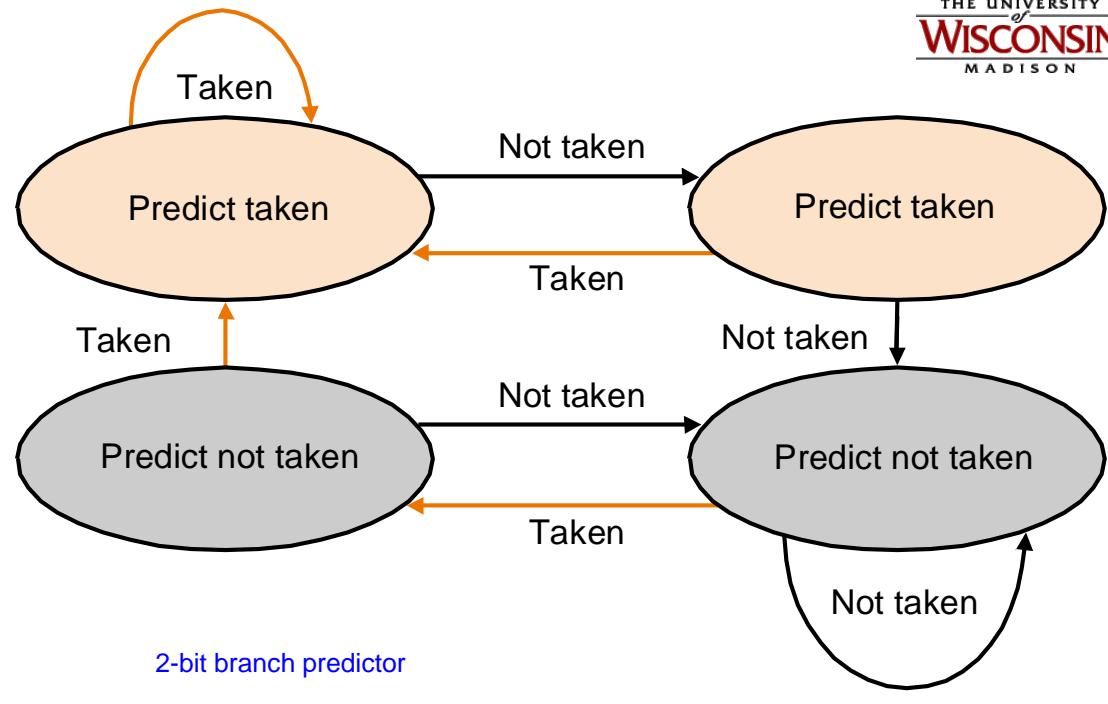
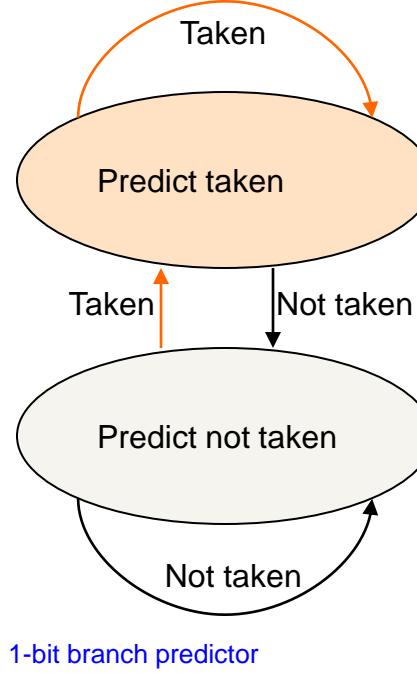
sub \$10, ...



Dynamic Branch Prediction

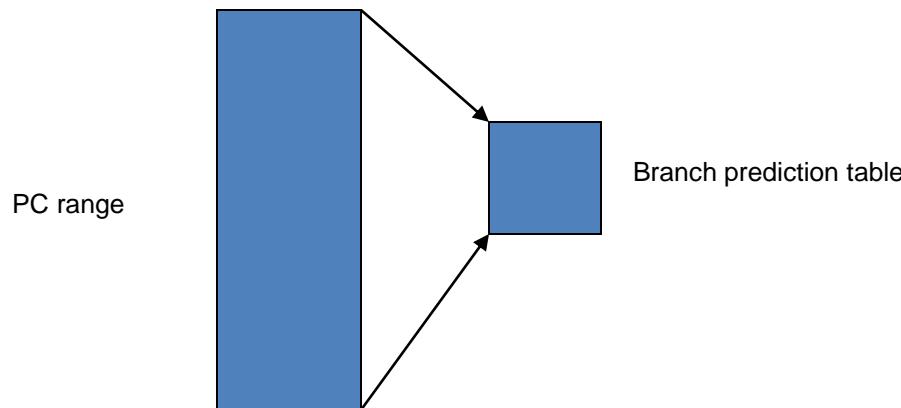
- Better than static prediction
 - Branches are predictable
 - ~90% of program execution time is spent in ~10% of code (inner loops)
 - Think of a program loop of N iterations
 - Taken $N-1$ times
 - Not taken last time

