
CSE 320
Computer Architecture
Fall 2009

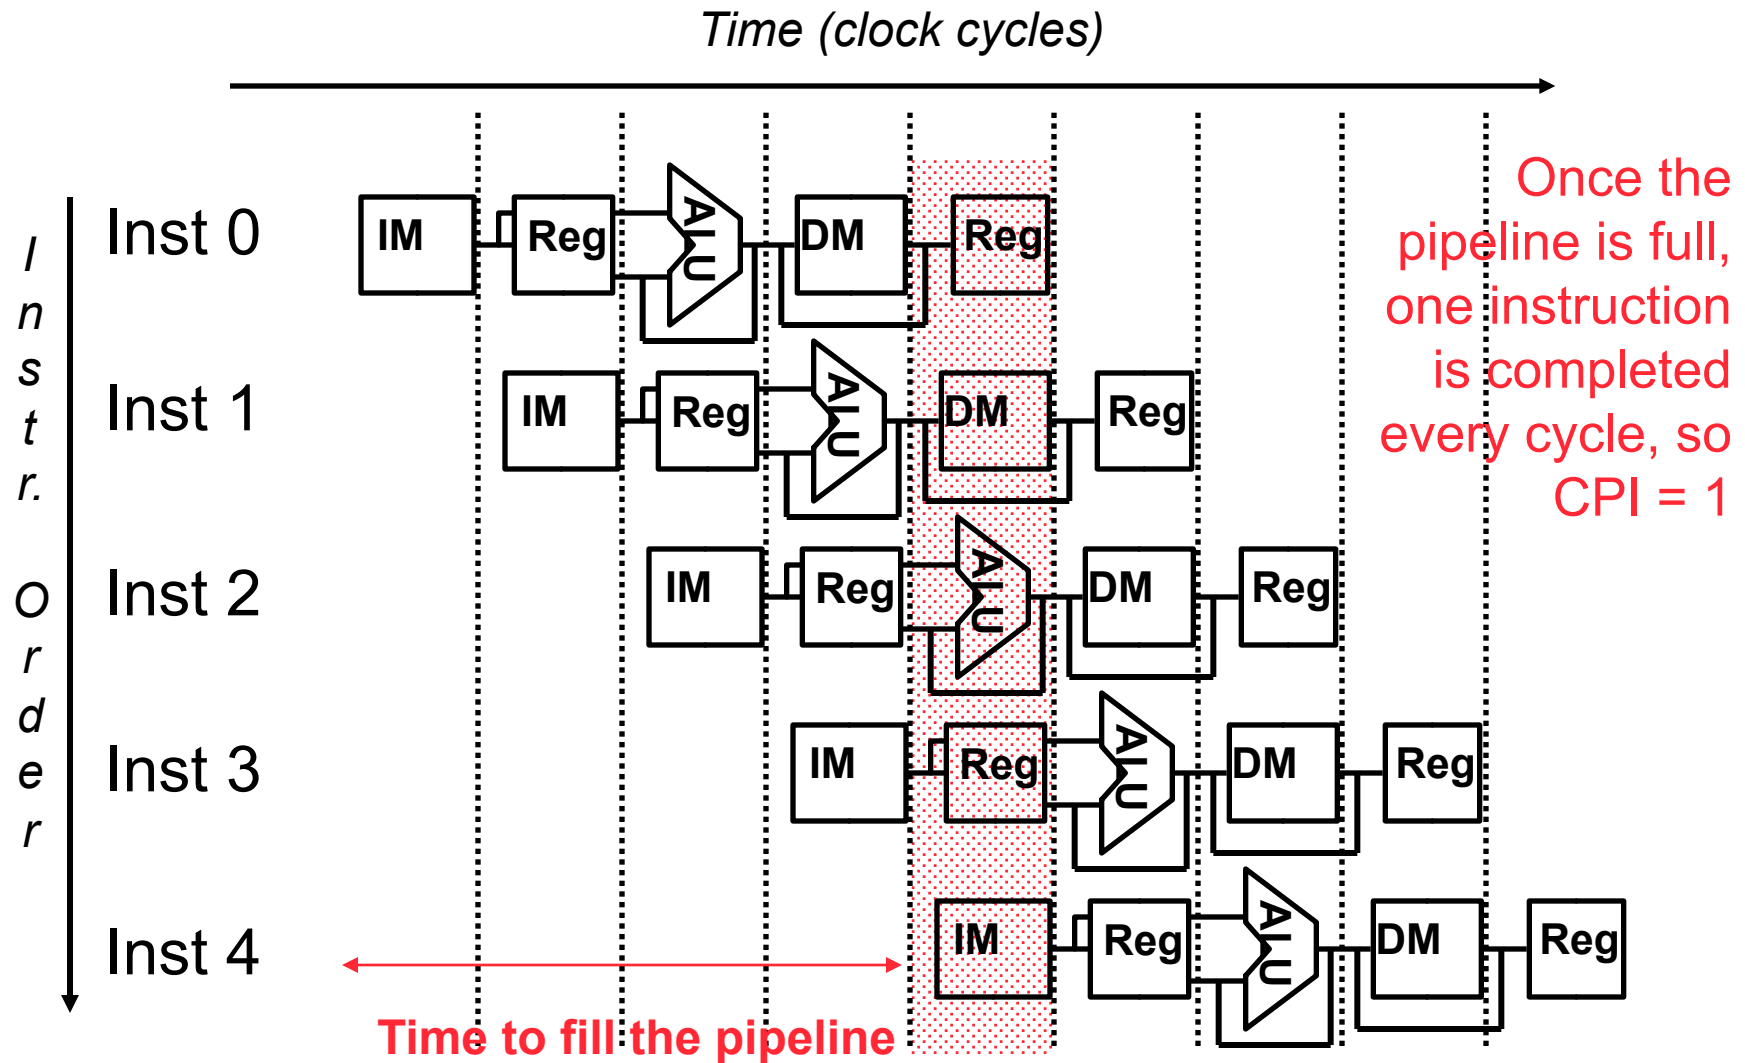
Chapter 4B: The Processor,
Part B

Larry Wittie
Stony Brook University

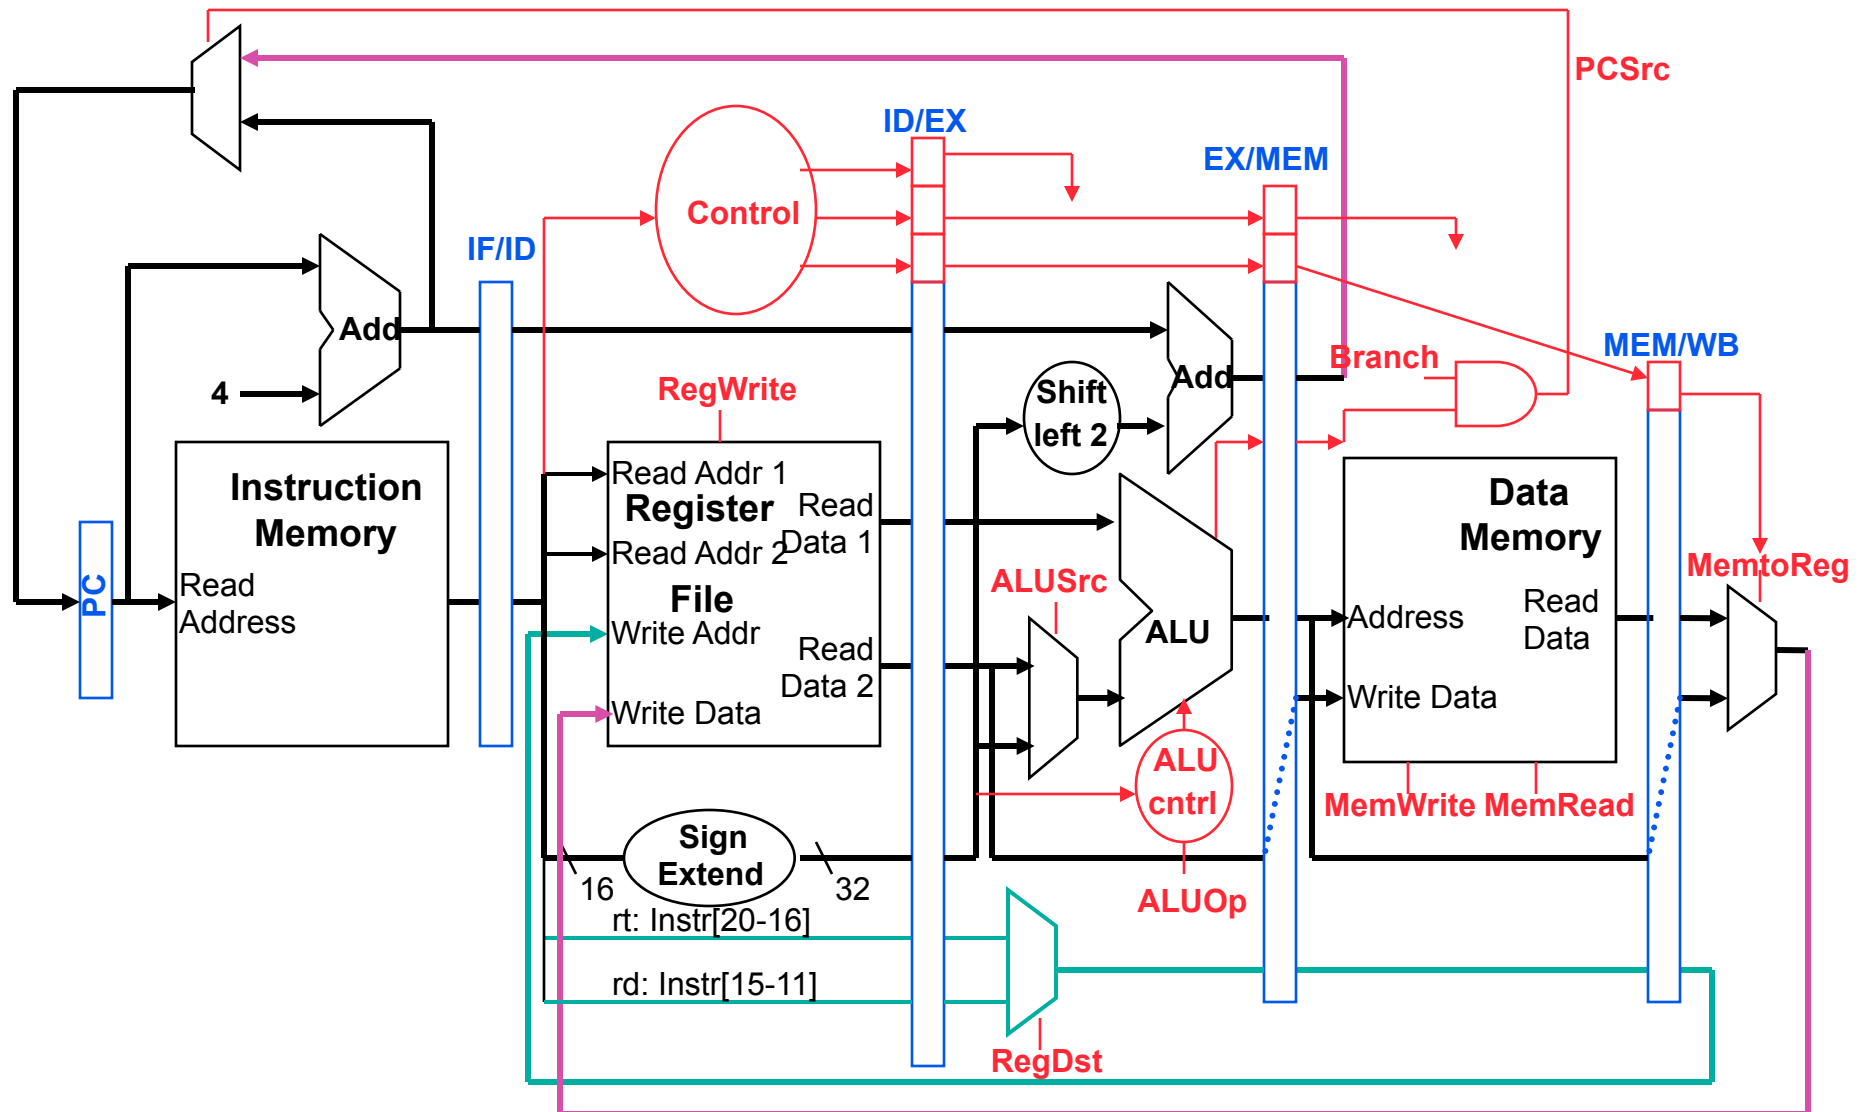
www.cs.sunysb.edu/~cse320

[Adapted from *Computer Organization and Design, 4th Edition*, Patterson & Hennessy, © 2008, MK, with many additions by Mary Jane Irwin, PennStateU]

Review: Why Pipeline? For Performance!



Review: MIPS Pipeline Data and Control Paths



Review: Can Pipelining Get Us Into Trouble?

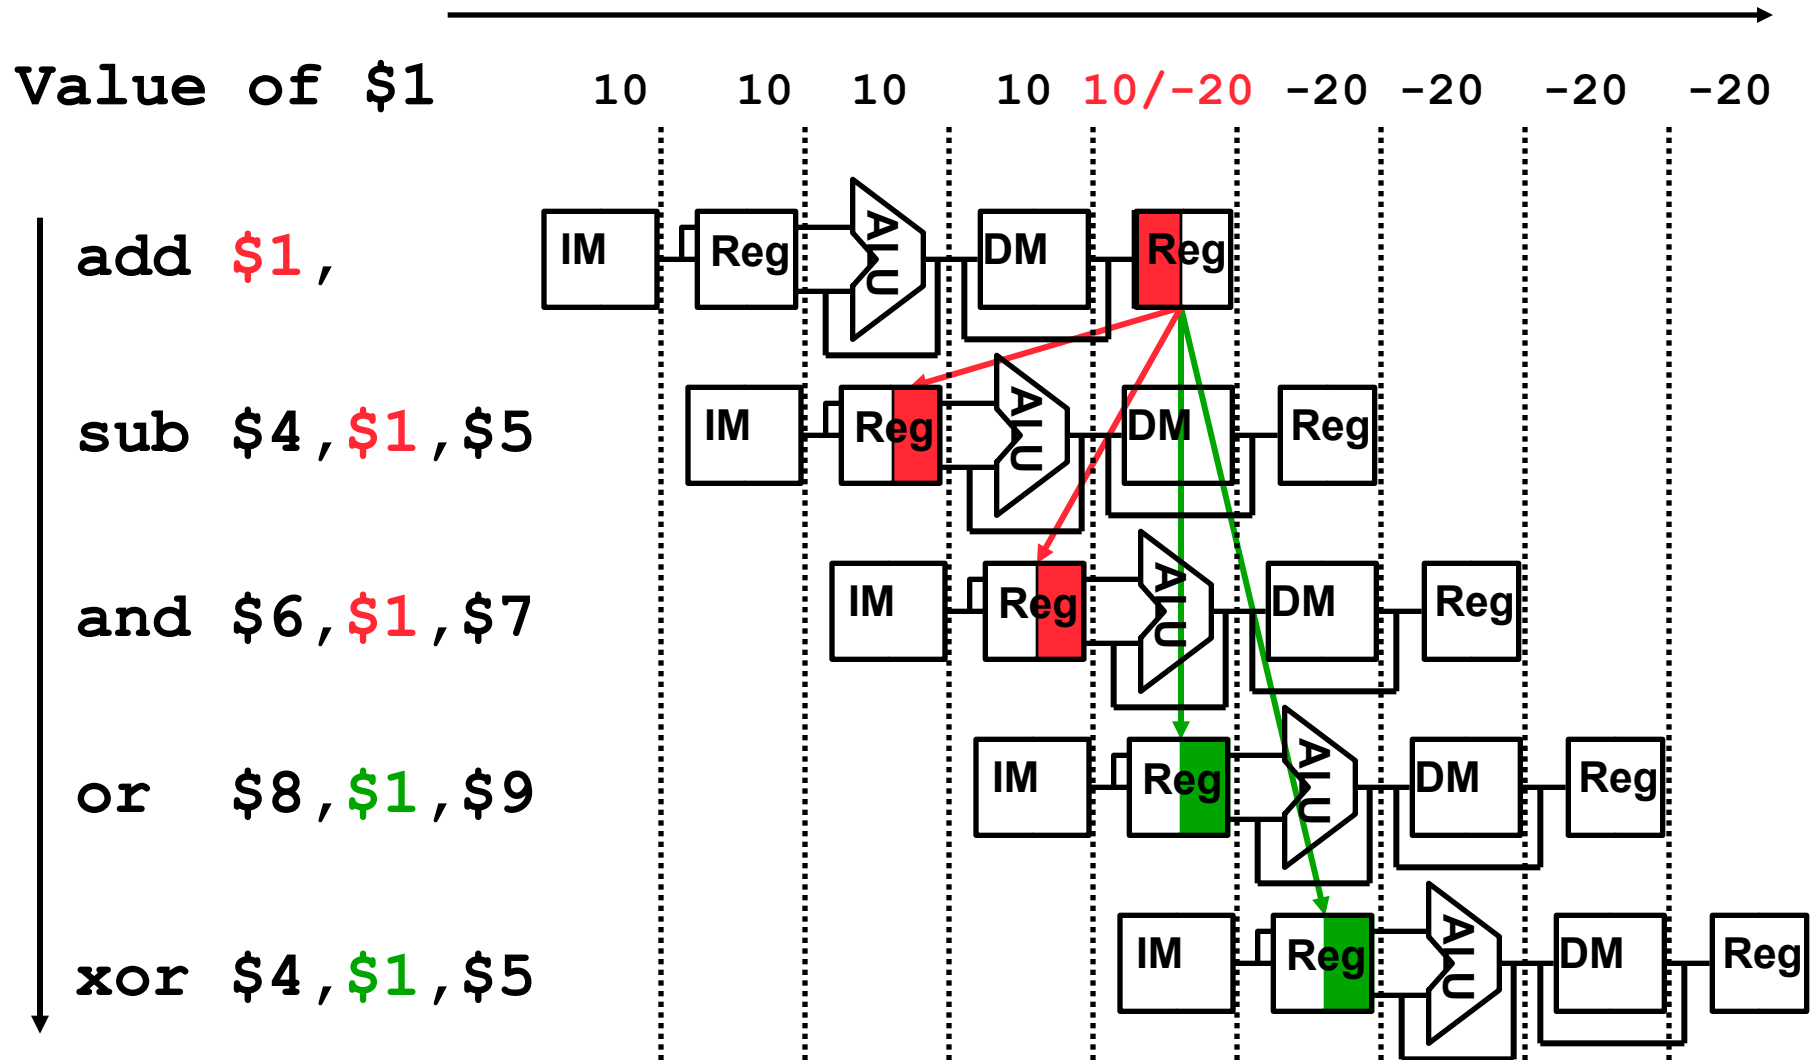
□ Yes: Pipeline Hazards

- **structural hazards**: attempt to use the same resource by two different instructions at the same time
- **data hazards**: attempt to use data before it is ready
 - An instruction's source operand(s) are produced by a prior instruction still in the pipeline
- **control hazards**: attempt to make a decision about program control flow before the condition has been evaluated and the new PC target address calculated
 - branch and jump instructions, exceptions

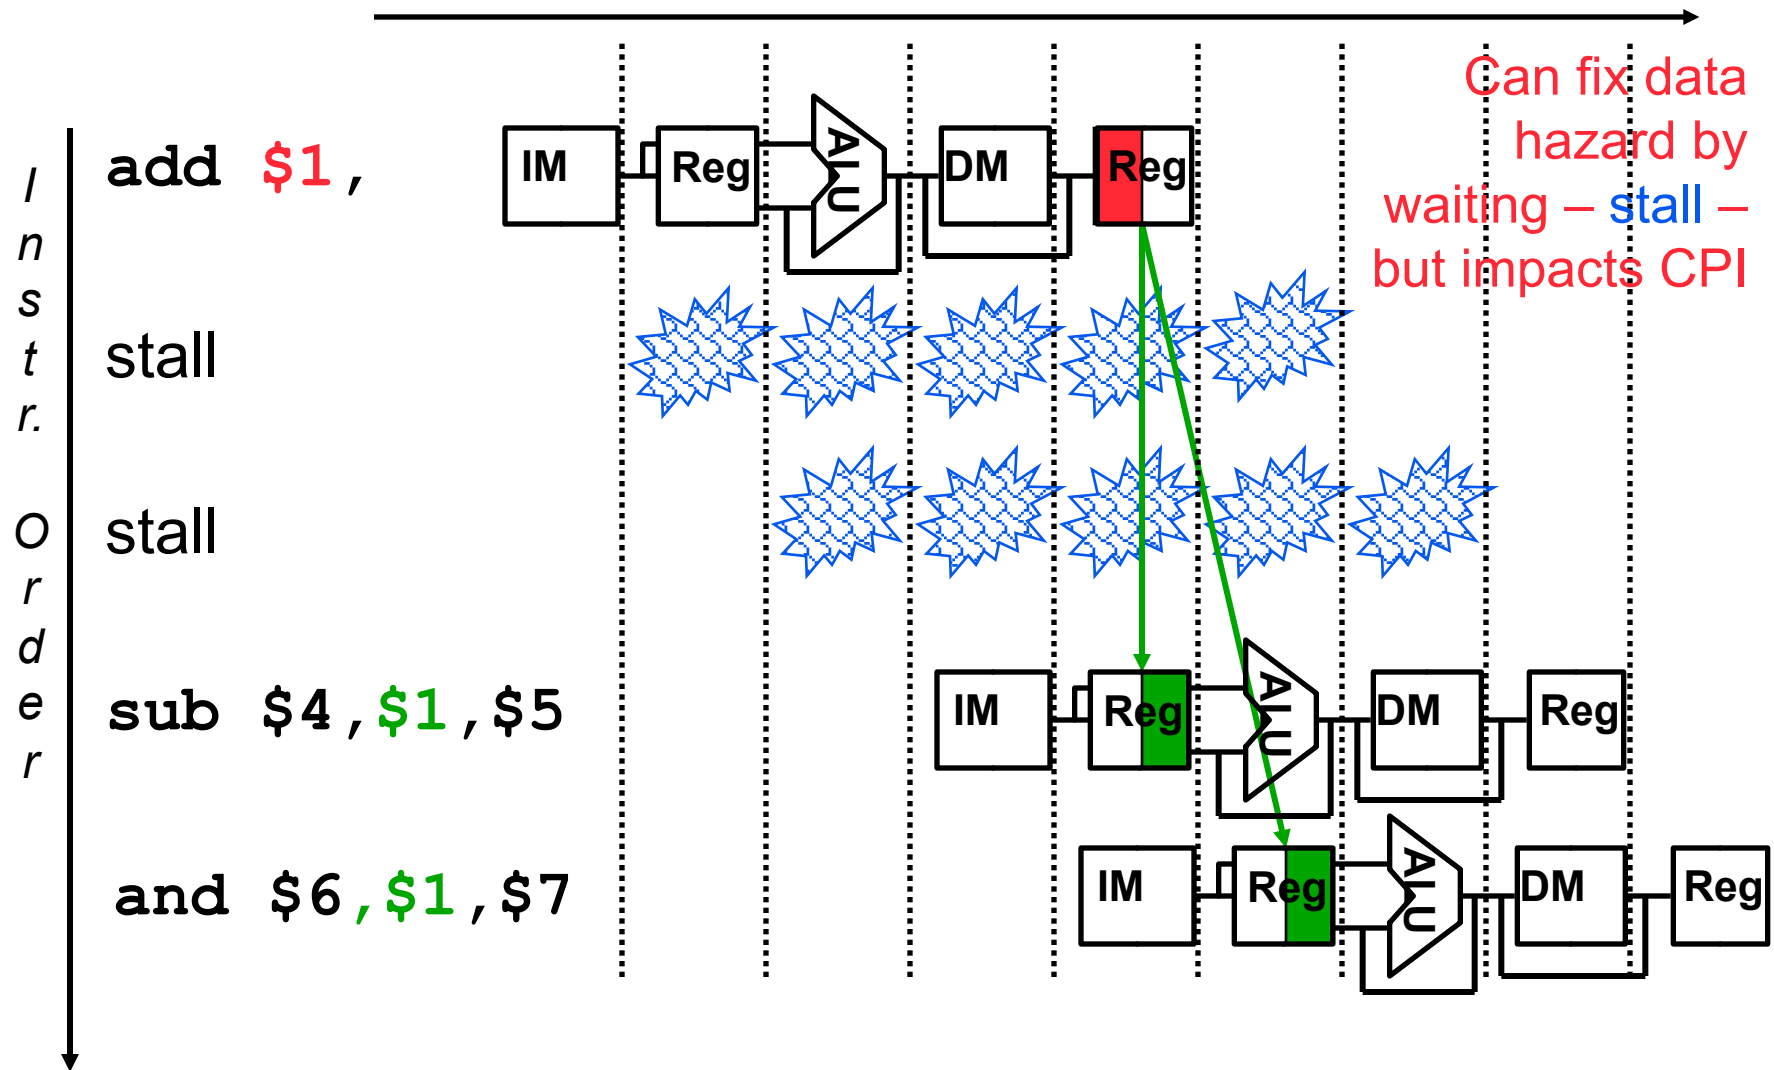
□ Pipeline control must **detect** the hazard and then take action to **resolve** hazards

Review: Register Usage Can Cause Data Hazards

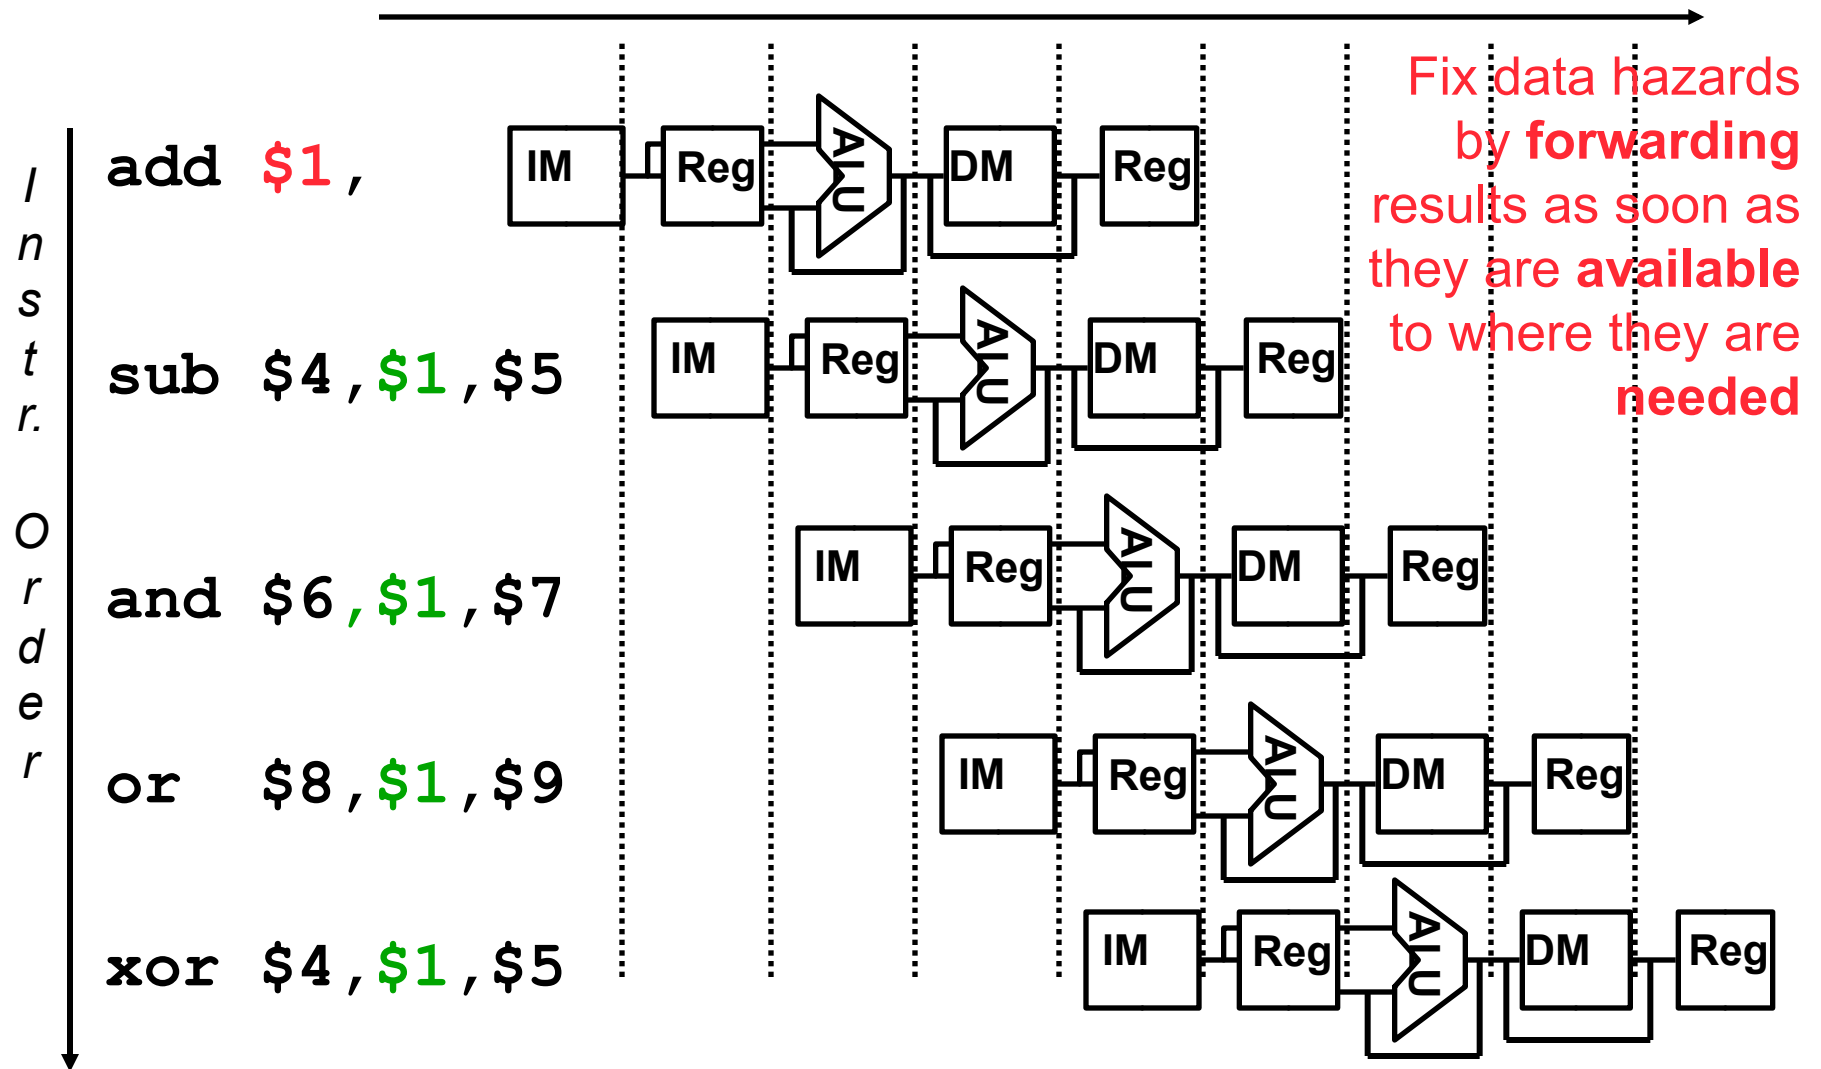
- Read after write (RAW) **data hazard**; error if **read-before-write**



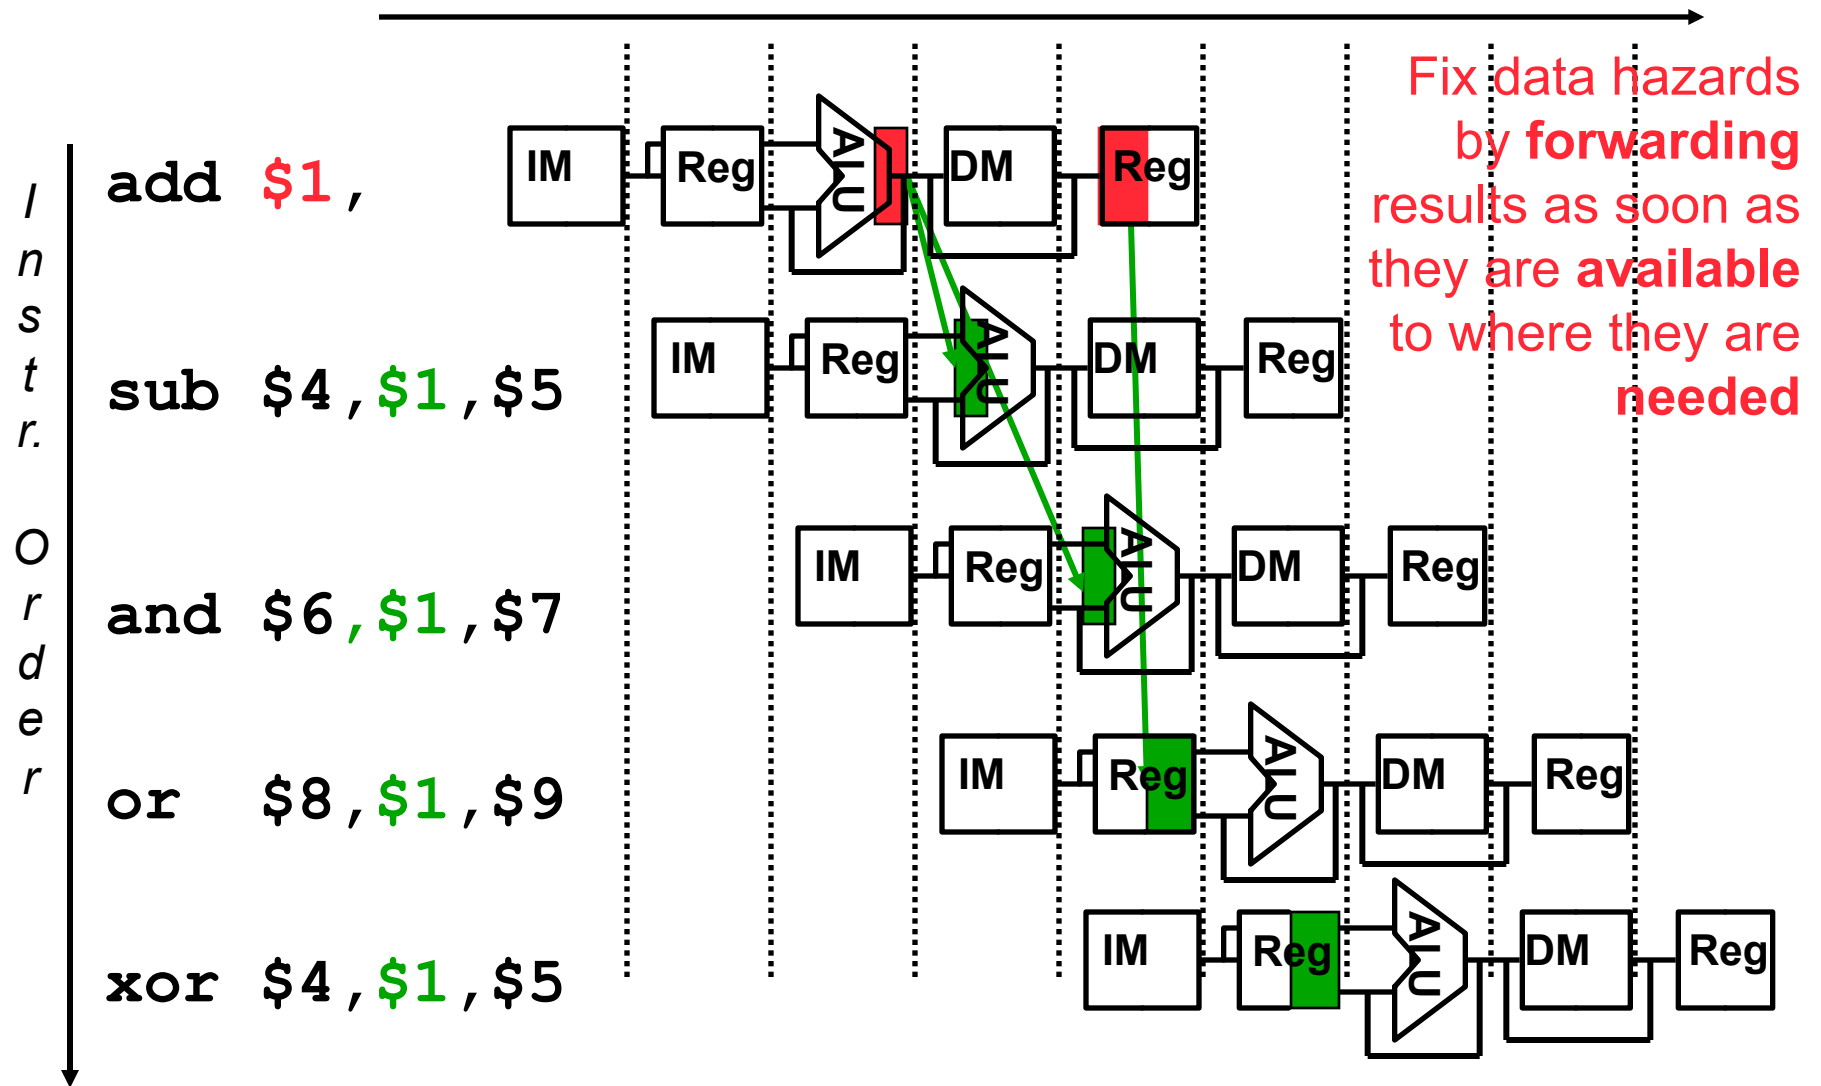
One Way to “Fix” a Data Hazard



A Better Way to "Fix" a Data Hazard



A Better Way to "Fix" a Data Hazard



Data Forwarding (aka 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

Data Forwarding Control Conditions

1. EX Forward Unit:

```
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
```

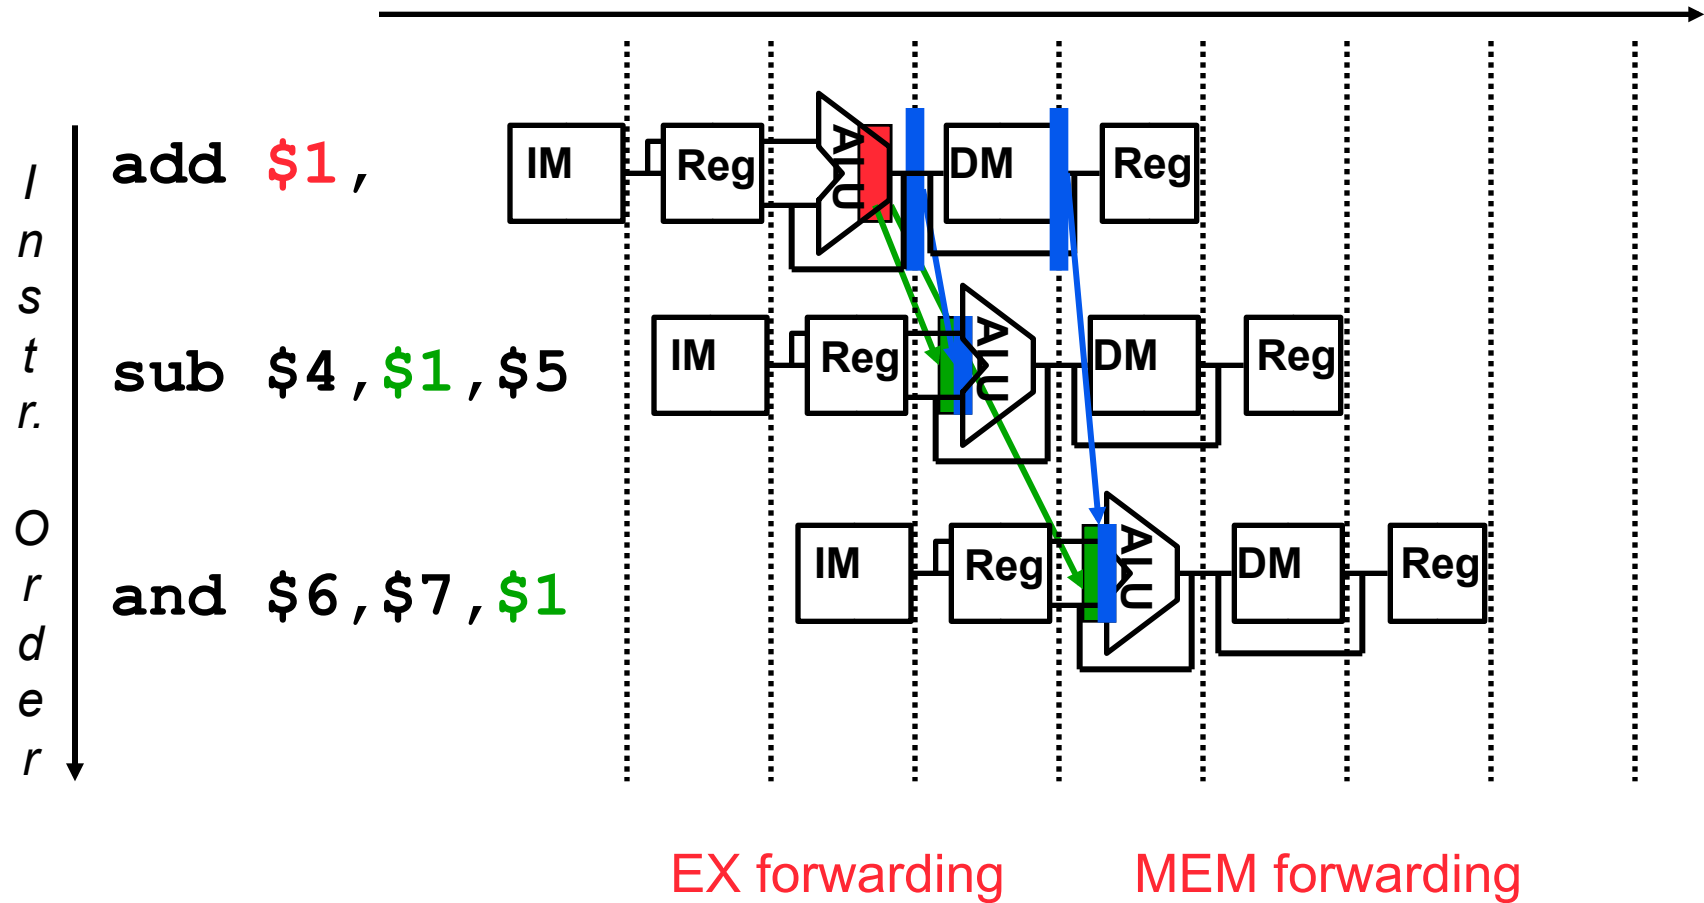
Forwards the result from the previous instr. to either input of the ALU

2. MEM Forward Unit:

```
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
```

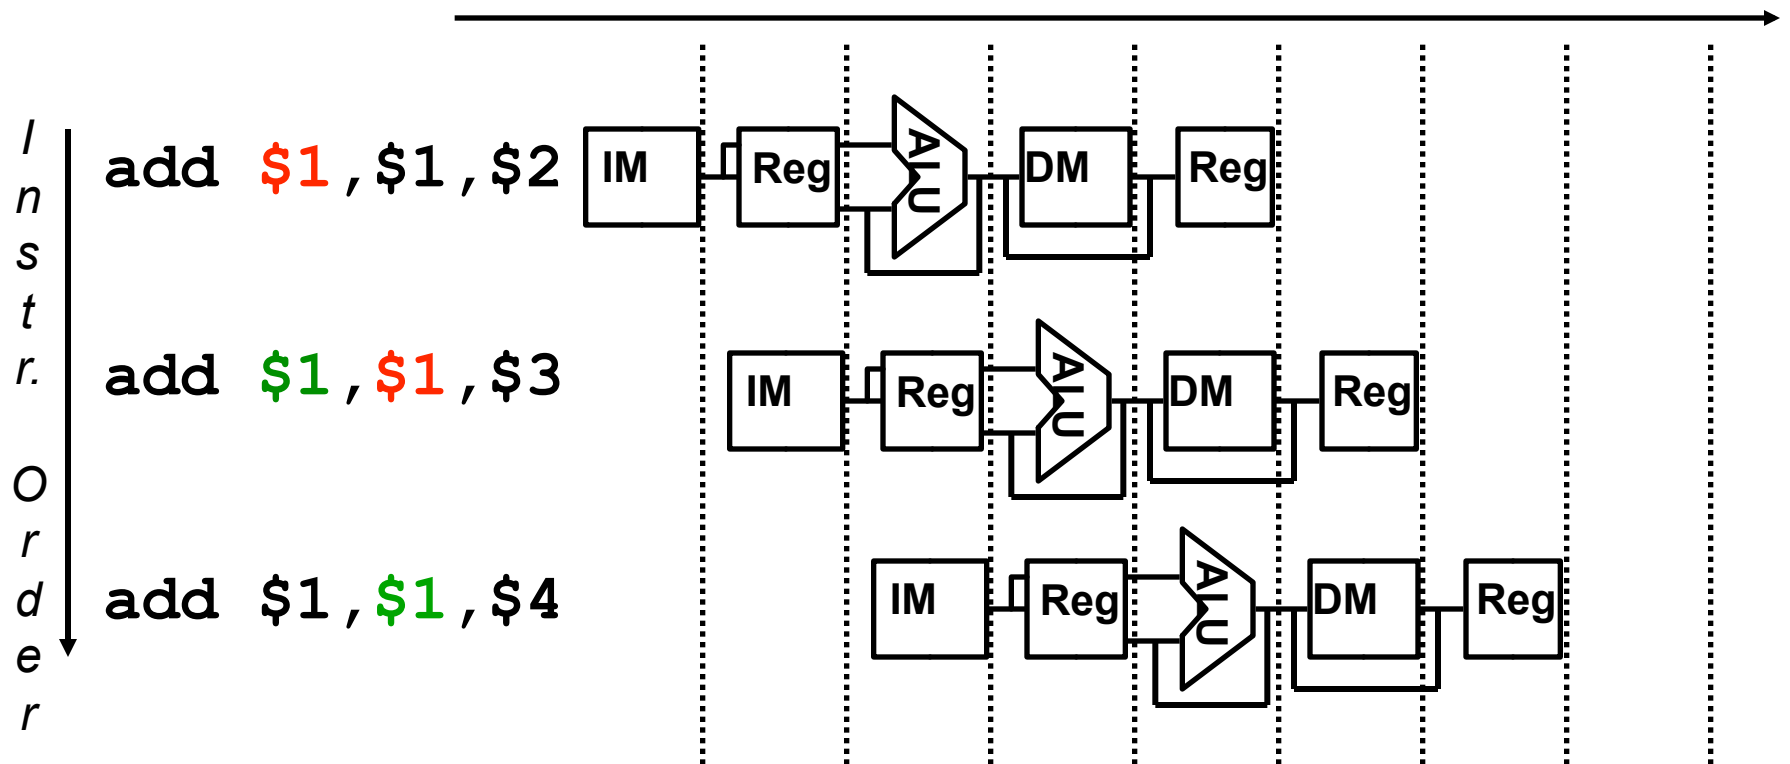
Forwards the result from the second previous instr. to either input of the ALU

Forwarding Illustration



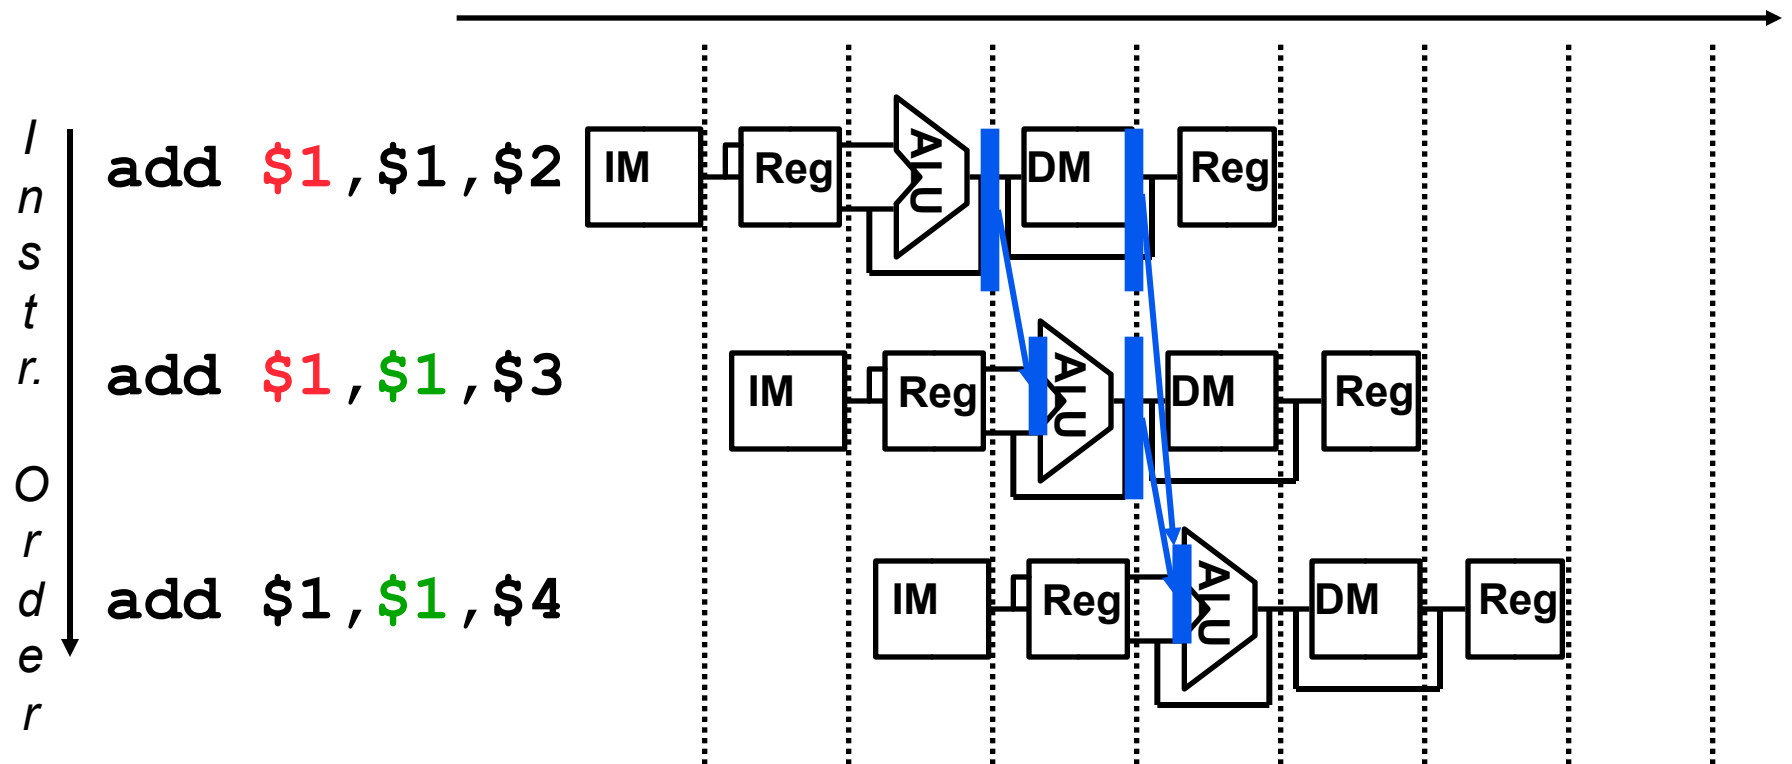
Yet Another Complication!