Dynamic Branch Prediction



- How does hardware “learn” branch behavior?
- Store each branch instruction’s history ***
 - If a branch was taken “recently”, predict taken
 - One bit saturating counter
 - Two bit counters

Branch Prediction



- Store each branch's history ***
 - Not really
- Keep a small table indexed by program counter
- PC is large (32 bit number)
- Mapping to number of table entries
 - E.g. 16-entry branch prediction table
 - Mapping: use last 4 bits of PC
- Problem: Multiple branches may map to same entry in table -- **Aliasing**

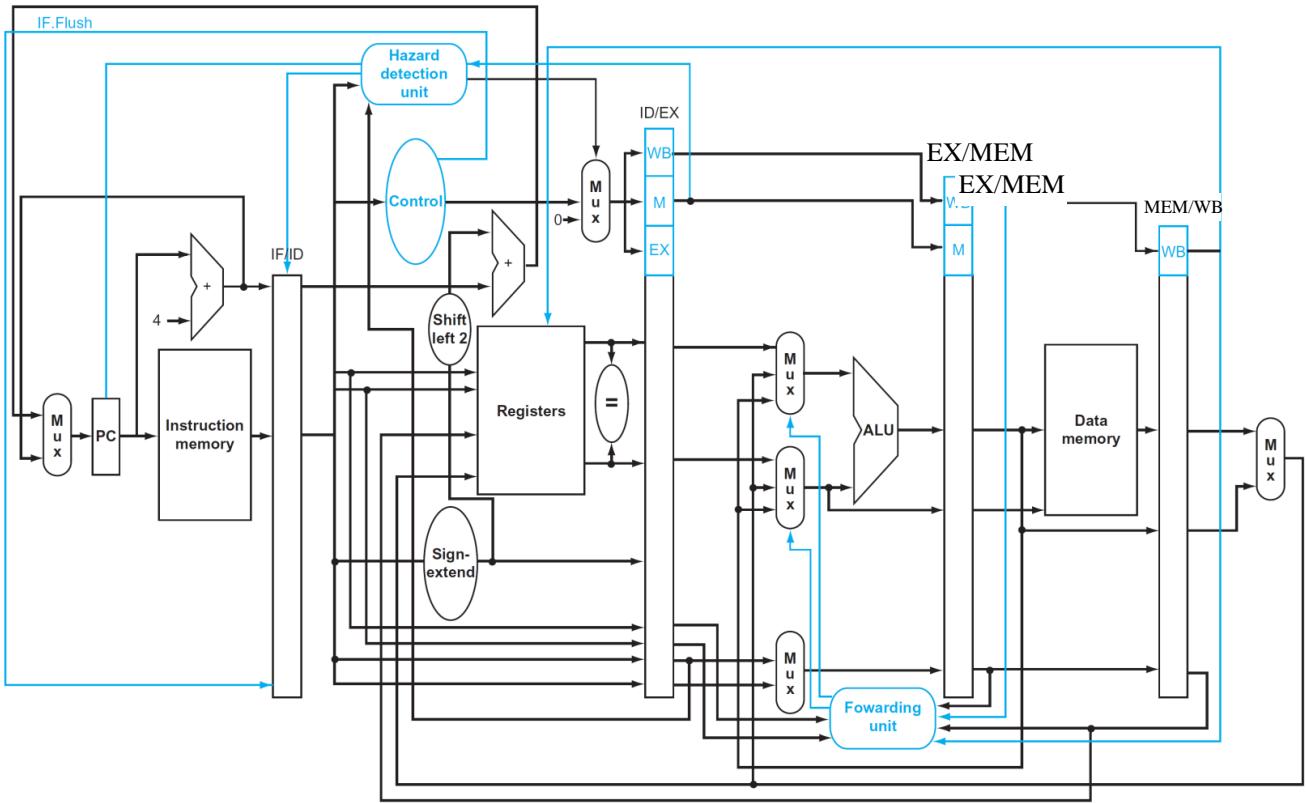
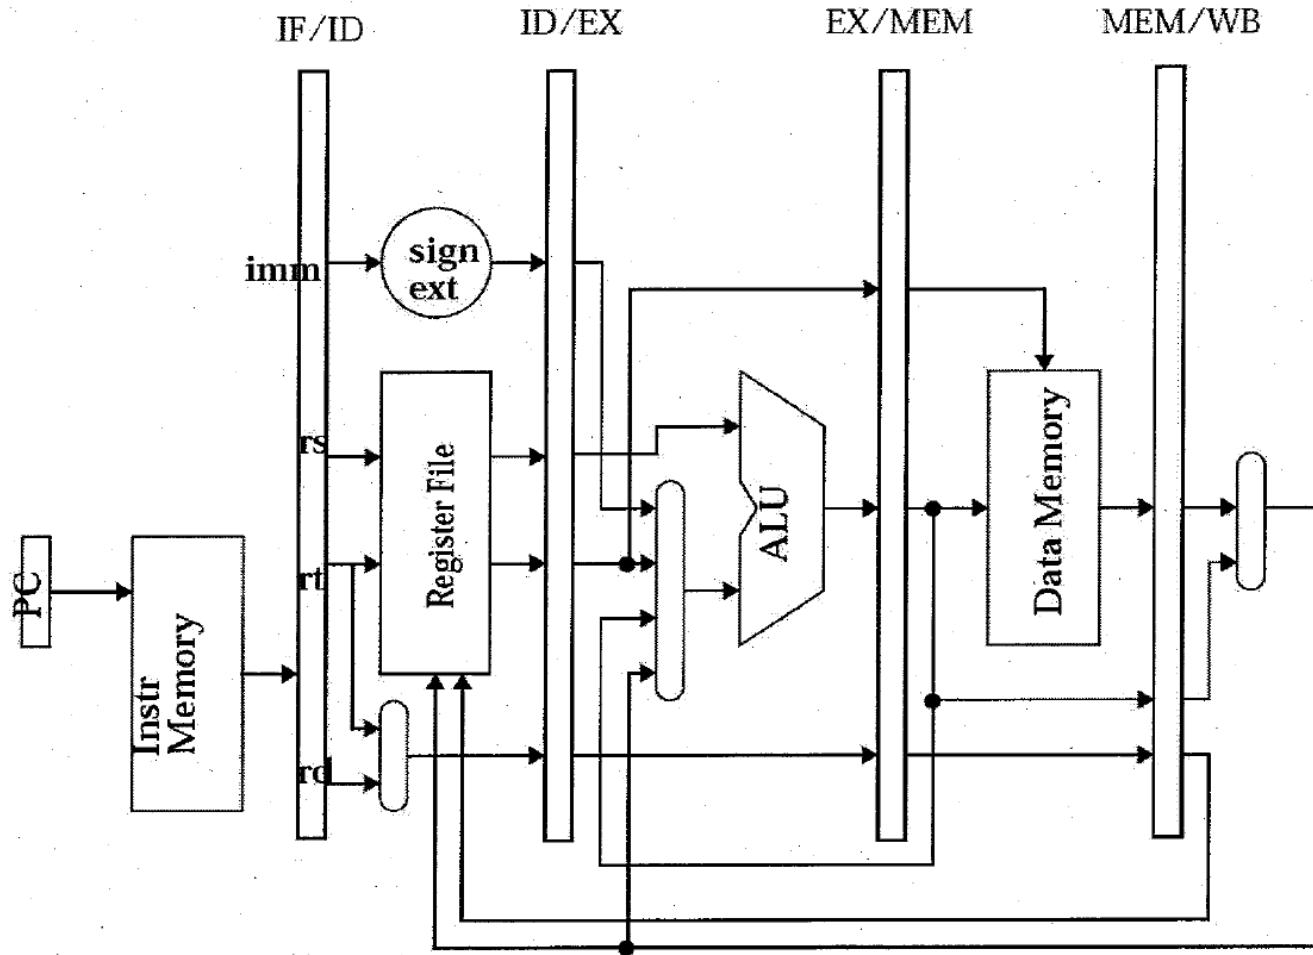