- Another potential data hazard can occur when there is a conflict between the result of the WB stage instruction and the MEM stage instruction – which should be forwarded?



Yet Another Complication!

- Another potential data hazard can occur when there is a conflict between the result of the WB stage instruction and the MEM stage instruction – which should be forwarded?



Corrected Data Forwarding Control Conditions

1. EX Forward Unit (result from ALU {add}):

```
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 EX
result from the
previous instr.
to A or B input
of the ALU

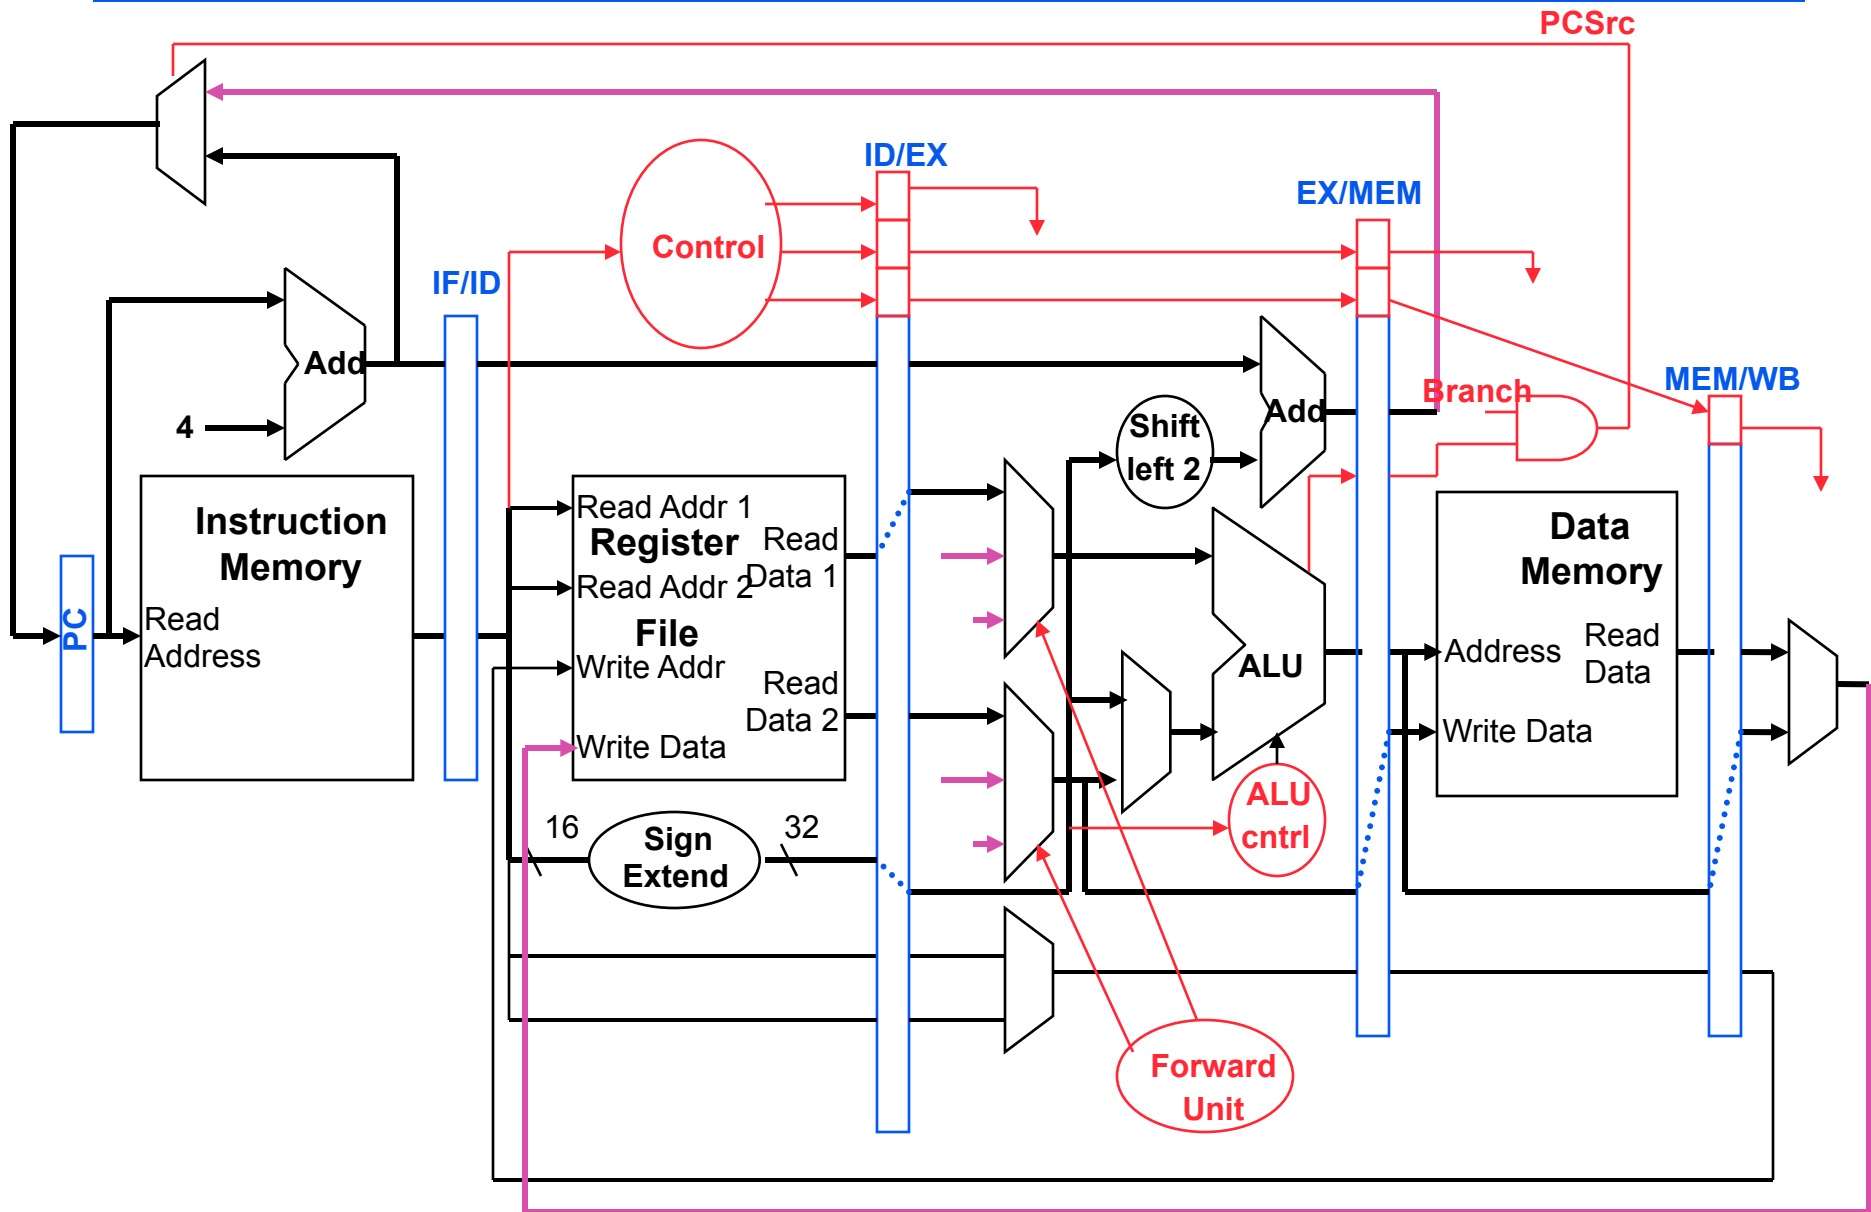
2. MEM Forward Unit (result from MEM {lw} or 2nd back ALU):

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

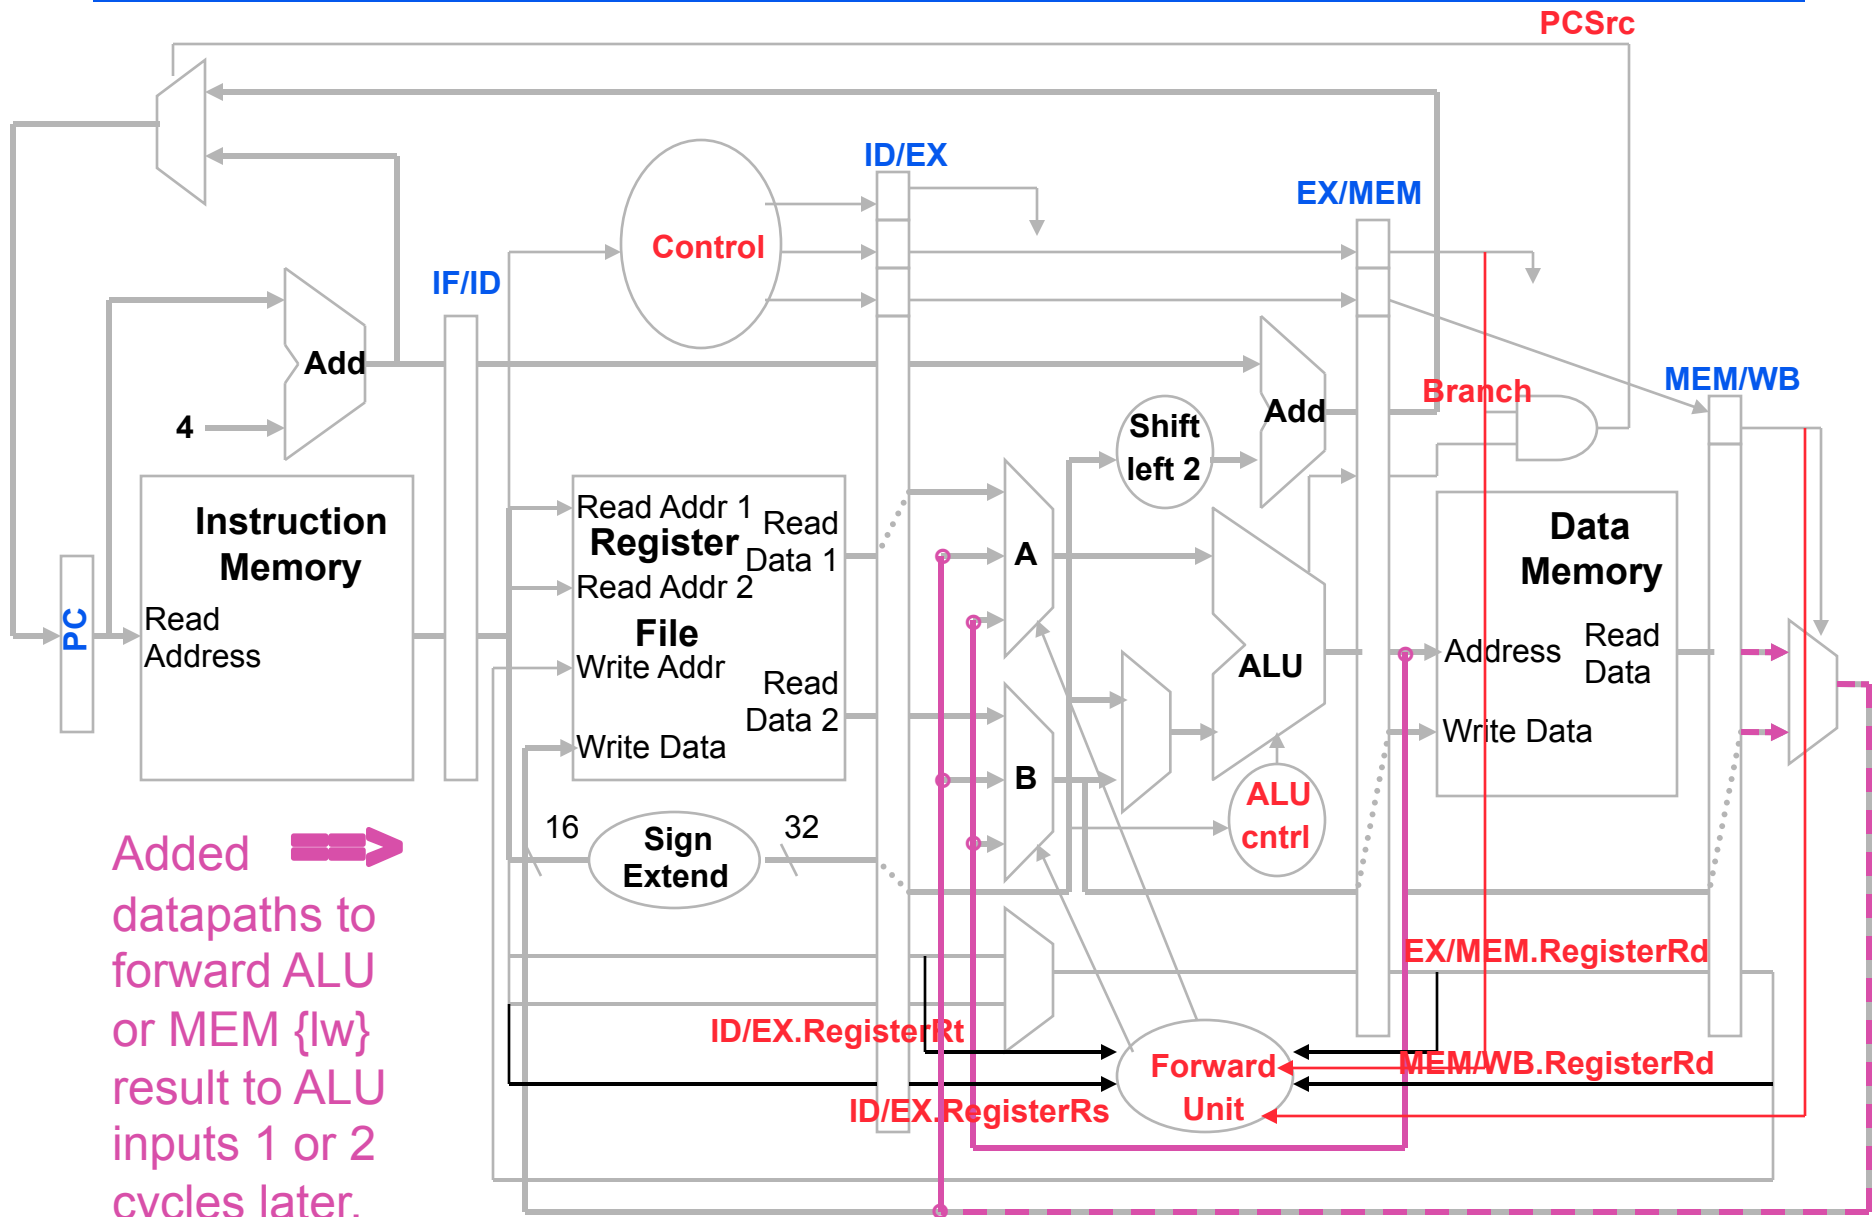
if (MEM/WB.RegWrite
and (MEM/WB.RegisterRd != 0)
and (EX/MEM.RegisterRd != ID/EX.RegisterRt))
    ForwardB = 01
```


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

Datapath with Forwarding Hardware



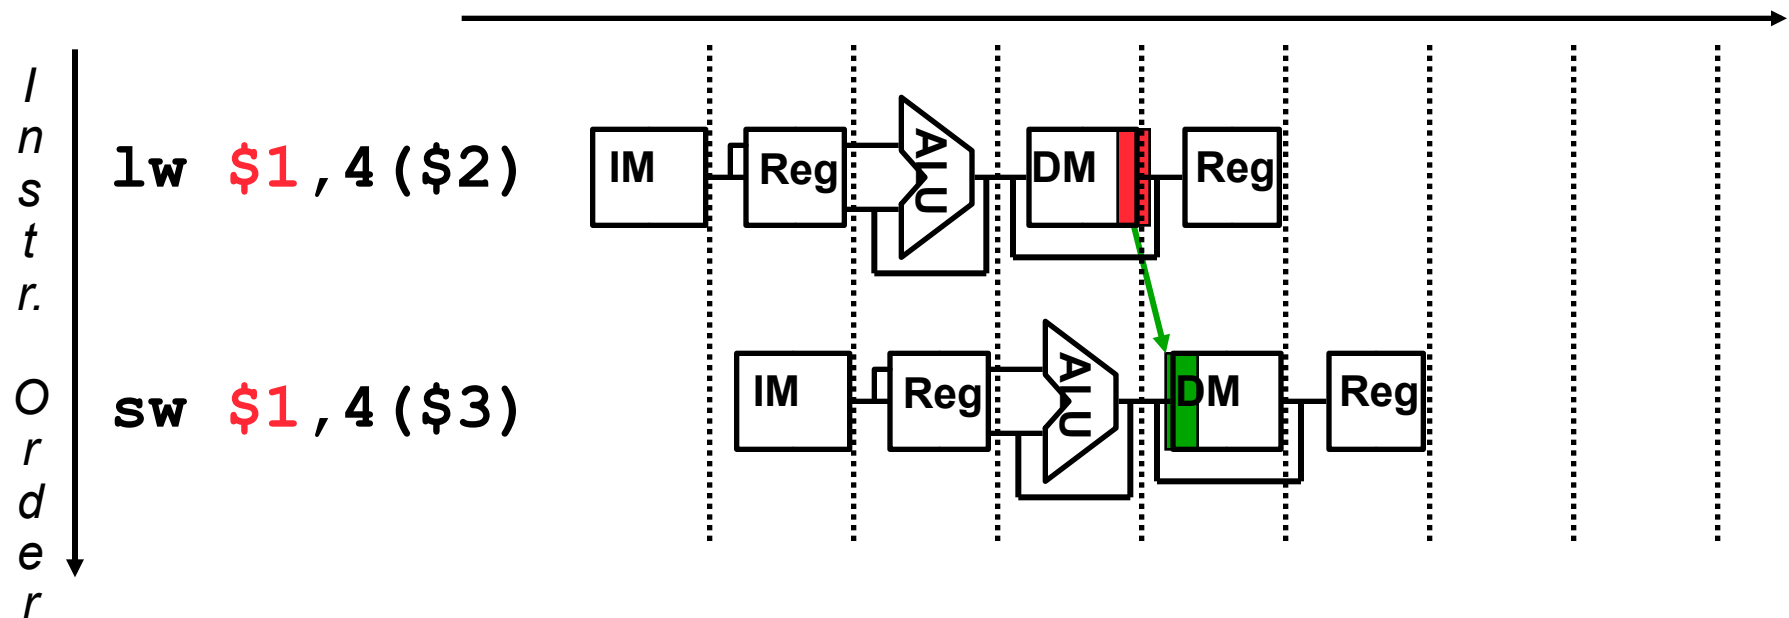
Datapath with Forwarding Hardware



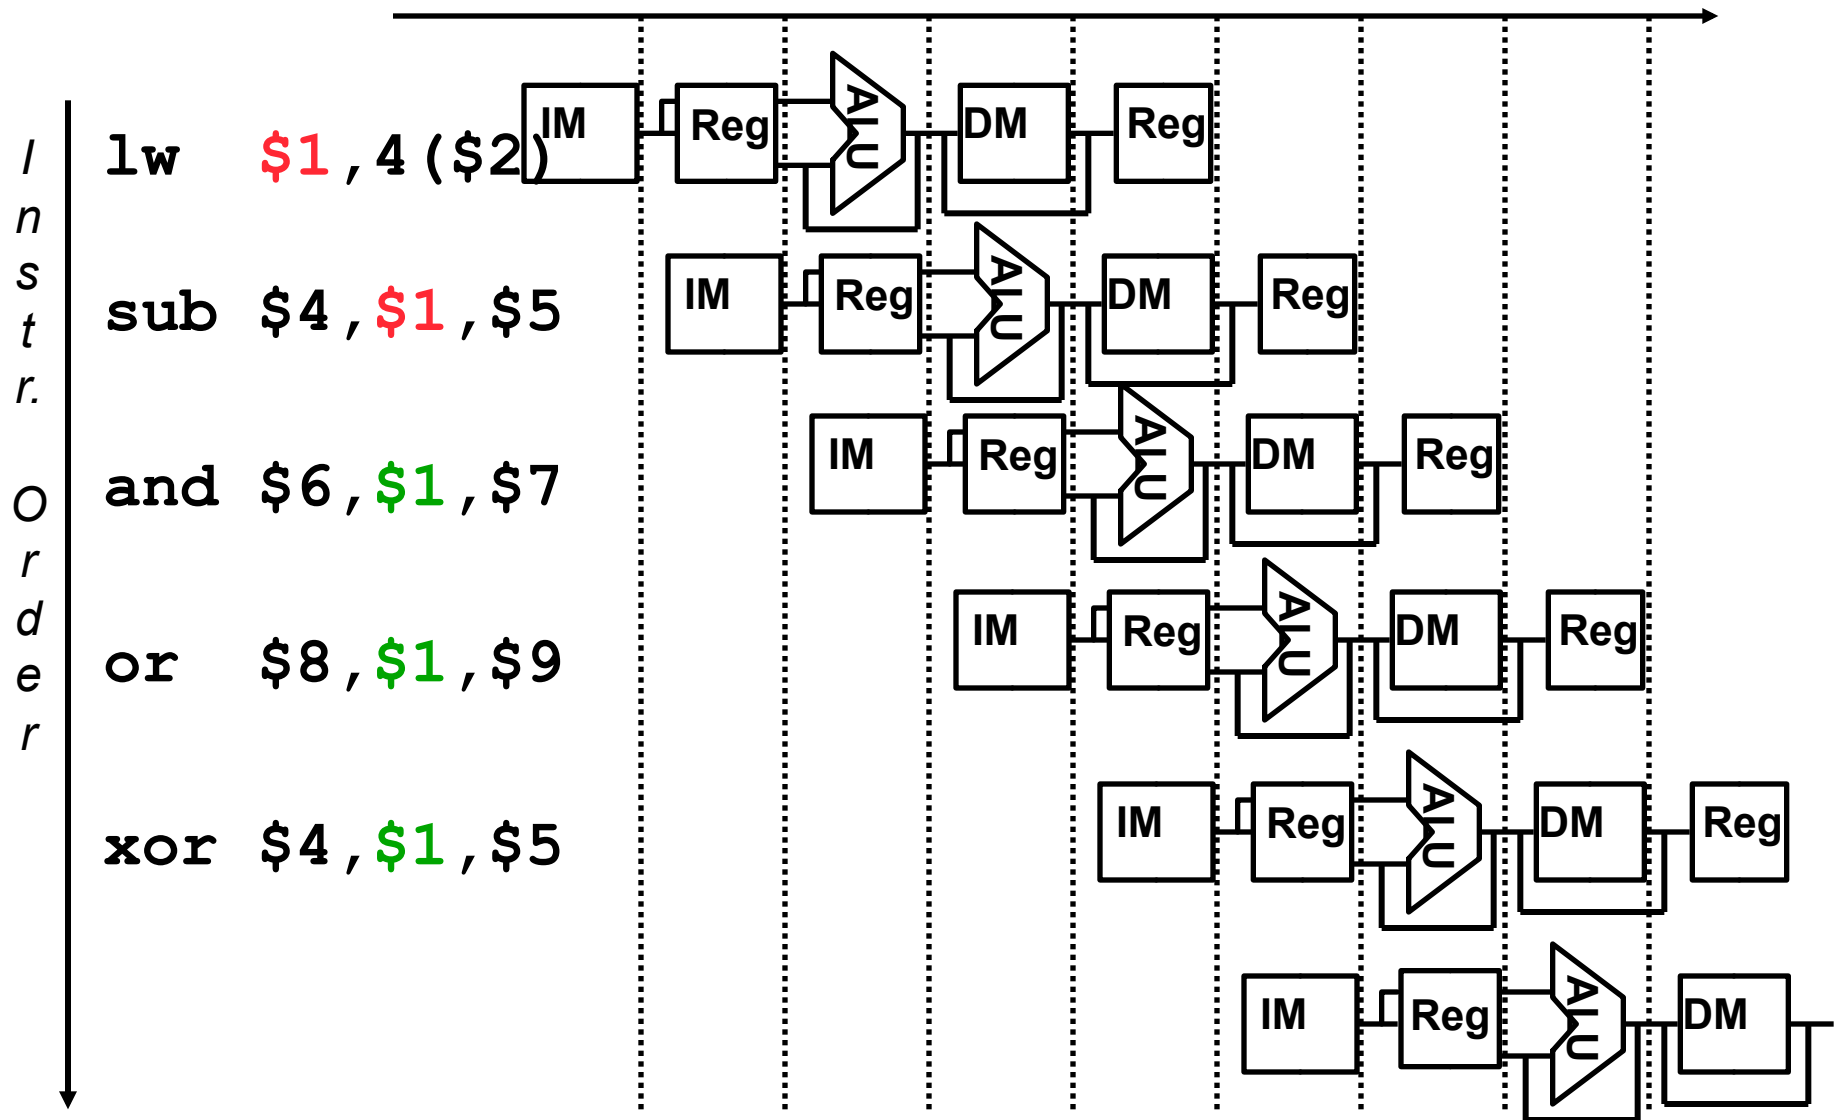
Added  datapaths to forward ALU or MEM {lw} result to ALU inputs 1 or 2 cycles later.

Memory-to-Memory Copies

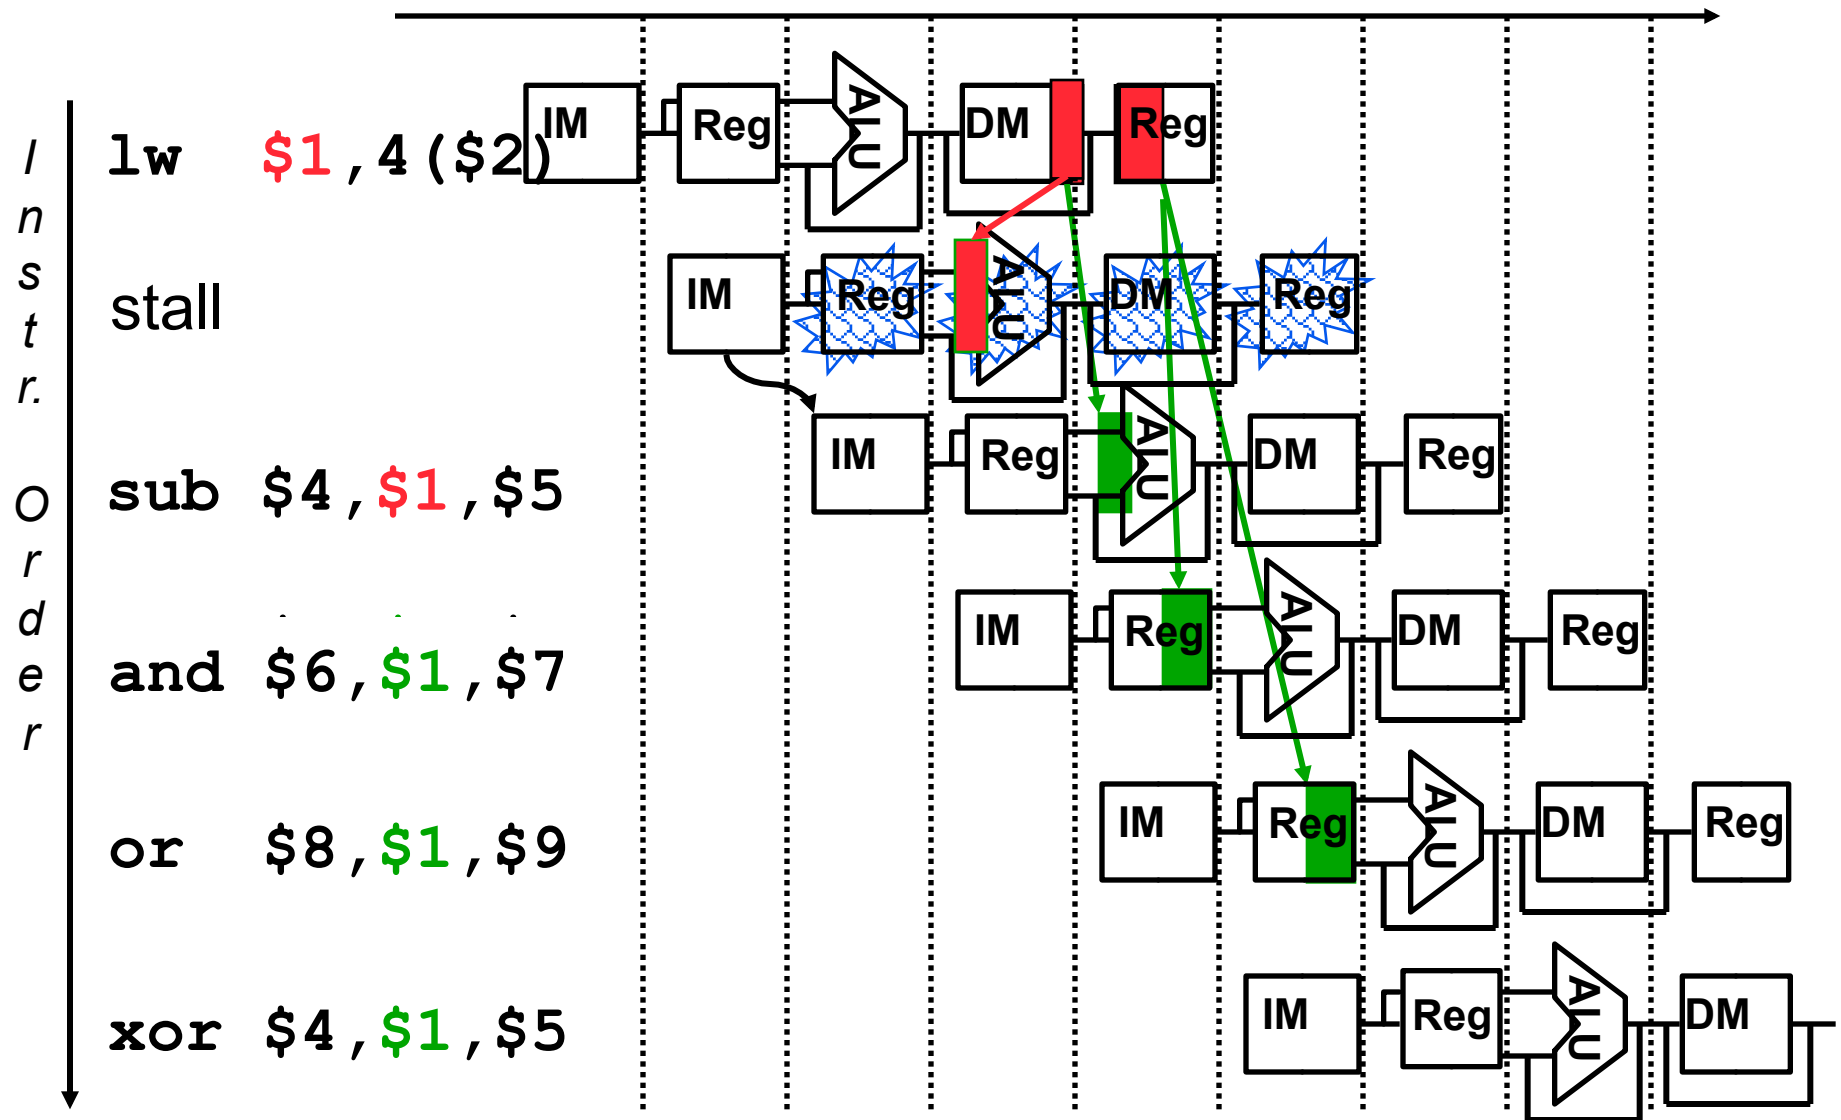
- ❑ For loads immediately followed by stores (memory-to-memory copies) can avoid a stall by adding forwarding hardware from the MEM/WB register to the data memory input.
 - Would need to add a Forward Unit and a mux to the MEM stage



Forwarding with Load-use RAW Data Hazards



Forwarding with Load-use RAW Data Hazards



❑ Will still need **one stall cycle** even with forwarding

Load-use Hazard Detection Unit

- ❑ Need a Hazard detection Unit in the ID stage that inserts a stall between the load and its use

1. ID Hazard detection Unit:

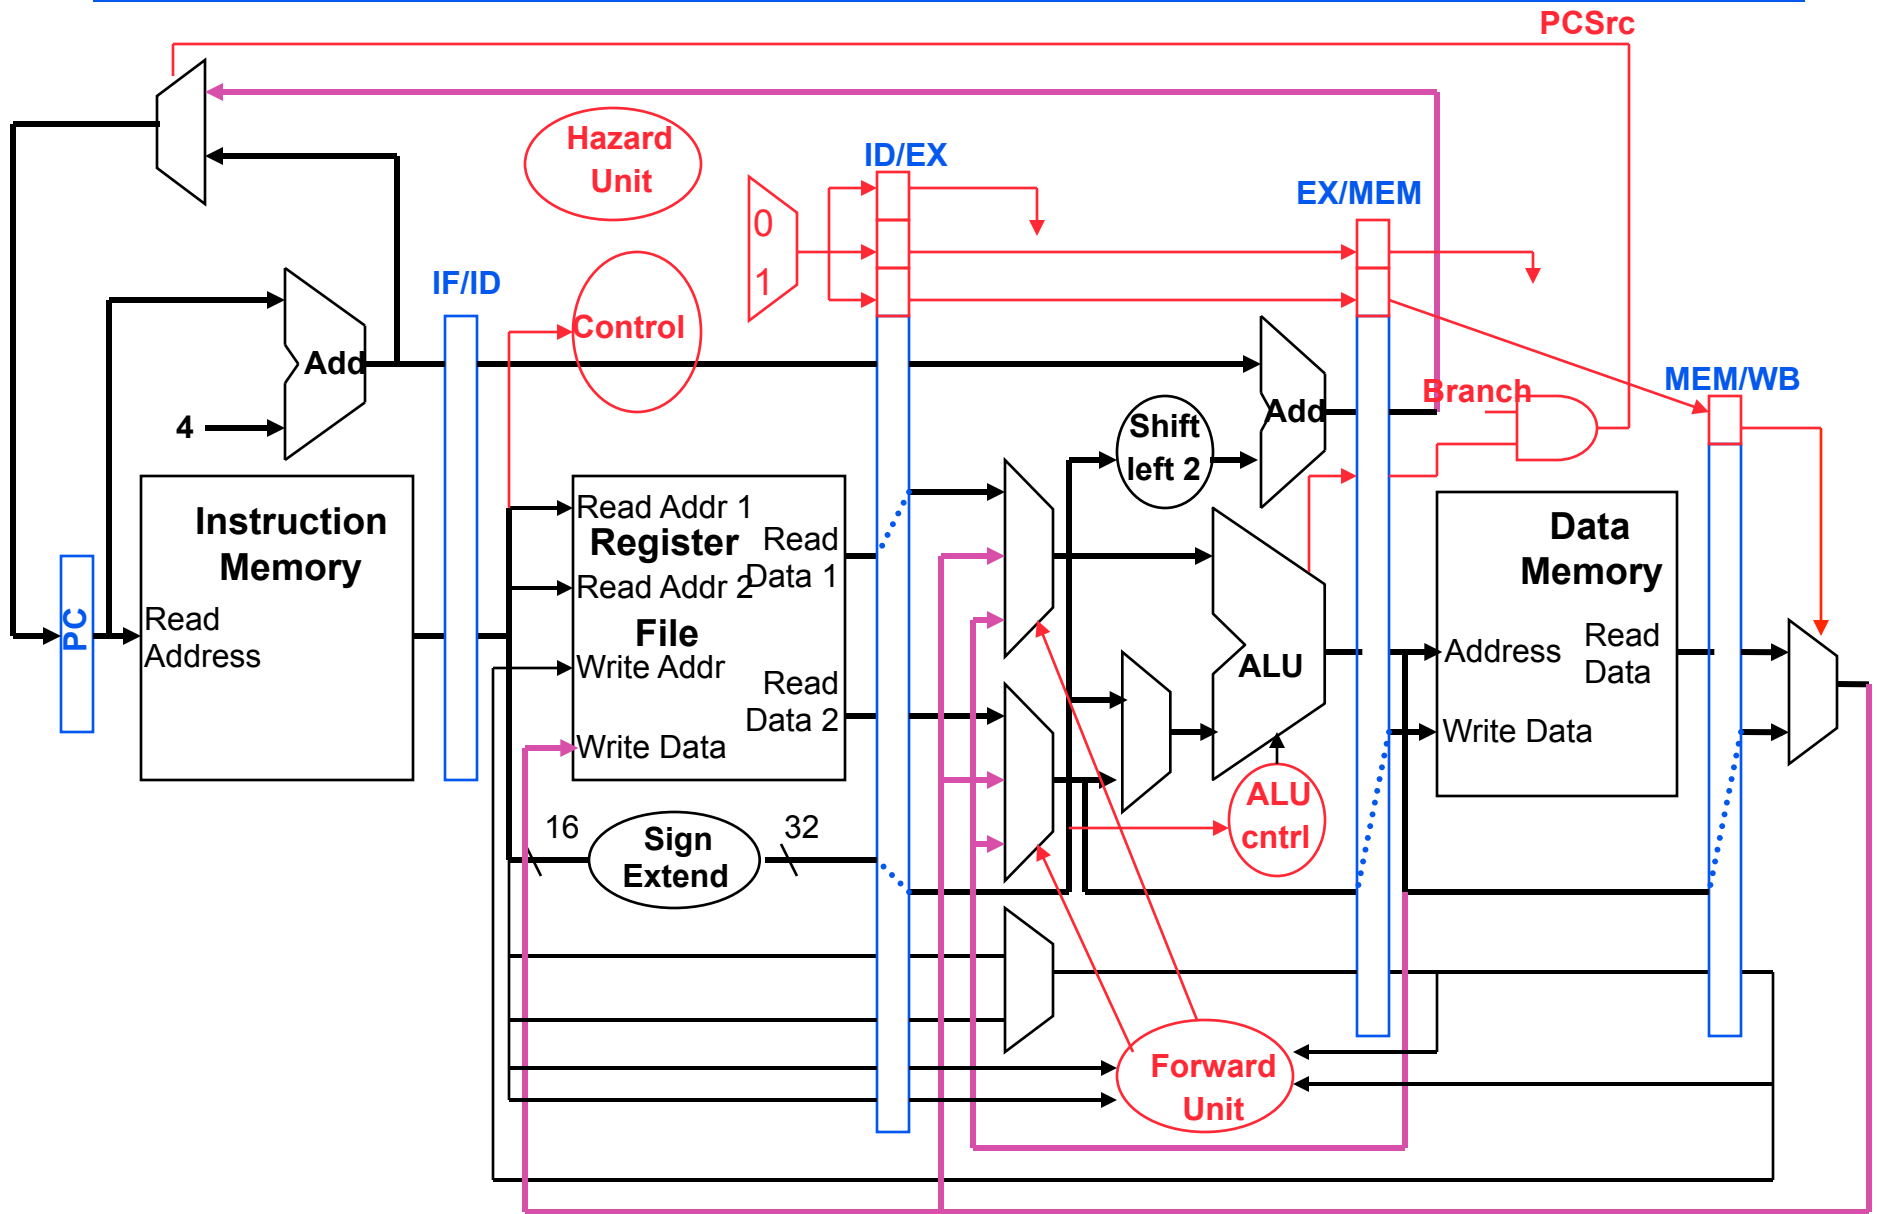
```
if (ID/EX.MemRead
and ((ID/EX.RegisterRt = IF/ID.RegisterRs)
or (ID/EX.RegisterRt = IF/ID.RegisterRt)))
stall the pipeline
```

- ❑ The first line tests $ID/EX.MemRead$ to see if the instruction now in the EX stage is a lw ; the next two lines check to see if the destination register $RegisterRt$ of the lw matches either source register $IF/ID.RegisterRs, \dots Rt$ of the instruction in the ID stage (the load-use instruction)
- ❑ After this one cycle stall, the forwarding logic can handle the remaining data hazards

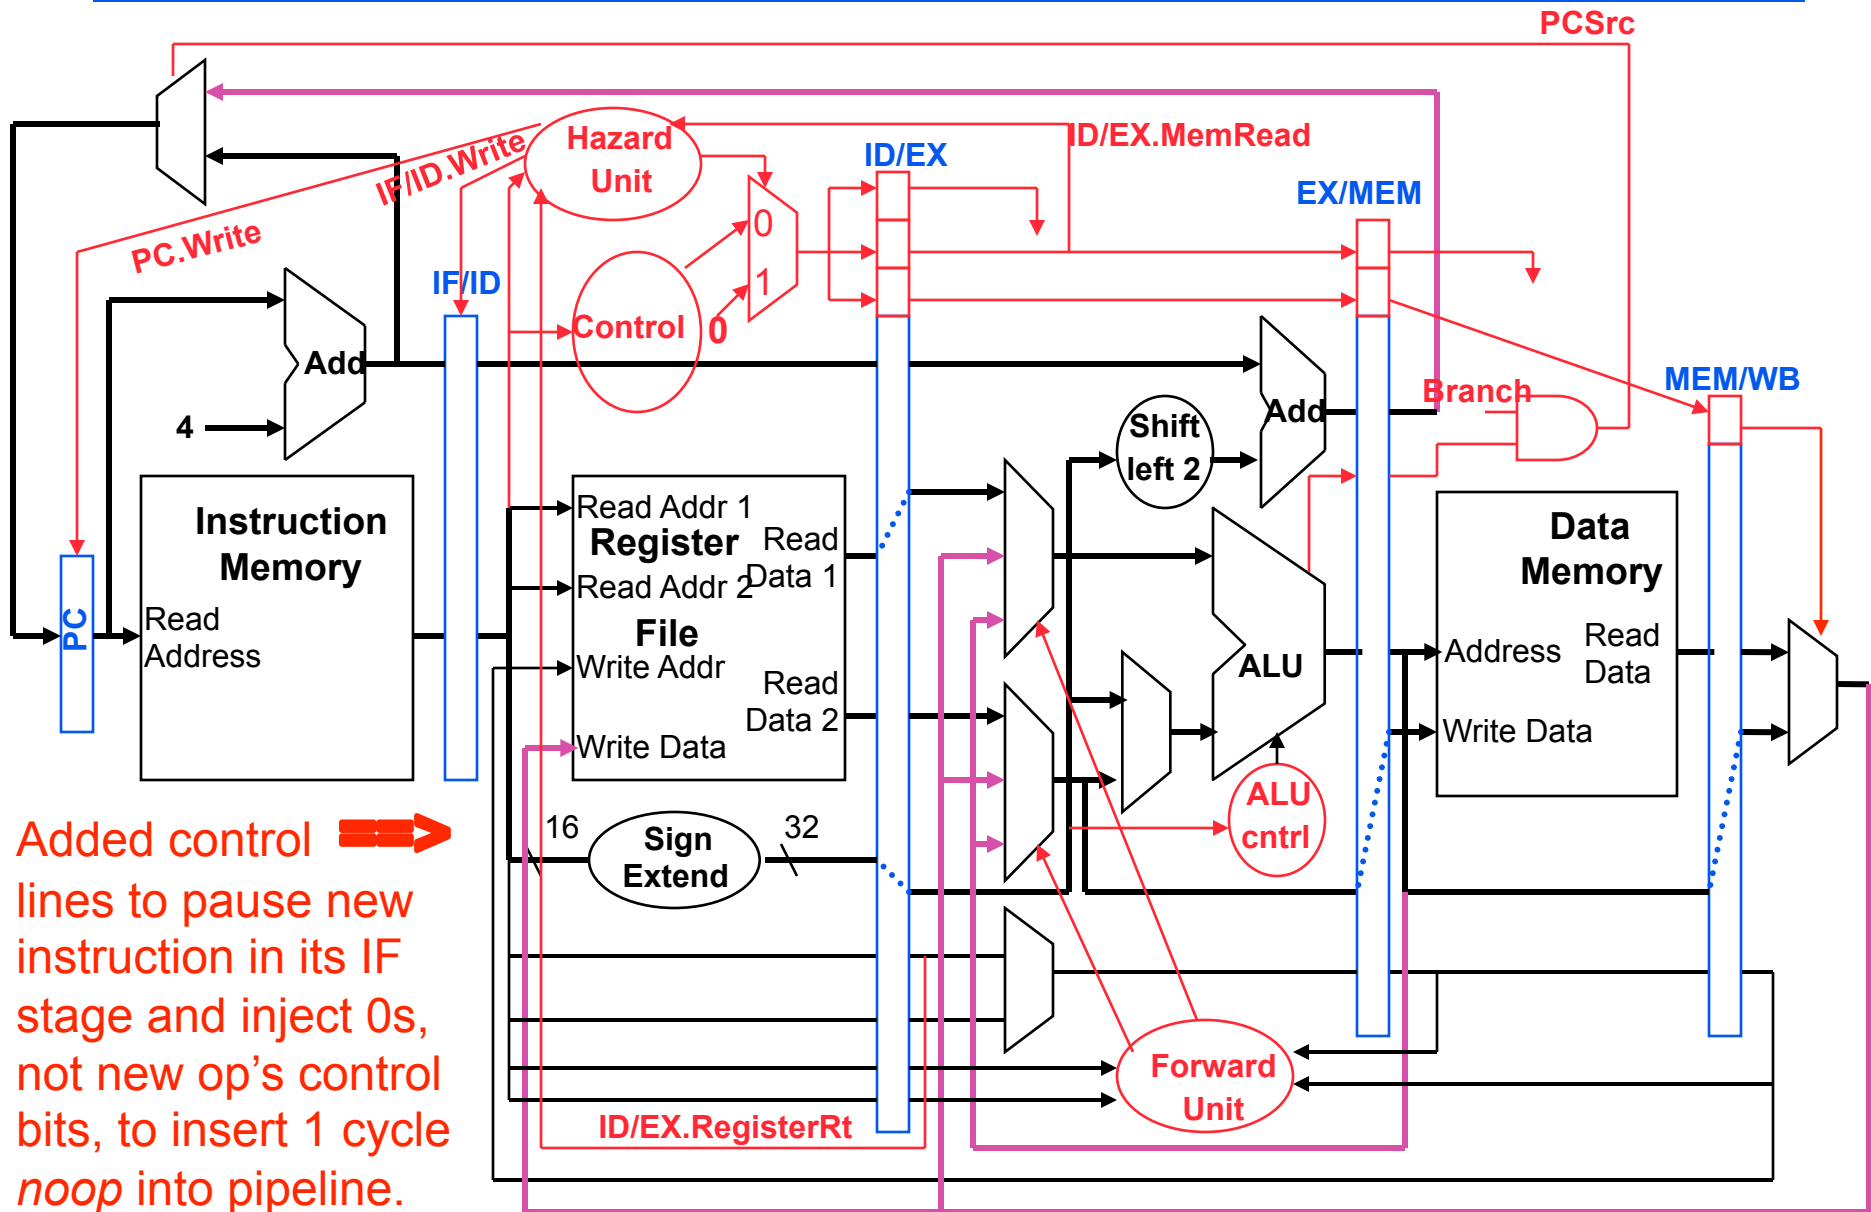
Hazard/Stall Hardware

- ❑ Along with the Hazard Unit, we have to implement the stall
- ❑ Prevent the instructions in the IF and ID stages from progressing down the pipeline – done by preventing the PC register and the IF/ID pipeline register from changing
 - Hazard detection Unit controls the writing of the PC (`PC.write`) and IF/ID (`IF/ID.write`) registers
- ❑ Insert a “bubble” **between** the `lw` instruction (in the EX stage) and the load-use instruction (in the ID stage) (i.e., insert a `noop` in the execution stream)
 - Set the control bits in EX, MEM, and WB control fields of the ID/EX pipeline register to 0 (`noop`). The Hazard Unit controls a mux that chooses either the instruction’s control values or the 0’s for a stall.
- ❑ Let the `lw` instruction and the instructions after it in the pipeline (before it in the code) proceed normally down the pipeline