FIGURE 4.65 The final datapath and control for this chapter. Note that this is a stylized figure rather than a detailed datapath, so it's missing the ALUsrc Mux from Figure 4.57 and the multiplexer controls from Figure 4.51

Instruction	1	2	3	4	5	6	7	8	9	10	11	12	13
i0: add \$1, \$2, \$3	F	D	X	M	W								
i1: sub \$4, \$1, \$5		F	D	X x-x fwd	M	W							
i2: or \$6, \$1, \$4			F	D	X x-x fwd(\$4) m-x fwd(\$1)	M	W						
i3: and \$7, \$4, \$8				F	D	X w-x fwd	M	W					
i4: lw \$9, 4(\$7)					F	D	X x-x fwd	M	W				
i5: add \$1, \$9, \$2						F	D	D	X m-x fwd	M	W		
i6: sw \$1, 4(\$7)							F	F	D	X	M	W	



Instruction	1	2	3	4	5	6	7	8	9	10	11	12	13
i0: add \$1, \$2, \$3	F	D	X	M	W								
i1: sub \$4, \$1, \$5		F	D	D	D	X	M	W					
i2: or \$6, \$1, \$4			F	F	F	D	X	M	W				
i3: and \$7, \$4, \$8					F	D	D	X	M	W			
i4: lw \$9, 4(\$7)						F	F	D	D	D	X	M	W
i5: add \$1, \$9, \$2							F	F	F	D	D	D	
i6: sw \$1, 4(\$7)													

Exceptions and Pipelining

- add \$1, \$2, \$3 overflows
- A surprise branch
 - Earlier instructions flow to completion
 - Kill later instructions
 - Save PC in EPC, set PC to EX handler, etc.
- Costs a lot of designer sanity
 - 554 teams that try this sometimes fail

Exceptions

- Even worse: in one cycle
 - I/O interrupt
 - User trap to OS (EX)
 - Illegal instruction (ID)
 - Arithmetic overflow
 - Hardware error
 - Etc.
- Interrupt priorities must be supported

Pipeline Hazards

- Program Dependences
- Data Hazards
 - Stalls
 - Forwarding
- Control Hazards
- Exceptions