Need Hazard_Detect & Stall HW for Load-ALU_Use (Here Have Jump & Branch Completion in Stage_4)



Adding Hazard/Stall Hardware for Load-ALU_Use



The Last and Worst Hazard Type - Control Hazards

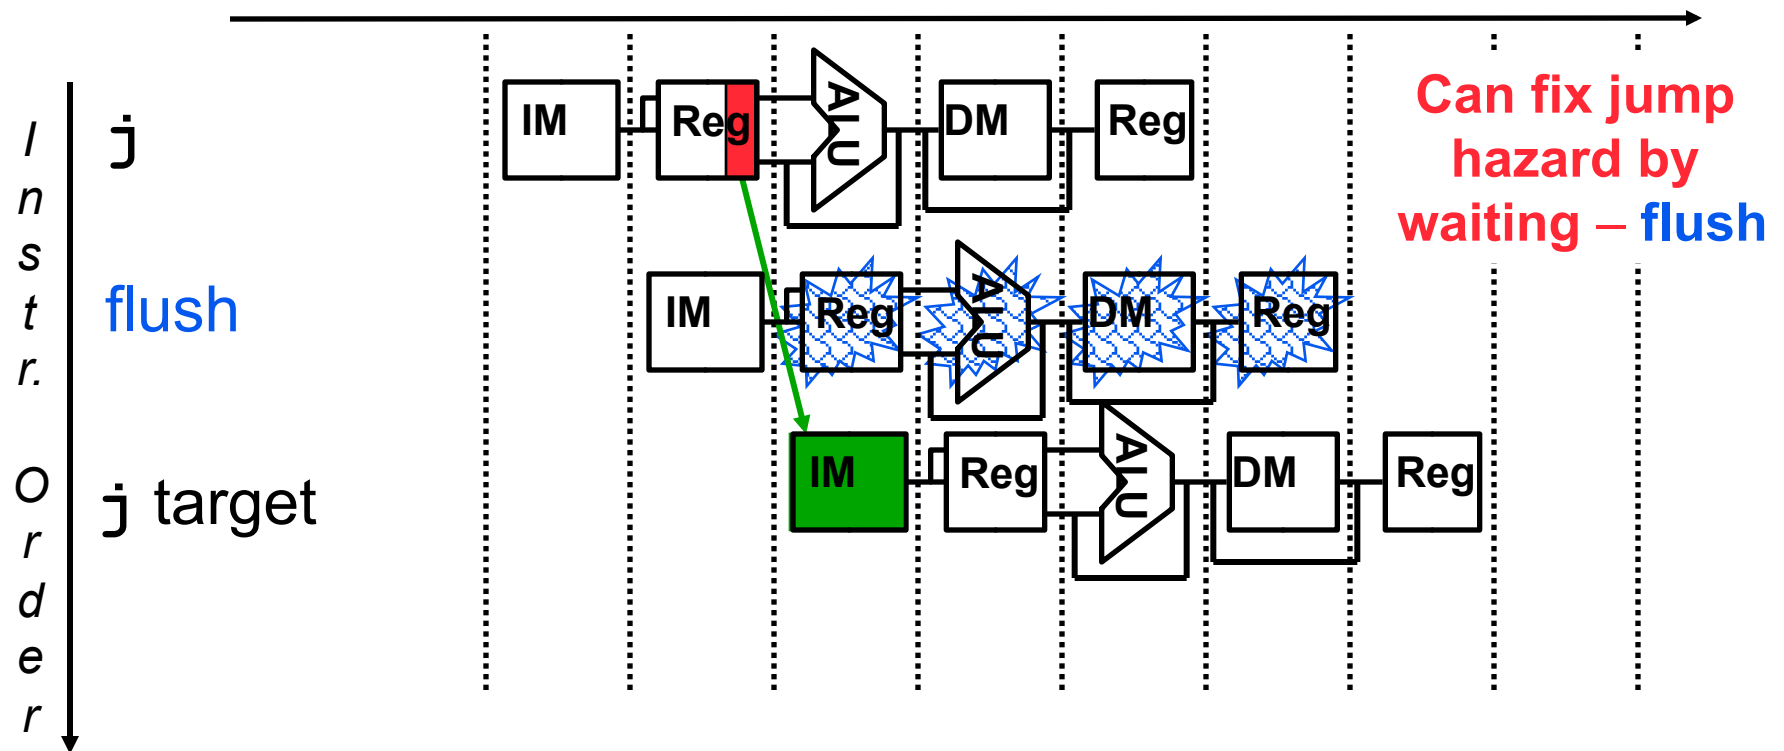
- ❑ Happen when flow of instruction addresses is non-sequential (i.e., not $PC = PC + 4$); caused by flow-changing instructions:
 - Unconditional branches (`j`, `jal`, `jr`)
 - Conditional branches (`beq`, `bne`)
 - Exceptions

- ❑ Possible approaches to handling control hazards
 - Stall (bad - impacts CPI & makes code execution take more time)
 - Move branch-or-not decision point as early in the pipeline as possible, thereby reducing the number of stall cycles
 - Delay branch decisions until after a following instruction so the fetch at $(PC + 4)$ is not wasted (needs compiler to fill “delay slot”)
 - Predict whether & where a branch will occur and hope for the best !

- ❑ Control hazards occur less frequently (15% of ops) than data hazards (~25%), but *no simple method* avoids control hazards as effectively as forwarding precludes data hazards

Simplest Control Hazards - Jumps Incur One Stall

- ❑ Jumps not decoded until ID, so at least 1 **flush** needed
 - To flush, set `IF.Flush` to zero the instruction field of the IF/ID pipeline register (turning it into a `noop`)

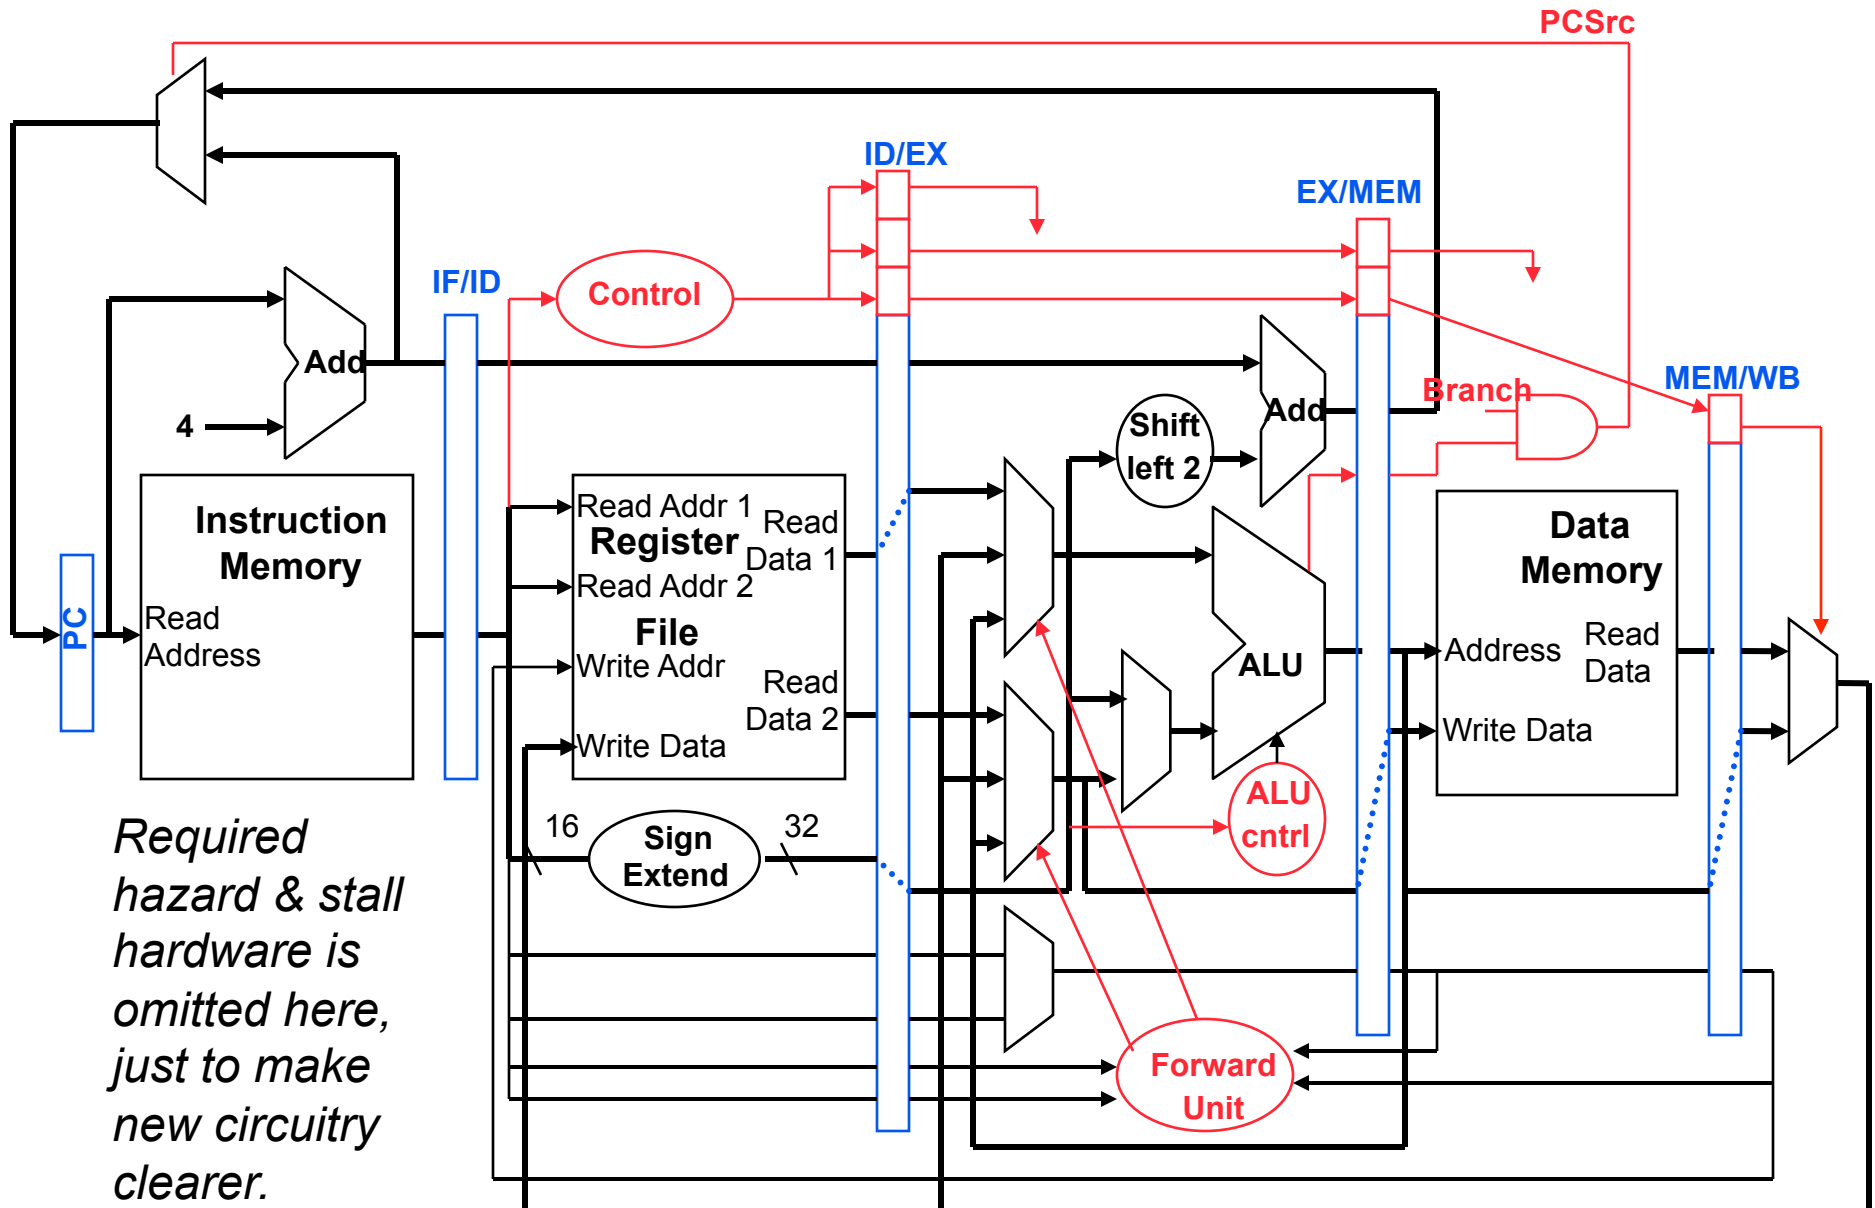


- ❑ Fortunately, jumps are very infrequently executed; only 3% of the SPECint instruction mix are jumps

Two “Types” of Stalls

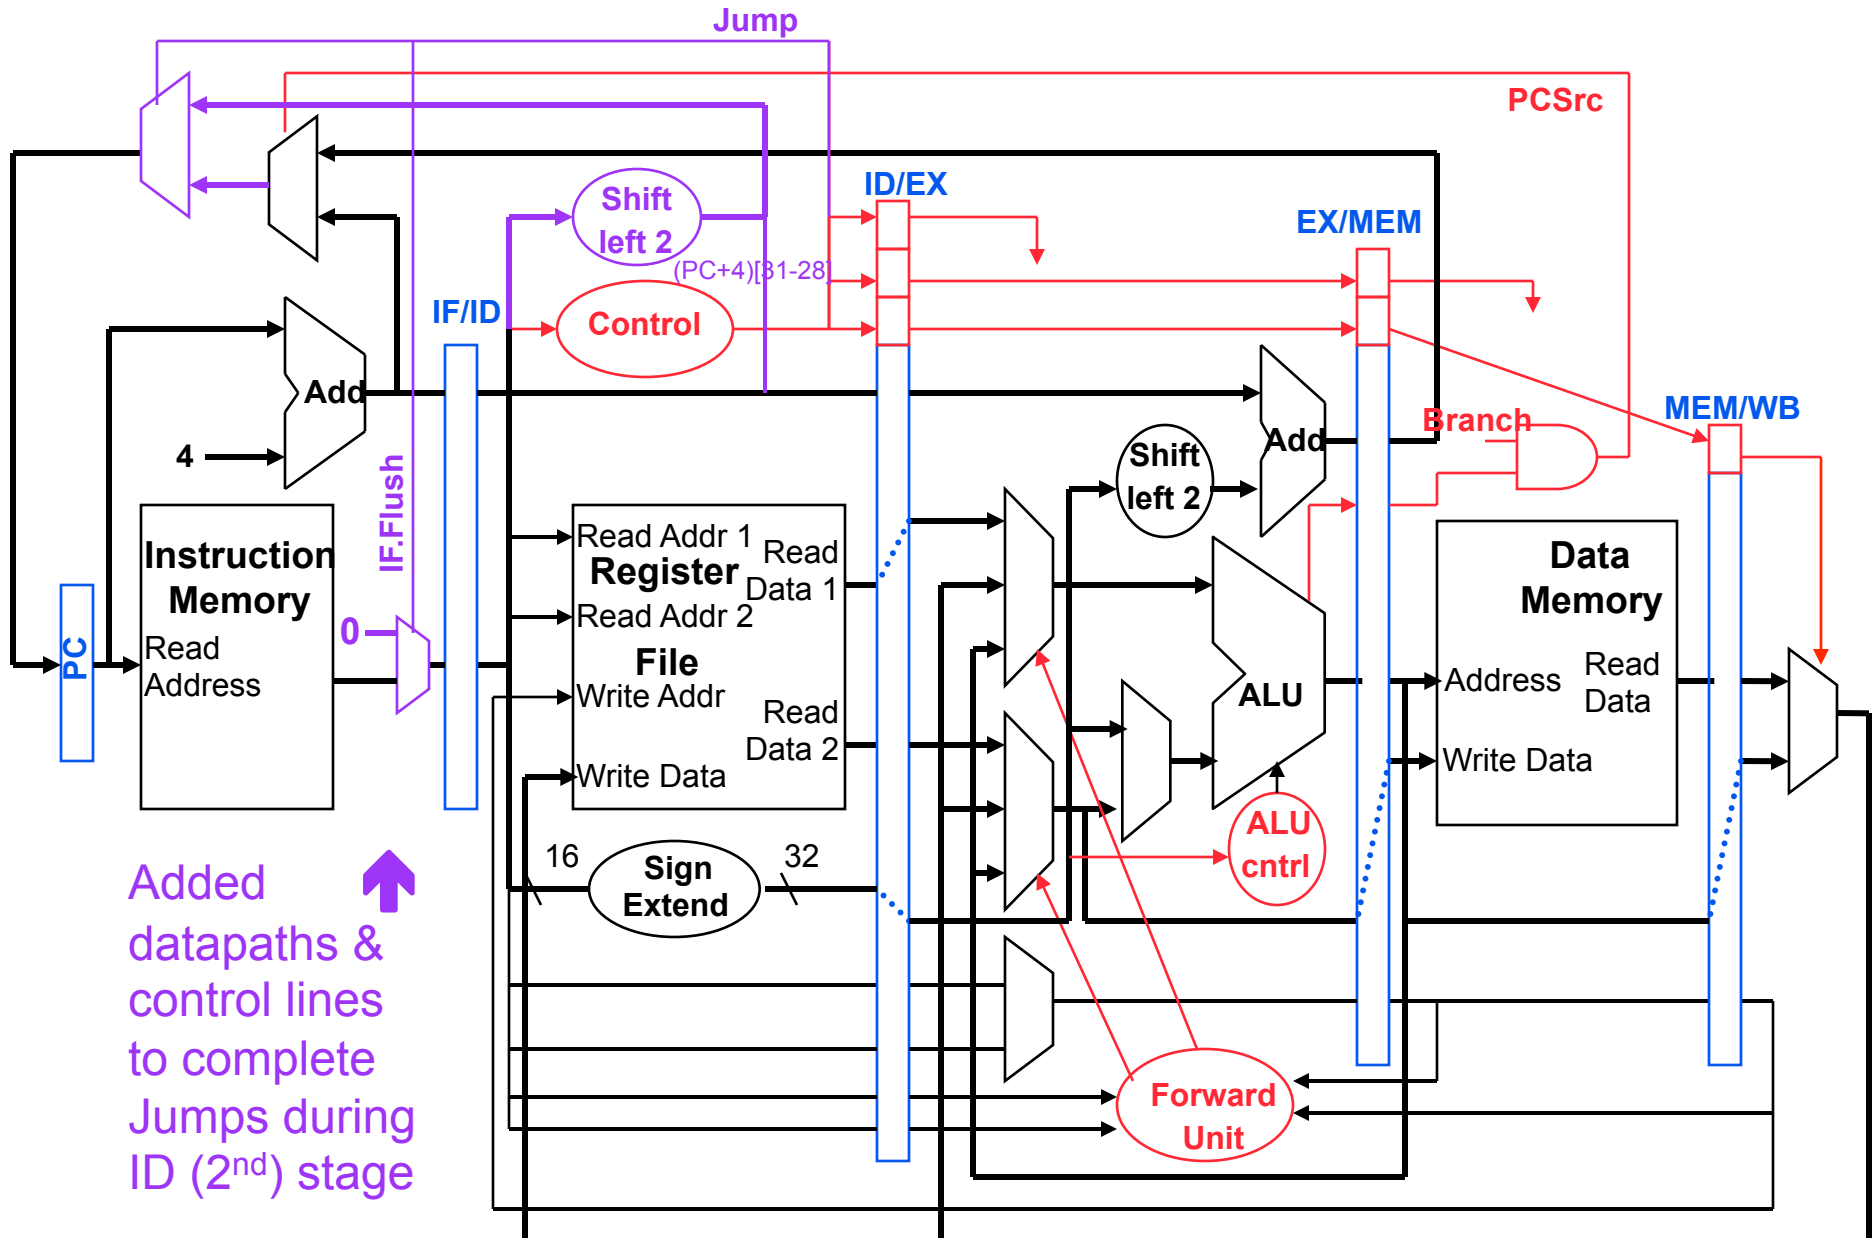
- ❑ Noop instruction (or bubble) **inserted** between two instructions in the pipeline (as done for load-use situations)
 - Keep the instructions *earlier* in the pipeline (later in the code) from progressing down the pipeline for one cycle (“bounce” them in place by using control signals that prevent changes to PC & IF/ID registers)
 - Insert `noop` by zeroing control bits in the pipeline register at the appropriate stage
 - Let the instructions later in the pipeline (earlier in the code) progress normally down the pipeline
- ❑ Flushes (or instruction squashing) where an instruction in the pipeline is **replaced** with a `noop` instruction (as done for instructions located sequentially after j instructions)
 - Zero the control bits that record results of each instruction to be flushed

MEM (4th) Stage Jump & Branch Hardware



Required hazard & stall hardware is omitted here, just to make new circuitry clearer.

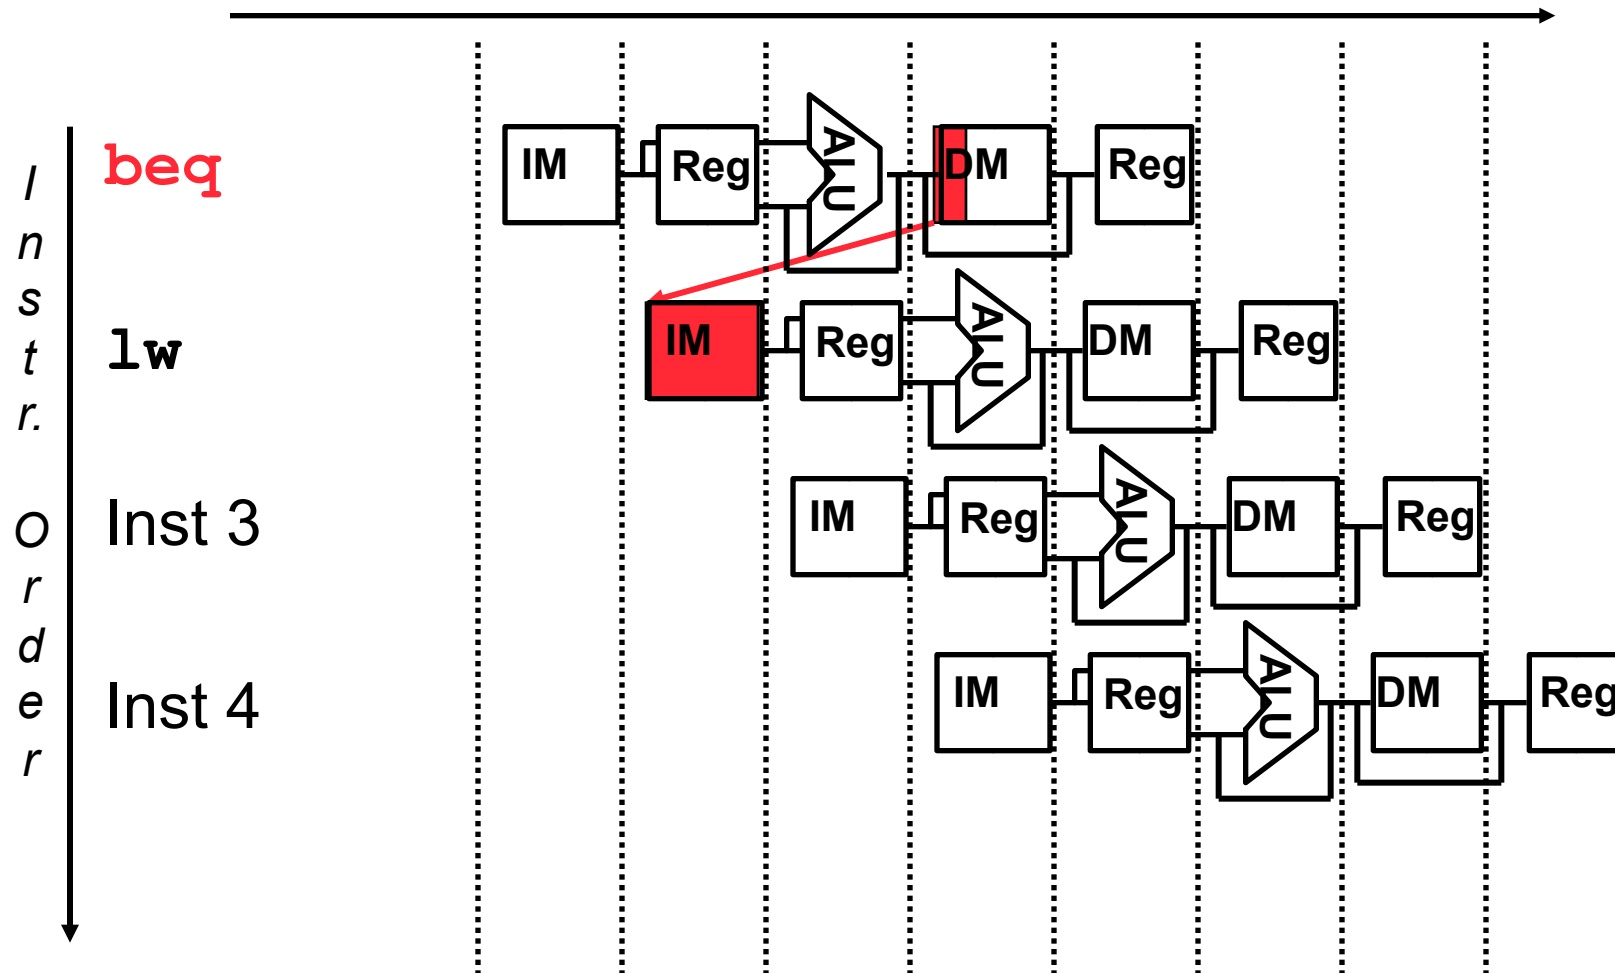
Added Hardware to Support ID (2nd) Stage Jumps



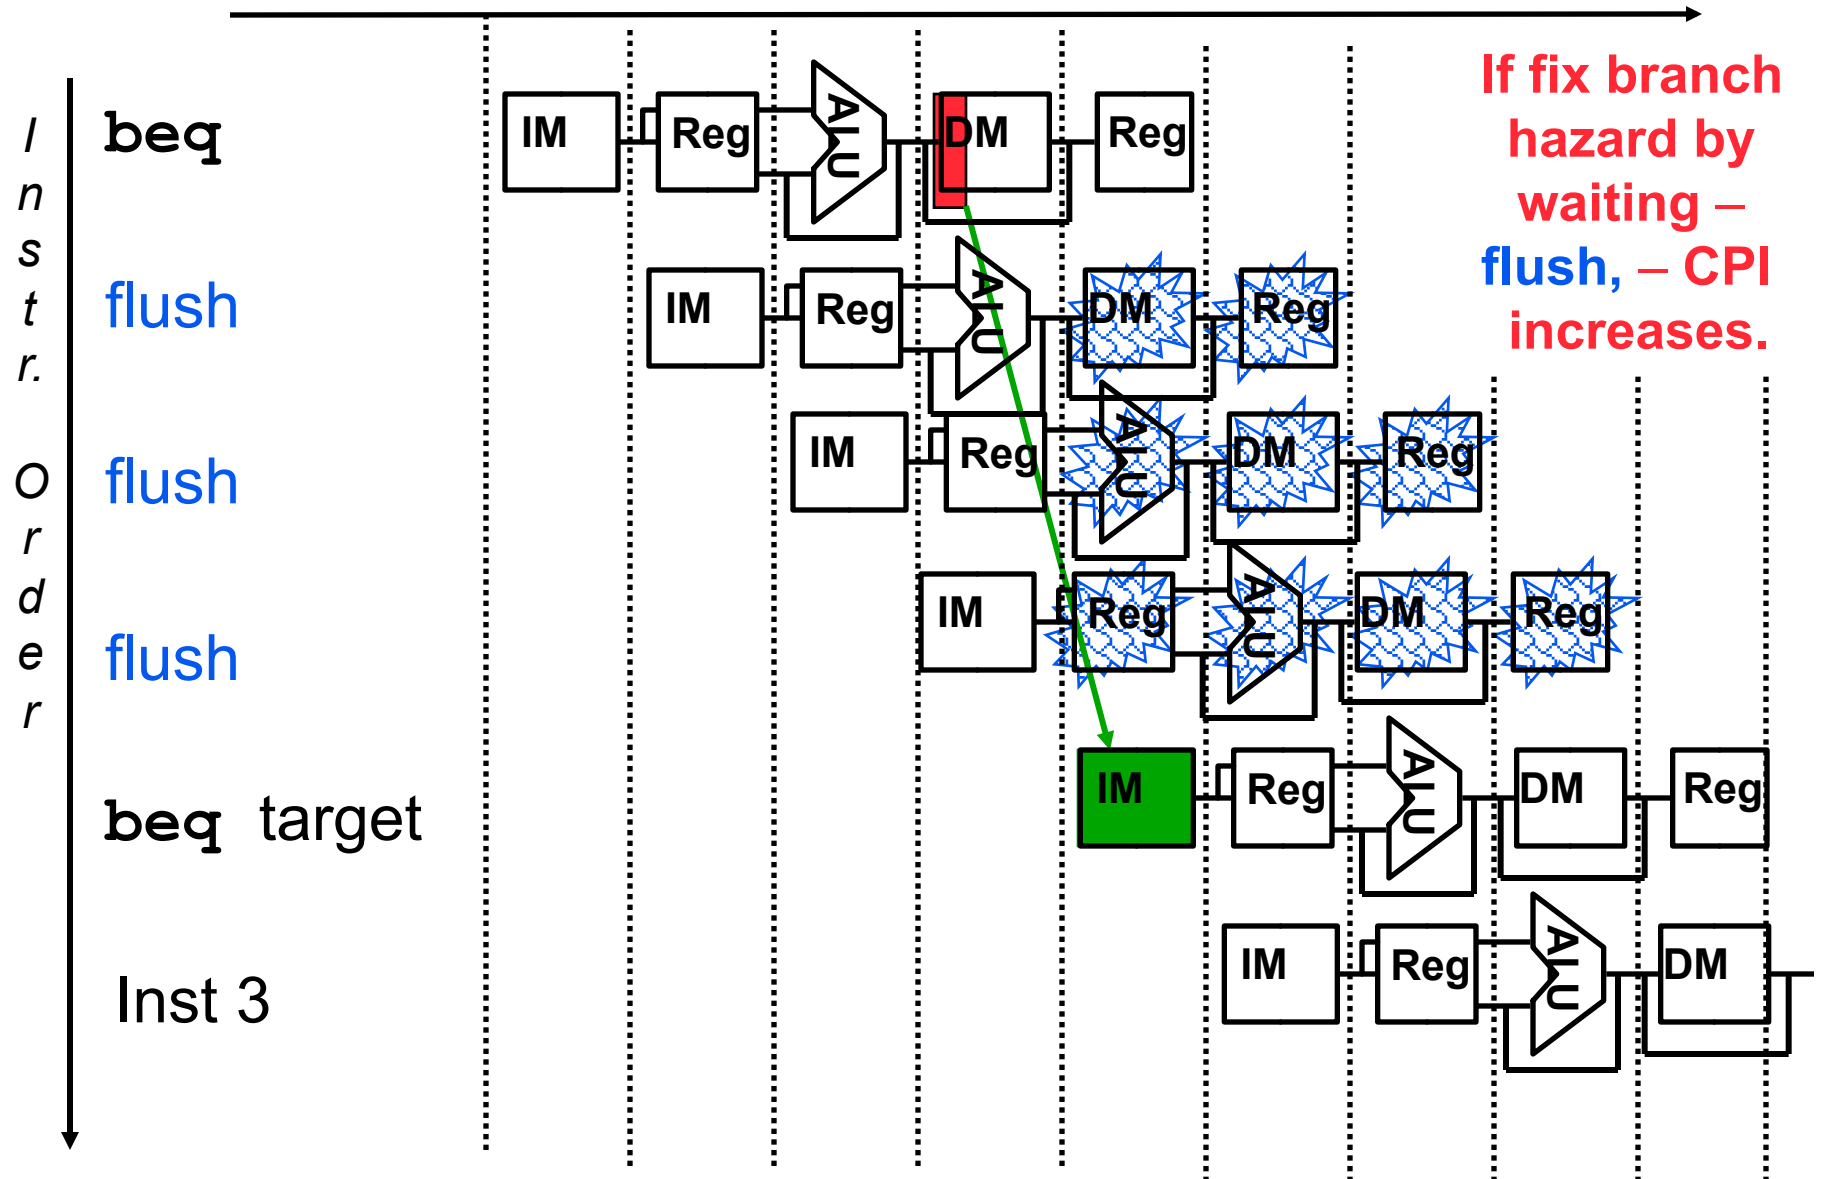
Added datapaths & control lines to complete Jumps during ID (2nd) stage

Review: Branch Instr's Cause Bad Control Hazards

- Dependencies backward in time cause **hazards**;
branches common, roughly 15% of all instructions

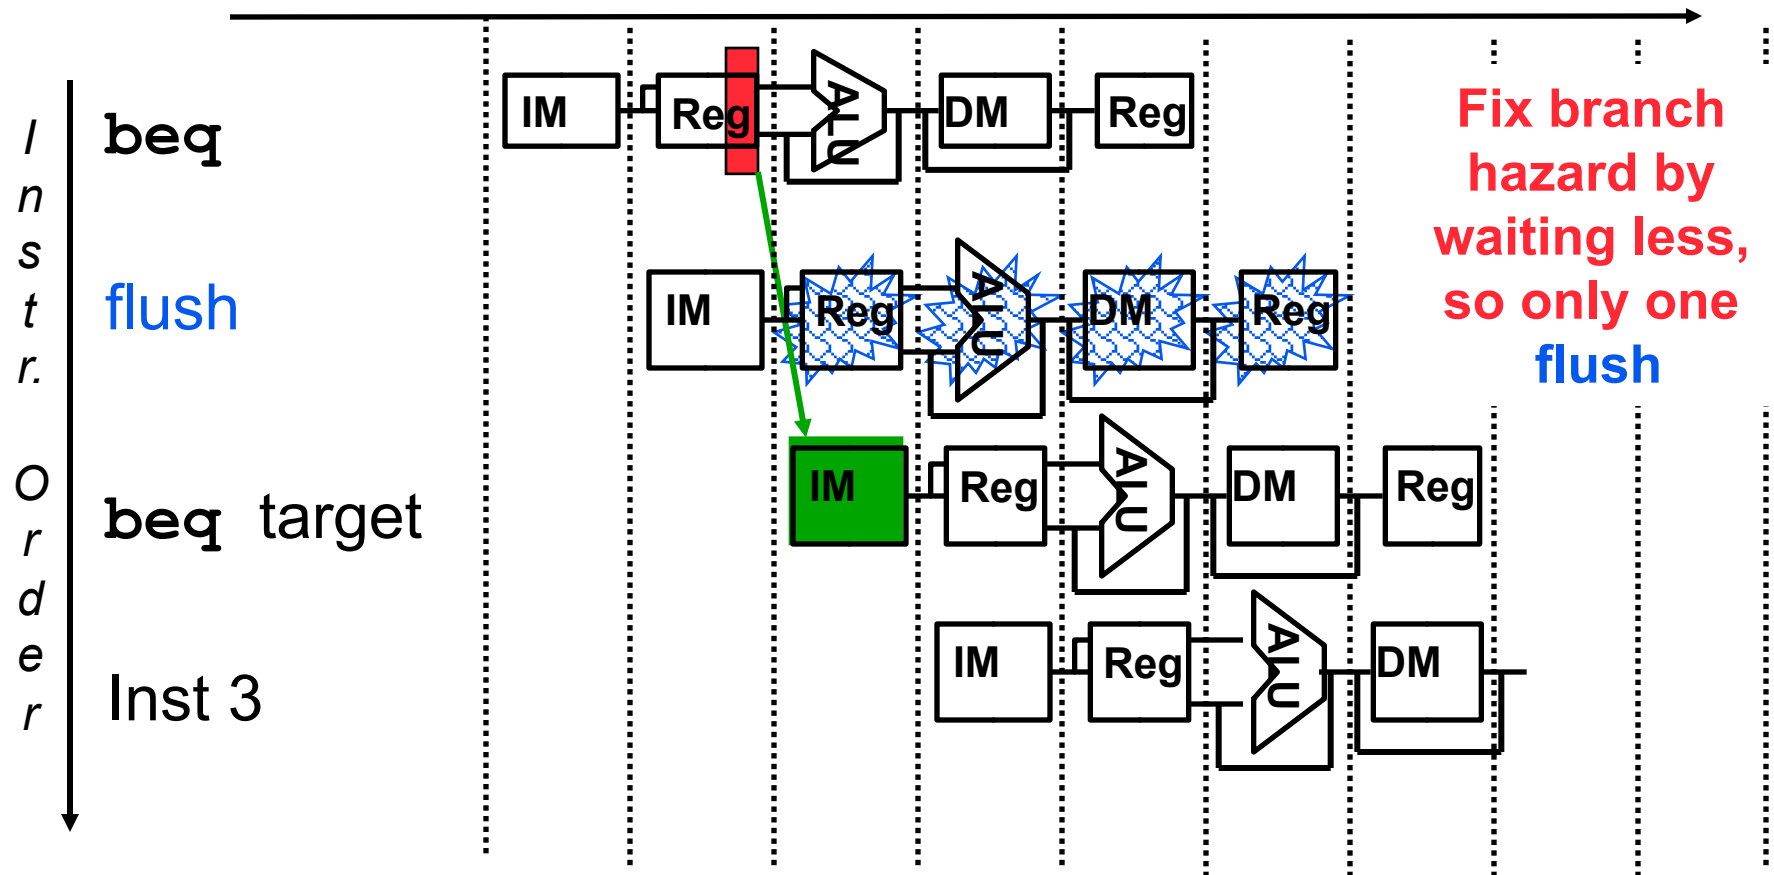


One Way to “Fix” a Branch Control Hazard



A Better Way to “Fix” a Branch Control Hazard

- Move branch decision hardware back to as **early** in the pipeline as possible – during the decode (2nd) stage



Reducing the Delay of Branches

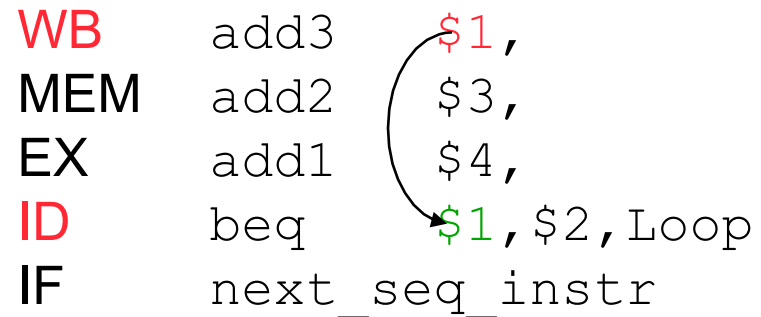
- ❑ Move the branch decision hardware back to the EX stage
 - Reduces the number of stall (flush) cycles to two, not early enough
 - Adds an `and` gate and a `2x1 mux` to the EX timing path

- ❑ Add hardware to compute the branch target address and evaluate the branch decision to the ID stage
 - Reduces the number of stall (flush) cycles to one (as already shown for jumps)
 - But now need to add **forwarding hardware** in ID stage
 - Computing branch target address can be done in parallel with RegFile read (done for all instructions – only used when needed)
 - Comparing the registers cannot be done until after RegFile read, so comparing and updating the PC adds a mux, a comparator, and an `and` gate to the ID timing path

- ❑ For deep pipelines, branch decision points are even *later* in the pipeline than stage 2, incurring more stalls

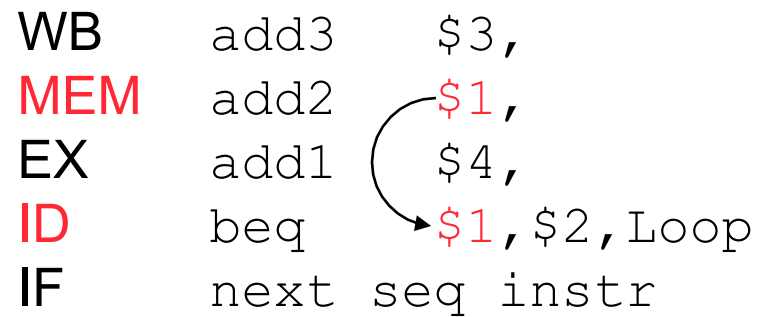
ID Branch Forwarding Issues

- MEM/WB “forwarding” is taken care of by the normal RegFile write-before-read sequence, if produce input to beq 3 instructions earlier.



- Need to forward from the EX/MEM pipeline latch to the ID stage comparison hardware if produce input only 2 instructions earlier.

Aligned \uparrow \downarrow stages



```

if (IDcontrol.Branch
and (EX/MEM.RegisterRd != 0)
and (EX/MEM.RegisterRd = IF/ID.RegisterRs))
    ForwardC = 1
if (IDcontrol.Branch
and (EX/MEM.RegisterRd != 0)
and (EX/MEM.RegisterRd = IF/ID.RegisterRt))
    ForwardD = 1
  
```

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

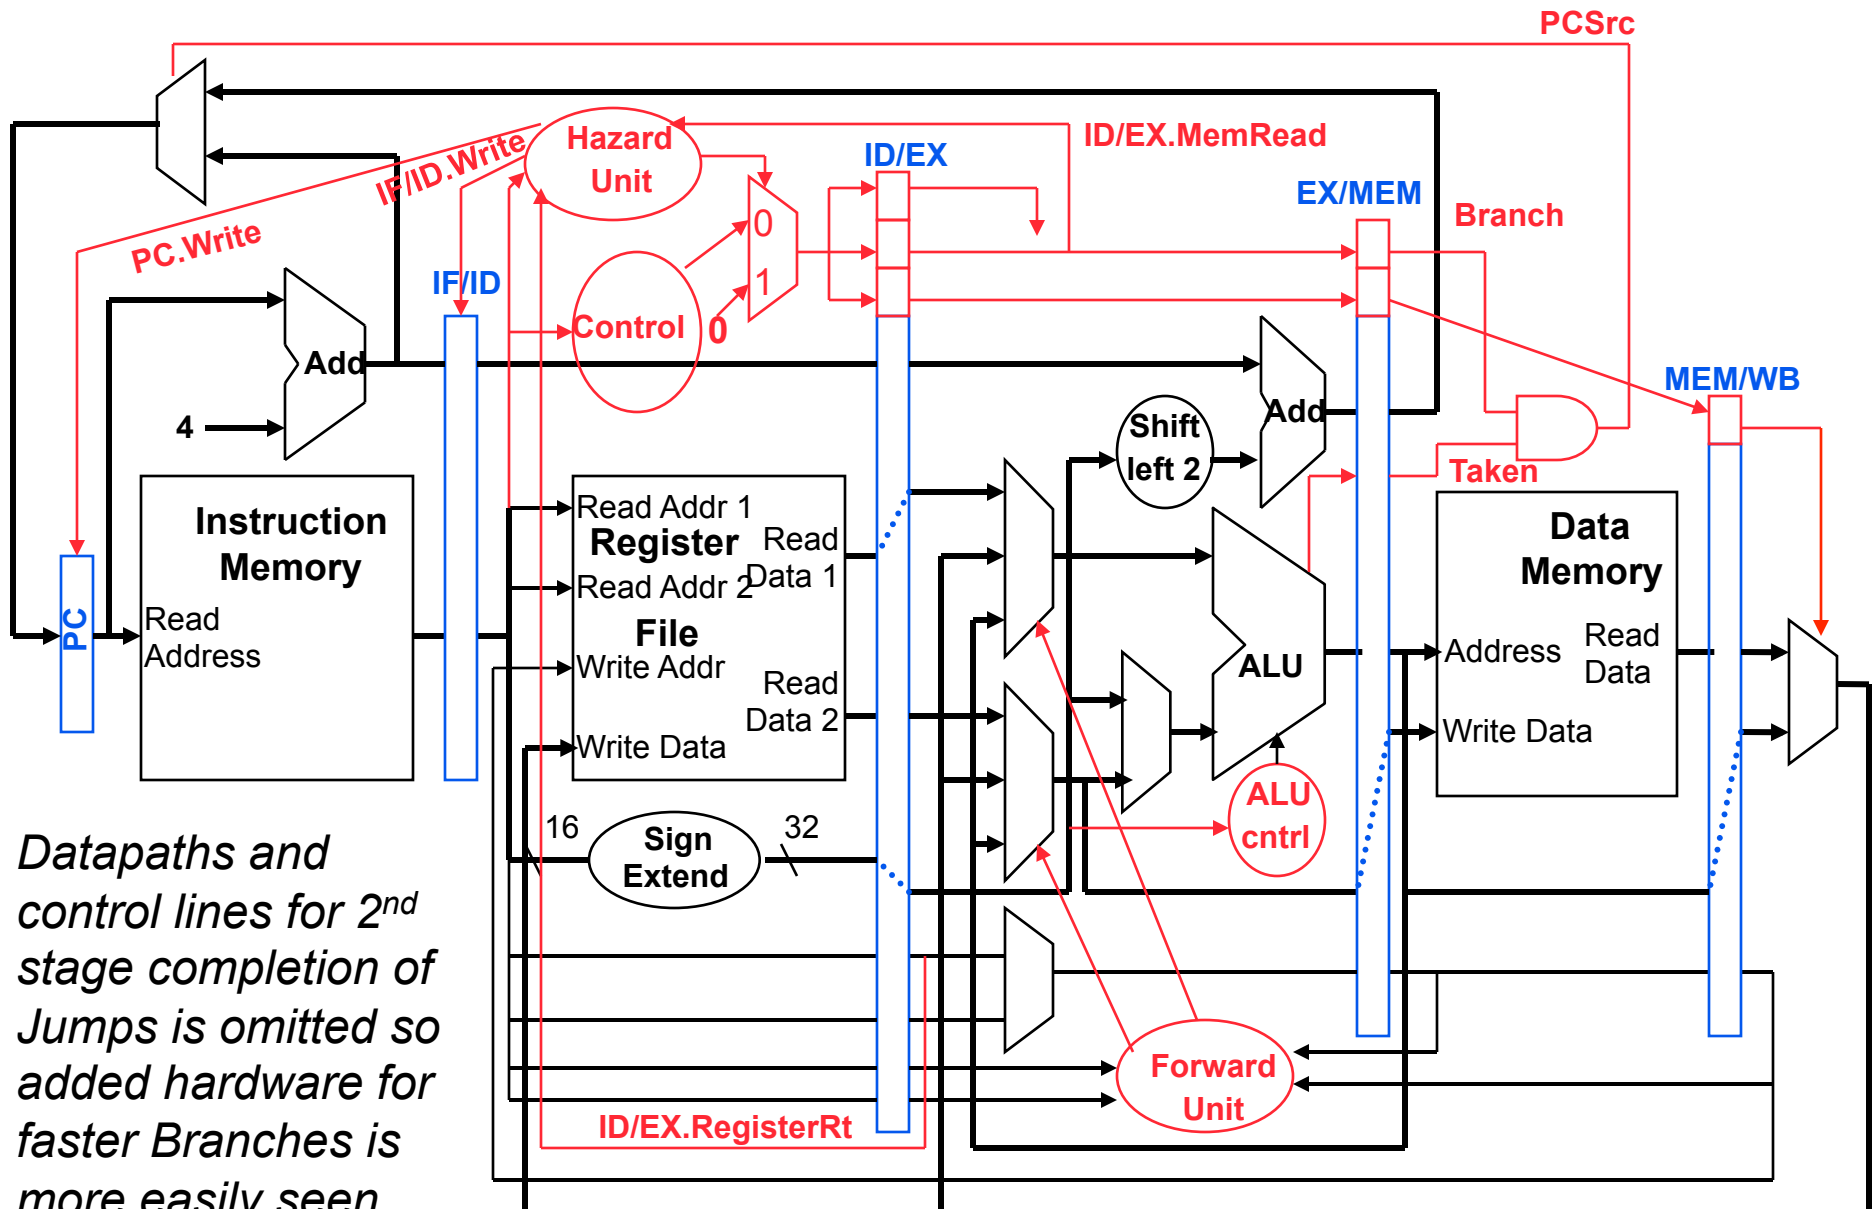
ID Branch Forwarding Issues, con't

- ❑ If the instruction immediately before the branch produces one of the branch source operands, then a **stall** must be inserted (between the `beq` and `add1`) since the EX stage ALU operation is occurring at the *same time* as the ID stage branch compare operation

| | | |
|-----|-----------------------------|-----------------------------|
| WB | <code>add3</code> | <code>\$3,</code> |
| MEM | <code>add2</code> | <code>\$4,</code> |
| EX | <code>add1</code> | <code>\$1,</code> |
| ID | <code>beq</code> | <code>\$1, \$2, Loop</code> |
| IF | <code>next_seq_instr</code> | |

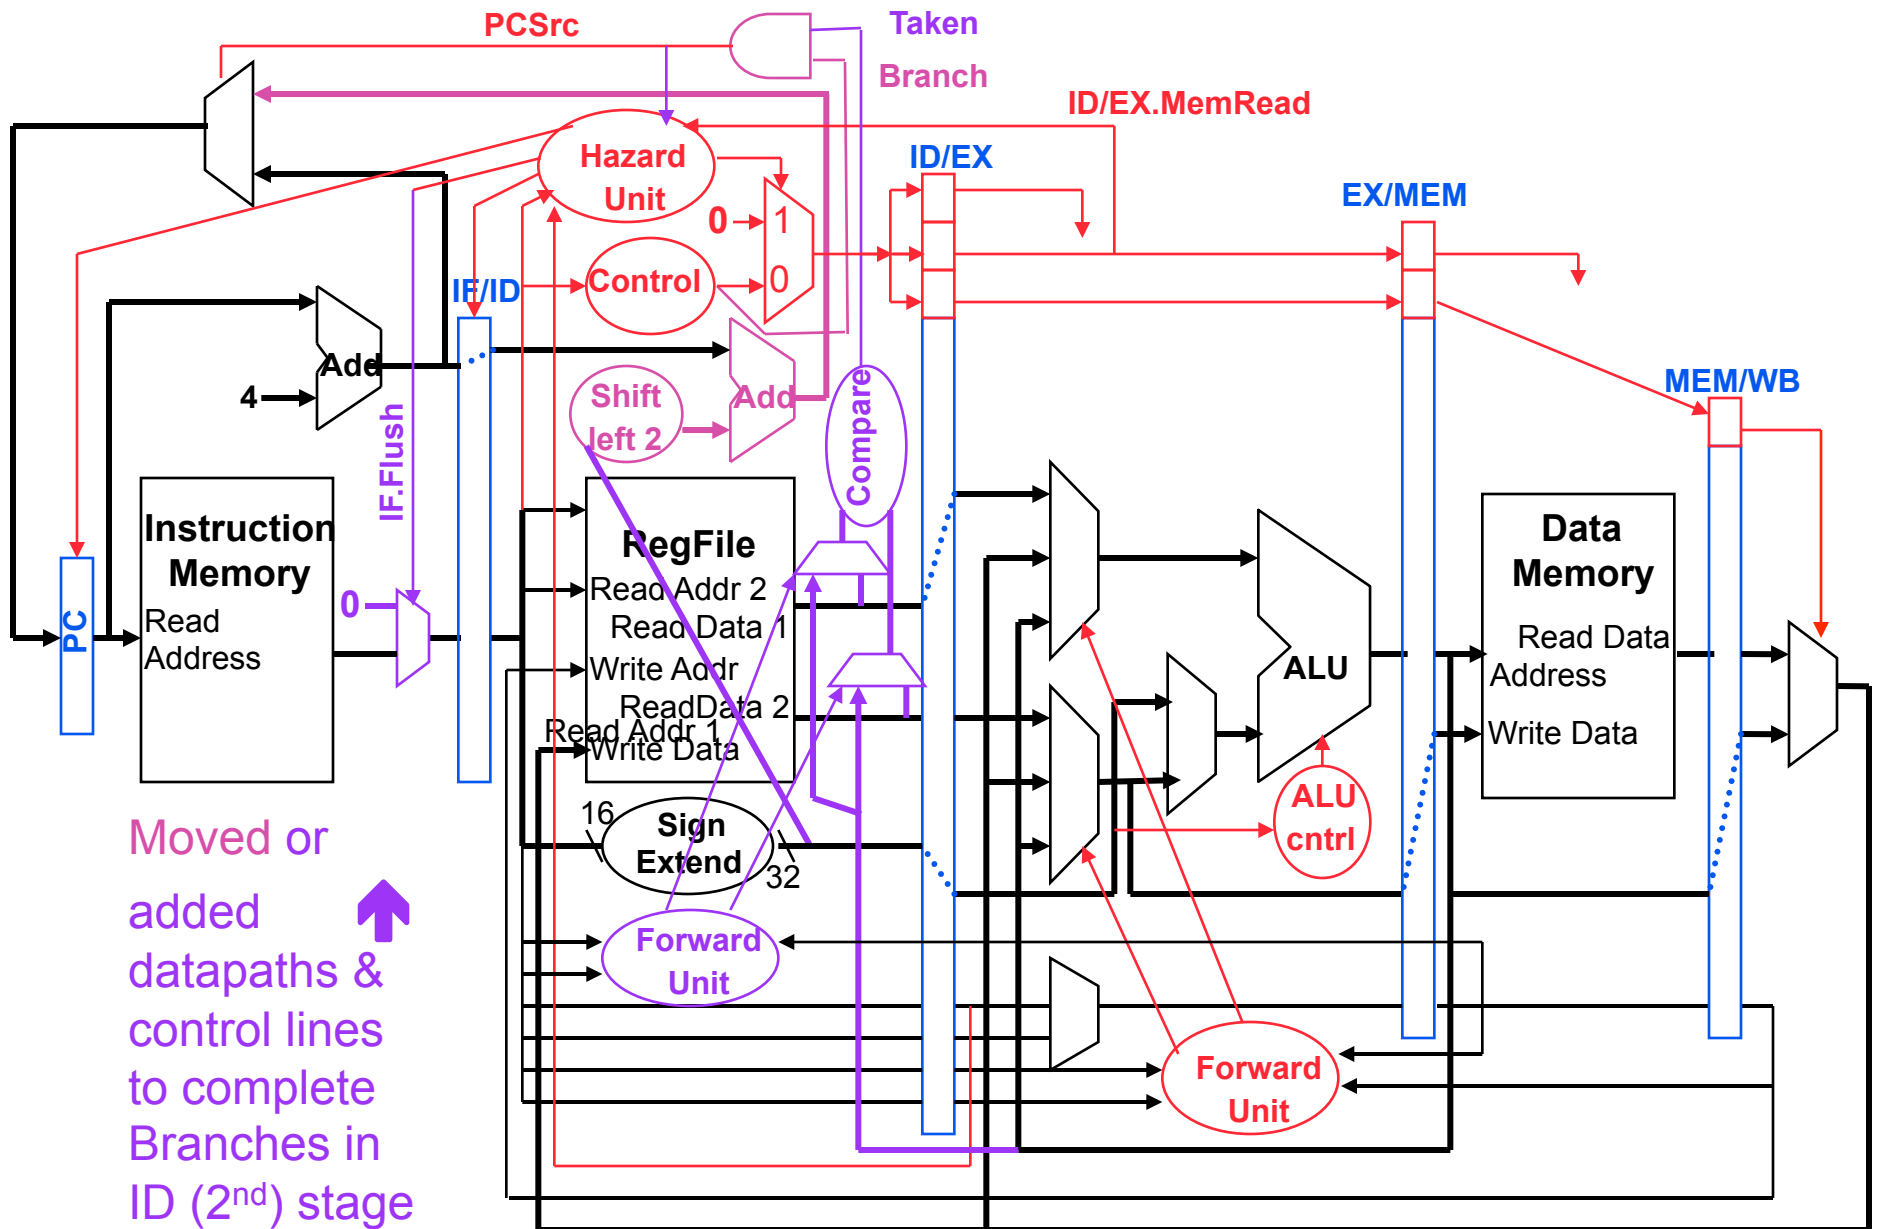
- “Bounce” the `beq` (in ID) and `next_seq_instr` (in IF) in place (ID Hazard Unit deasserts `PC.Write` and `IF/ID.Write`)
 - Insert a stall between the `add` in the EX stage and the `beq` in the ID stage by zeroing the control bits going into the ID/EX pipeline register (done by the ID Hazard Unit)
- ❑ If the branch is found to be taken, then flush the instruction currently in IF (**IF.Flush**)

MEM Stage-4 Branch Hardware with Hazard/Stall HW



Datapaths and control lines for 2nd stage completion of Jumps is omitted so added hardware for faster Branches is more easily seen.

Adding HW to Support ID 2nd Stage Branches



Moved or added datapaths & control lines to complete Branches in ID (2nd) stage

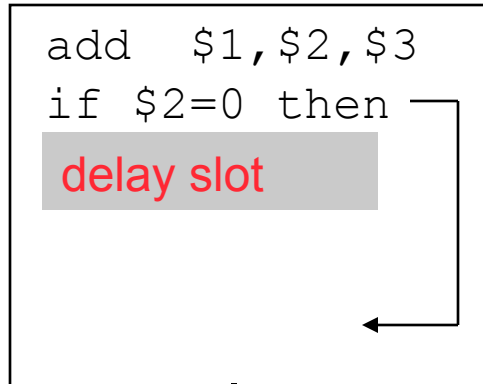
Delayed Branches

- ❑ If the branch hardware has been moved to the ID stage, then we can eliminate all branch stalls with **delayed branches**, which are defined as always executing the next sequential instruction after the branch instruction – the branch takes effect *after* that next instruction
 - The original MIPS compiler moves an instruction that is not affected by the branch immediately after the branch (a **safe** instruction) thereby **hiding** the branch delay

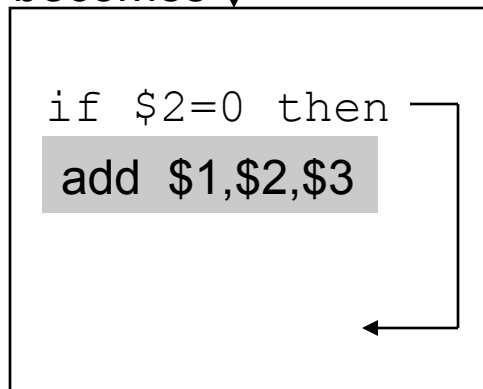
- ❑ With deeper pipelines, the branch delay grows requiring many more than one delay slot
 - Delayed branches have lost popularity compared to more expensive but more flexible (dynamic) hardware branch prediction
 - Growth in available transistors has made hardware branch prediction relatively cheaper

Scheduling Branch Delay Slots (only early MIPS)

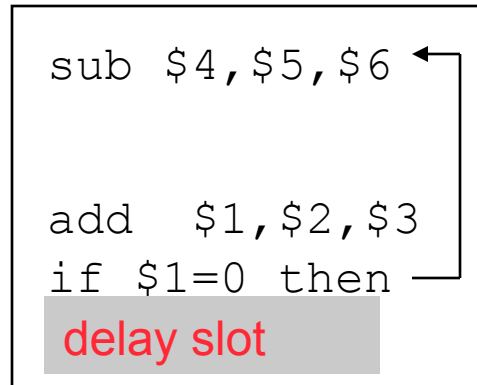
A. From before branch



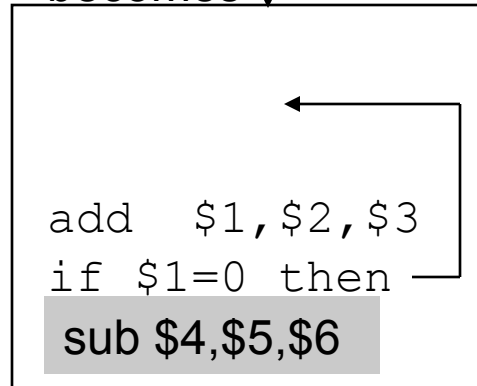
becomes



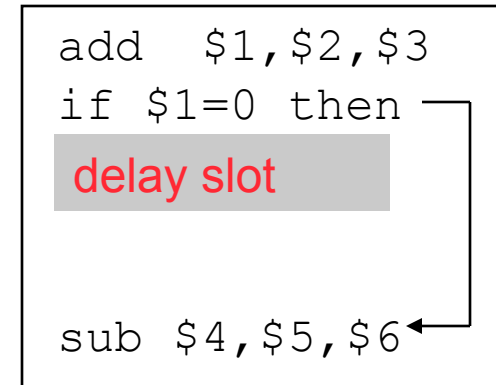
B. From branch target



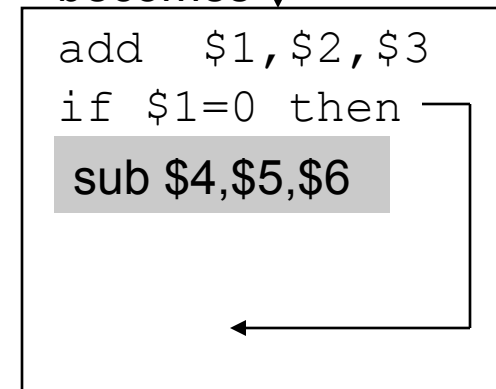
becomes



C. From fall through



becomes

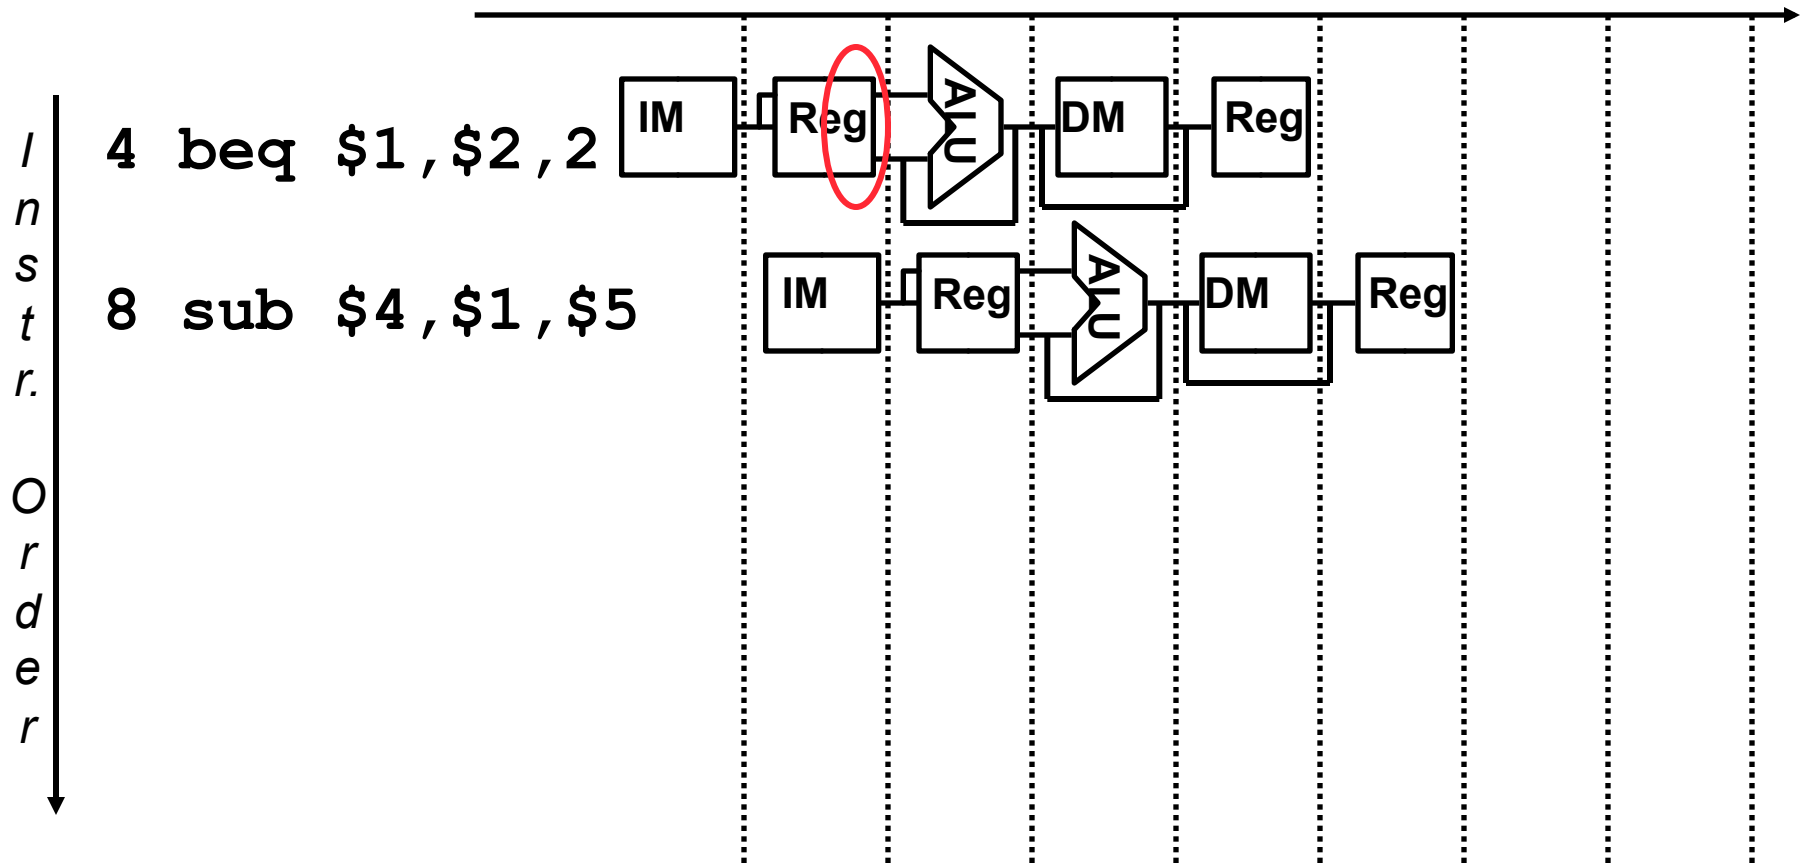


- ❑ A is the best choice, fills delay slot and reduces IC
- ❑ In B and C, the `sub` instruction may need to be copied, increasing IC
- ❑ In B and C, must be okay to execute `sub` when branch fails

Static (Compile-Time) Branch Prediction

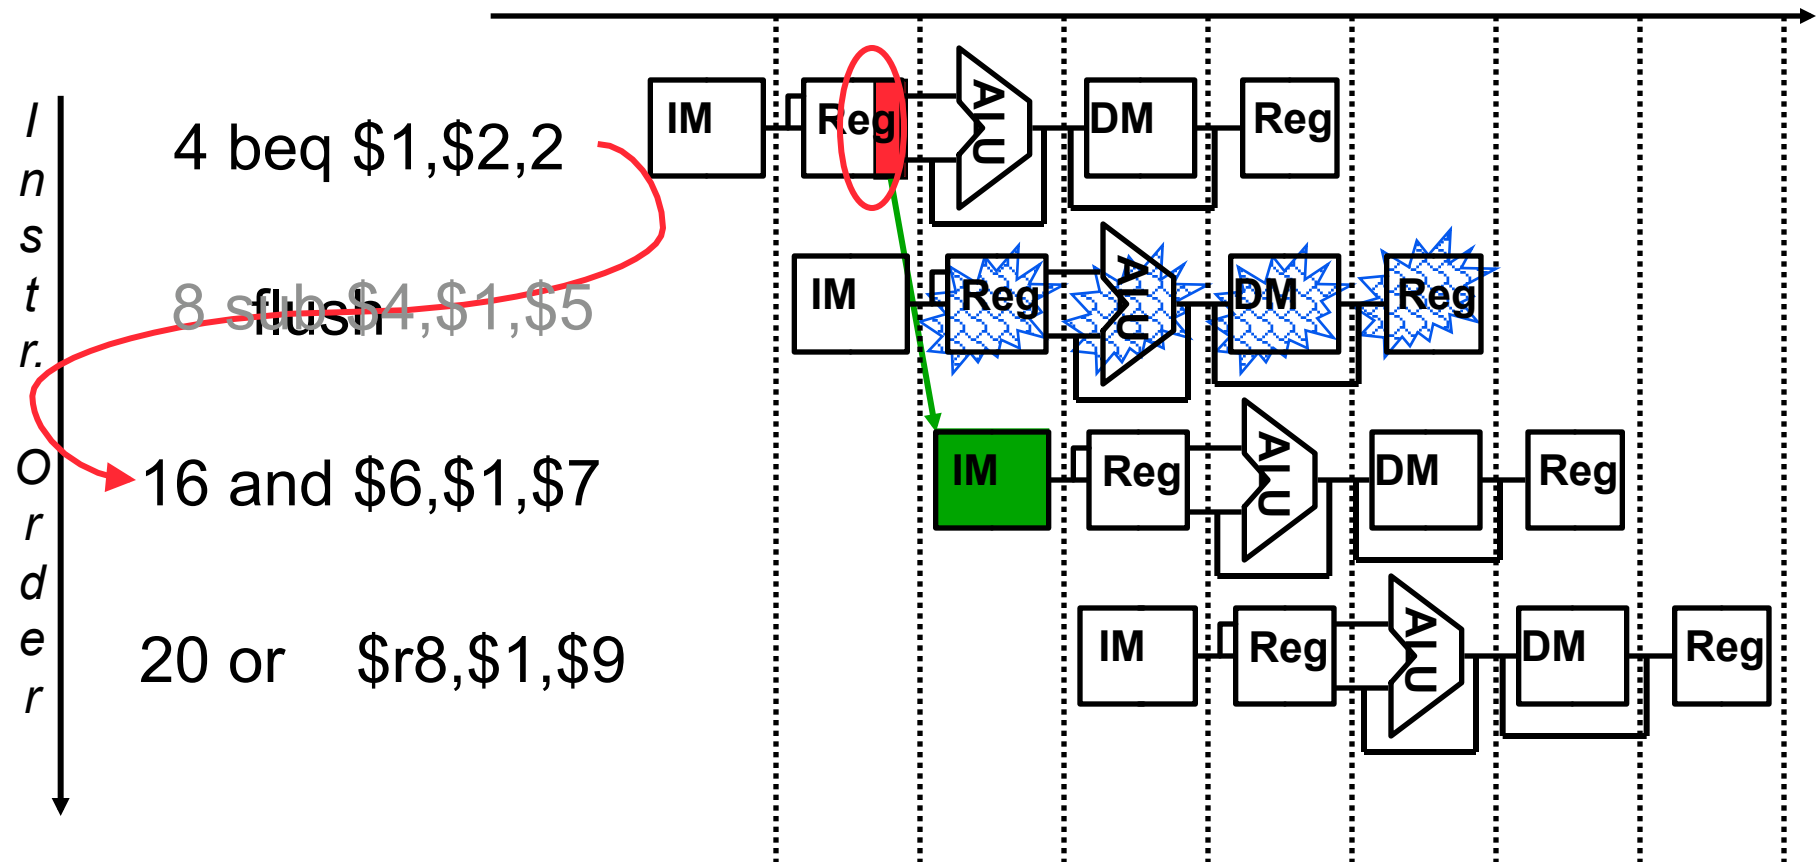
- ❑ Resolve branch hazards by assuming a given outcome and proceeding without waiting to see the actual branch outcome
- 1. **Predict not taken** – always predict branches will **not** be taken, continue to fetch from the sequential instruction stream, only when branch *is* taken does the pipeline stall
 - If taken, **flush** instructions **after** the branch (earlier in the pipeline)
 - in IF, ID, and EX stages if branch logic in MEM – **three** stalls
 - In IF and ID stages if branch logic in EX – **two** stalls
 - in IF stage if branch logic in ID – **one** stall
 - ensure that those flushed instructions have not changed the machine state – automatic in the MIPS pipeline since machine state changing operations are at the tail end of the pipeline (MemWrite (in MEM) or RegWrite (in WB))
 - restart the pipeline at the branch destination

Flushing After a Misprediction (Not Taken is wrong)



- ❑ To flush the IF stage instruction, assert `IF.Flush` to zero the instruction field of the IF/ID pipeline register (transforming it into a `noop`)

Flushing with Misprediction (Not Taken was wrong)

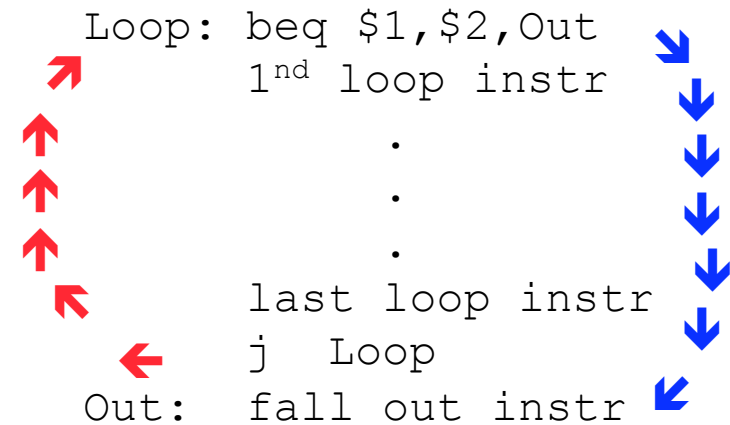


- ❑ To flush the IF stage instruction (sub \$4,\$1,\$5), assert `IF.Flush` (see next slide) to zero the instruction field of the IF/ID pipeline register (transforming it into a `noop`)

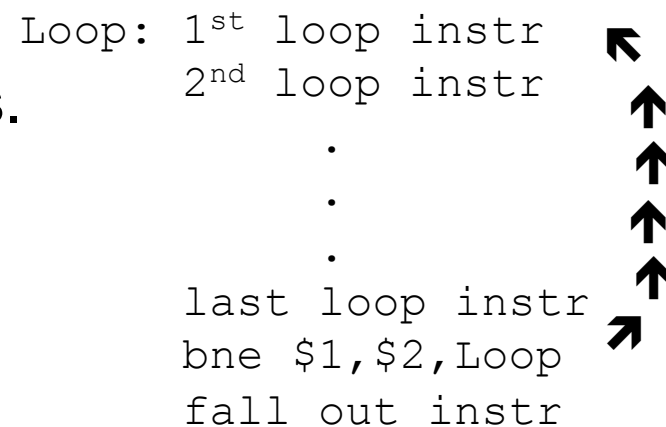
Branching Structures

- ❑ Predicting not-taken works well for “top of the loop” branching structures

- But such loops have jumps at the bottom of the loop to return to the top of the loop – and incur the jump stall overhead



- ❑ Predict not taken does not work well for “bottom of the loop” branching structures since most loops last about 10 iterations.



Static Branch Prediction, continued

- ❑ Resolve branch hazards by assuming a given outcome and proceeding
- 2. **Predict taken** – predict branches will always be taken
 - Predict taken *always* incurs at least one stall cycle (just one, if branch destination hardware has been moved to the ID stage)
 - Is there a way to “cache” the address of the branch target instruction ??
- ❑ As the branch penalty increases (for deeper pipelines), a simple static prediction scheme will hurt performance. With more hardware, it is possible to try to predict branch behavior **dynamically** during program execution
- 3. **Dynamic branch prediction** – predict branches at run-time using *run-time* information

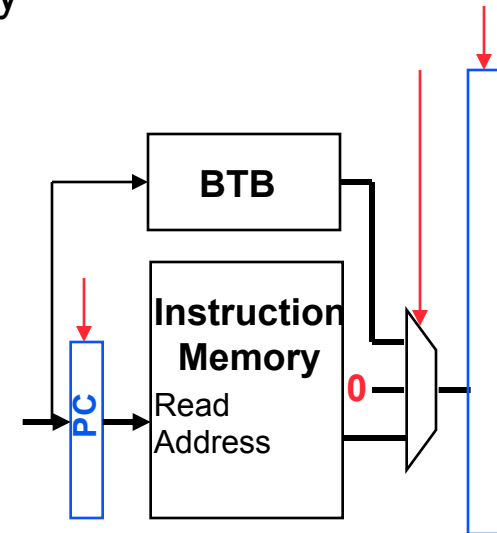
Dynamic (Run-Time) Branch Prediction

- A **branch prediction buffer** (aka branch history table (**BHT**)) in the IF stage addressed by the lower bits of the PC, contains bit(s) passed to the ID stage through the IF/ID pipeline register that tells whether the branch was taken the last time it was executed
 - Prediction bit may predict incorrectly (may be a wrong prediction for this branch for this iteration or may be for a different branch with the same low order PC bits) but a bad prediction does not affect **correctness**, just **performance**
 - Branch decision occurs in the ID stage after determining that the fetched instruction is a branch and checking the prediction bit(s)
 - If the prediction is wrong, flush the incorrect instruction(s) in pipeline, restart the pipeline with the right instruction, and invert the prediction bit(s)
 - A 4096 entry BHT with a 1-bit prediction entry varies from 1% misprediction (nasa7, tomcatv) to 18% (eqntott)

Branch Target Buffer

- ❑ The BHT predicts *when* a branch is taken, but does not tell *where* it is taken to!
 - A **branch target buffer (BTB)** in the IF stage can cache the branch target address, but the CPU also needs to fetch the next sequential instruction. The prediction bit in IF/ID selects which “next” instruction will be loaded into IF/ID at the next clock edge
 - Needs two read-ports for instruction memory

- Or the BTB can cache the branch-taken **instruction** while the instruction memory is fetching the next sequential instruction



- ❑ If the prediction is correct, a stall can be avoided no matter which direction the branch goes

1-bit Prediction Accuracy

- ❑ A 1-bit predictor will be incorrect twice when not taken

- Assume `predict_bit = 0` to start (indicating branch not taken) and loop control is at the bottom of the loop code

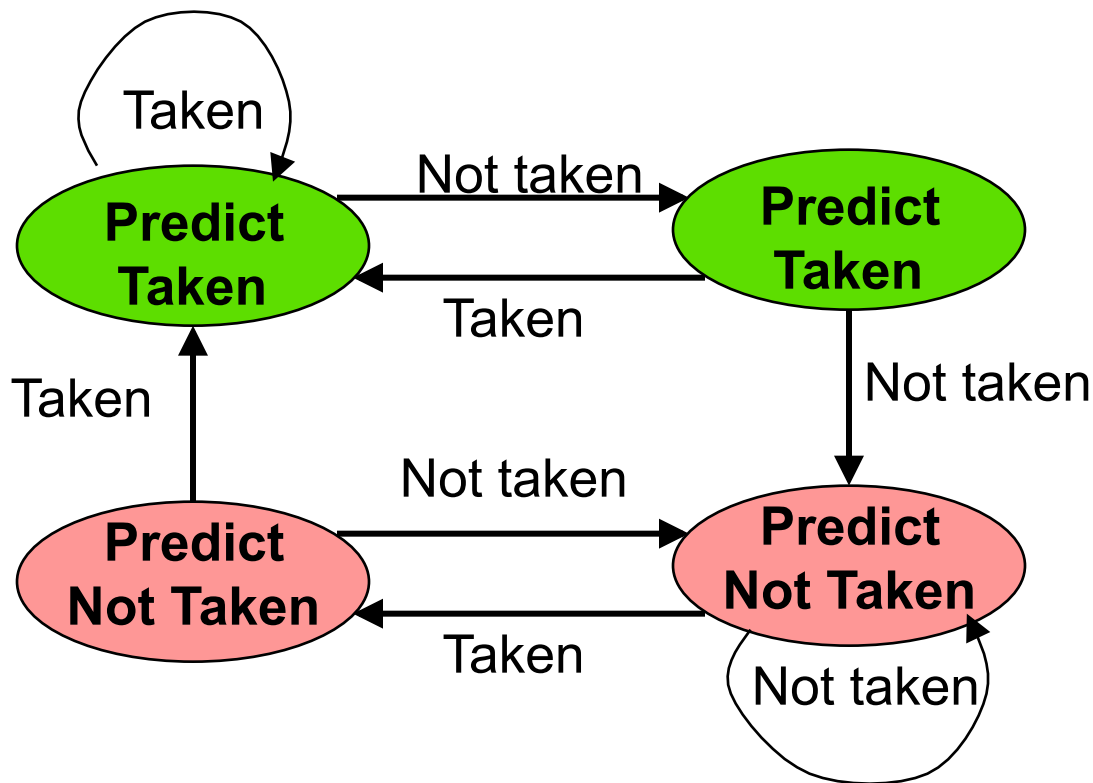
1. The first time through the loop, the predictor mispredicts the branch since the branch is taken back to the top of the loop; invert prediction bit (`predict_bit = 1`)
2. As long as branch is taken (looping), the prediction will be correct
3. Exiting the loop, the predictor again mispredicts the branch since this time the branch is not taken, falling out of the loop; invert prediction bit (`predict_bit = 0`)

```
Loop: 1st loop instr
      2nd loop instr
      .
      .
      .
      last loop instr
      bne $1,$2,Loop
      fall out instr
```

- ❑ For 10 times through the loop we have a 80% prediction accuracy for a branch that is taken 90% of the time

2-Bit Branch Direction Predictors

- A 2-bit scheme can give 90% accuracy by waiting until a prediction is wrong twice in a row before the prediction bit is changed

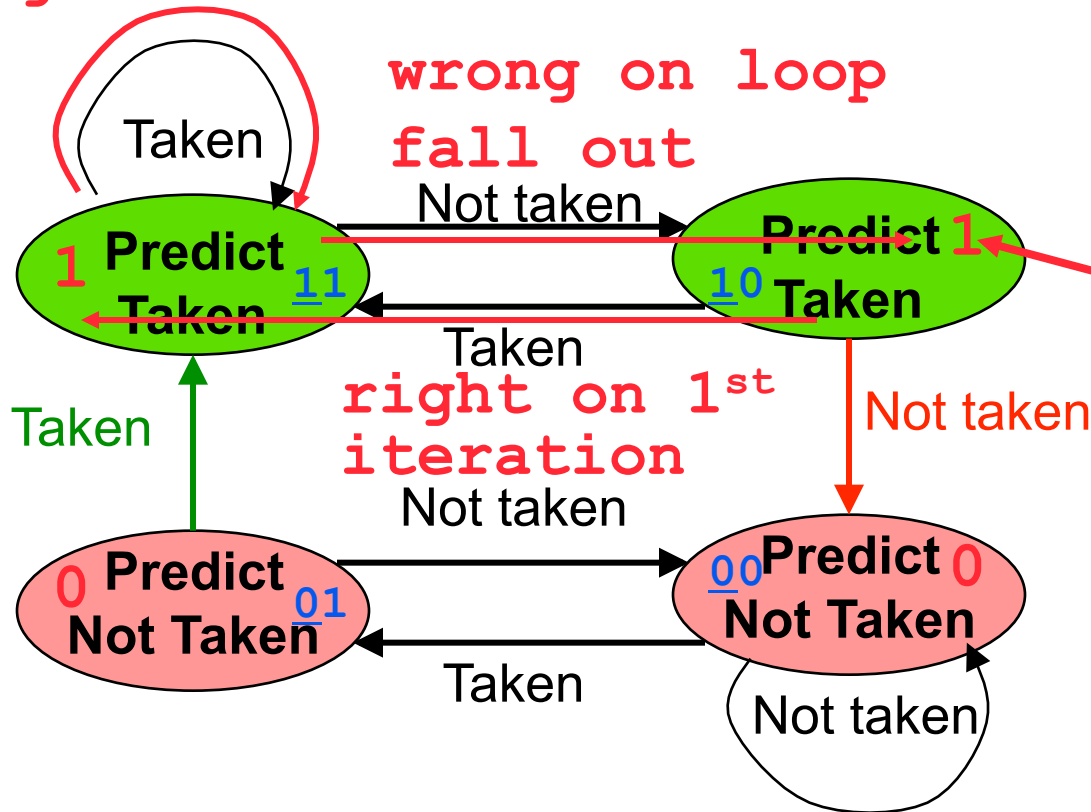


```
Loop: 1st loop instr
      2nd loop instr
      .
      .
      .
      last loop instr
      bne $1,$2,Loop
      fall out instr
```

2-Bit Branch Predictor (Branch History Table)

- A 2-bit scheme can give 90% accuracy by waiting until a prediction is wrong twice in a row before the prediction bit is changed

right 9 times



Loop: 1st loop instr
 2nd loop instr
 .
 .
 .
 last loop instr
 bne \$1,\$2,Loop
 fall out instr

- BHT also stores the initial state of the FSM (finite state machine)

Finite State Machine for 2-bit branch predictor

Dealing with Rare Control Hazards - Exceptions

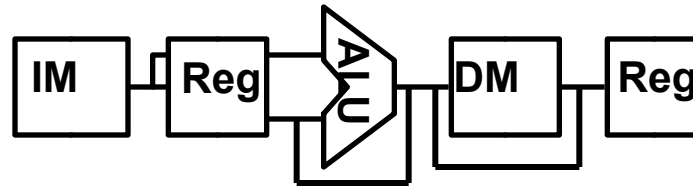
- ❑ Exceptions (aka interrupts) are just another form of control hazard. Exceptions arise from
 - R-type arithmetic overflow
 - Trying to execute an undefined instruction
 - An I/O device request
 - An OS service request (e.g., a page fault, TLB exception)
 - A hardware malfunction
- ❑ The pipeline has to stop executing the offending instruction in midstream, let all prior instructions complete, flush all following instructions, set a register to show the cause of the exception, save the address of the offending instruction, and then jump to a prearranged address (the address of the exception handler code)
- ❑ The software (OS) looks at the cause of the exception and “deals” with it

Two Types of Exceptions (Interrupts and Traps)

- ❑ Interrupts – asynchronous to program execution
 - caused by **external events**
 - may be handled **between** instructions, so can let the instructions currently active in the pipeline *complete* before passing control to the OS interrupt handler
 - simply suspend and resume user program

- ❑ Traps (A.K.A. Exceptions) – synchronous to program execution
 - caused by **internal events**
 - condition must be remedied by the trap handler for **that** instruction, so much stop the offending instruction *midstream* in the pipeline and pass control to the OS trap handler
 - the offending instruction may be retried (or simulated by the OS) and the program may continue or it may be aborted

Where in the Pipeline Exceptions Occur

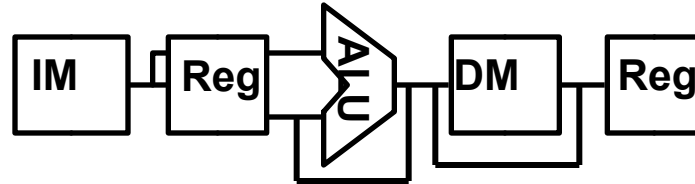


Stage(s)?

Synchronous?

- Arithmetic overflow
- Undefined instruction
- TLB or page fault
- I/O service request
- Hardware malfunction

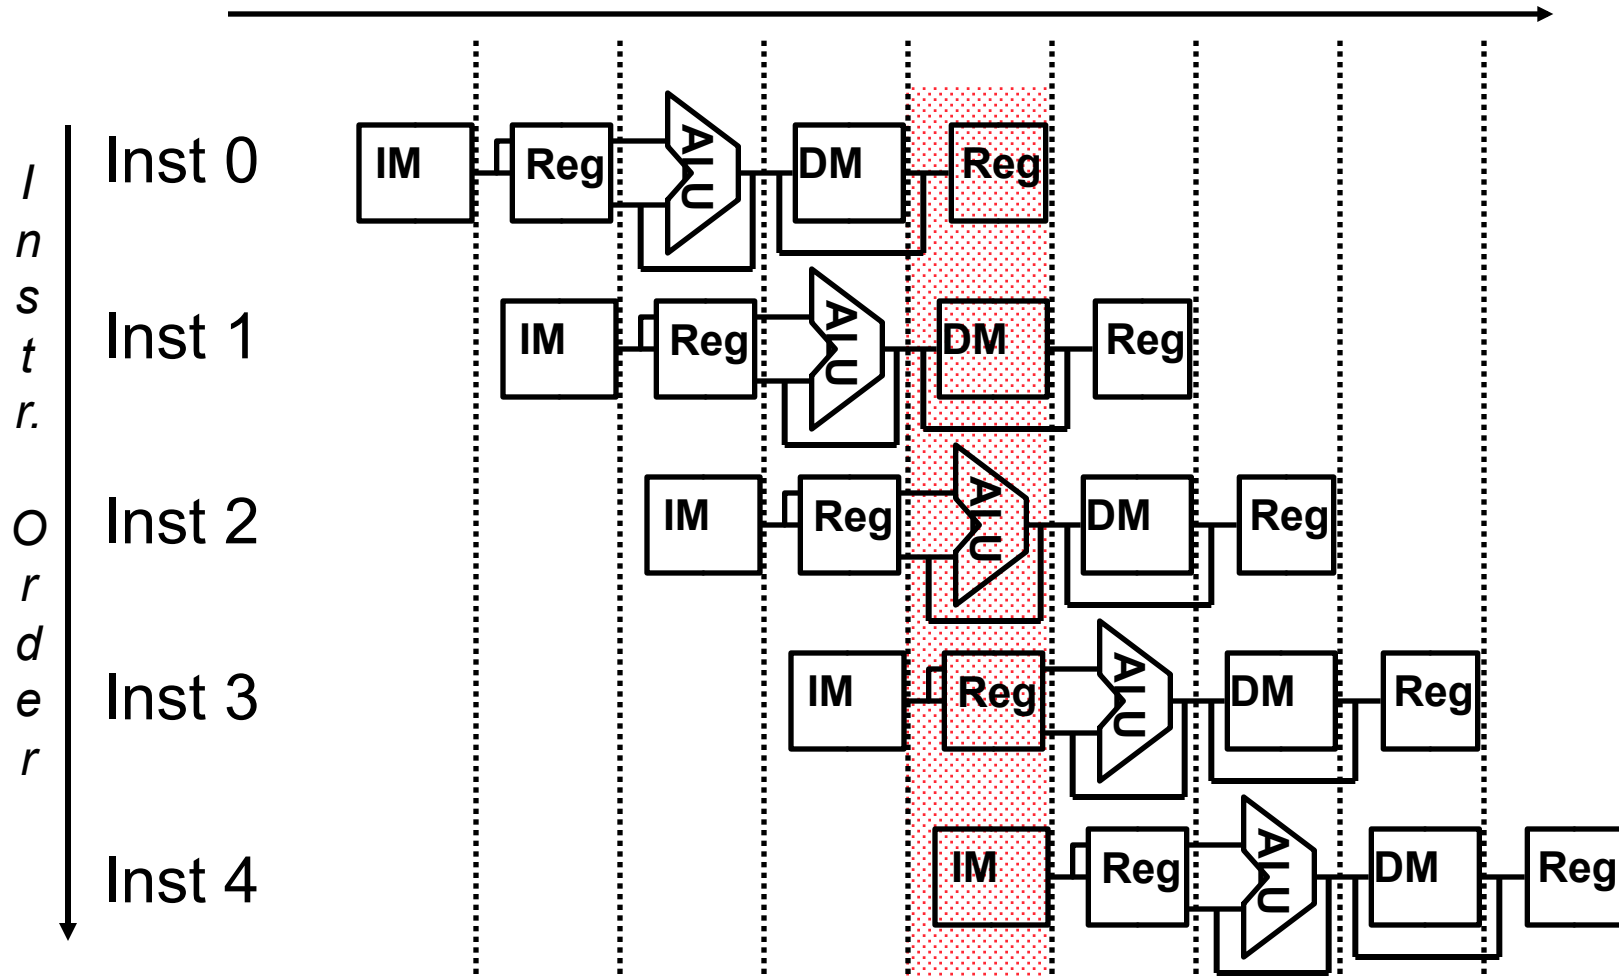
Where in the Pipeline Exceptions Occur



| | Stage(s)? | Synchronous? |
|--|-----------|--------------|
| <input type="checkbox"/> Arithmetic overflow | EX | yes |
| <input type="checkbox"/> Undefined instruction | ID | yes |
| <input type="checkbox"/> TLB or page fault | IF, MEM | yes |
| <input type="checkbox"/> I/O service request | any | no |
| <input type="checkbox"/> Hardware malfunction | any | no |

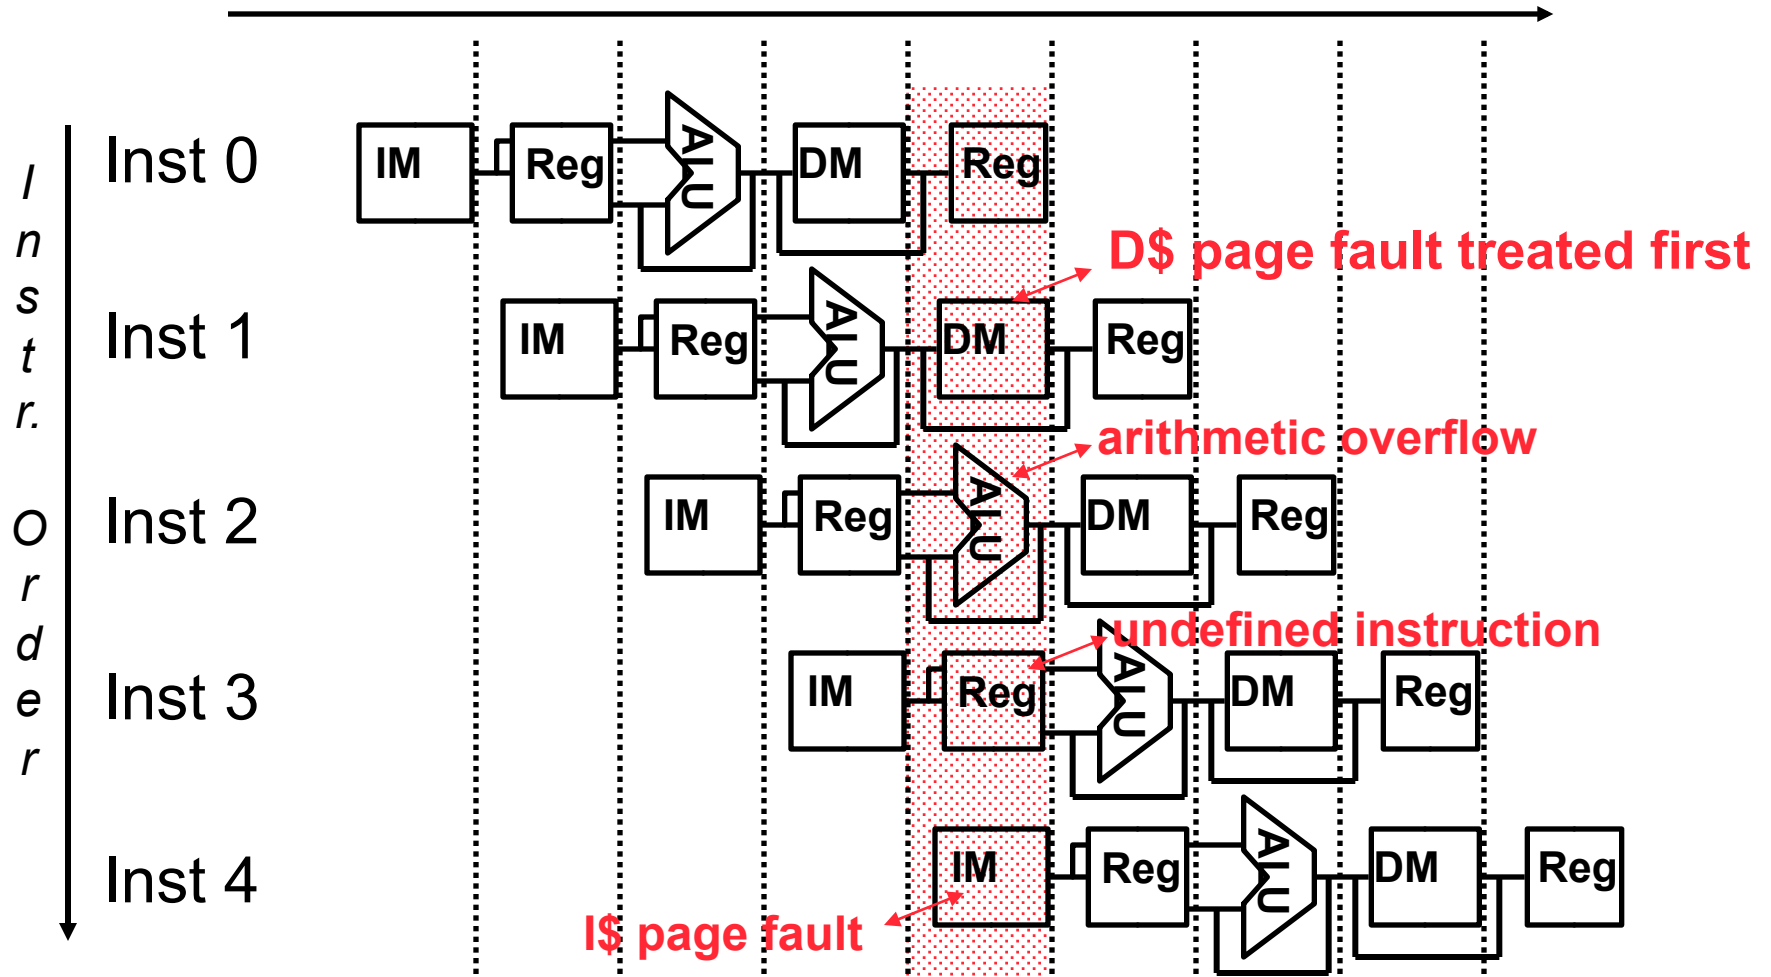
Beware that multiple exceptions can occur simultaneously in a *single* clock cycle

Multiple Simultaneous Exceptions



- ❑ Hardware sorts the exceptions so that the earliest instruction is the one interrupted first

Multiple Simultaneous Exceptions

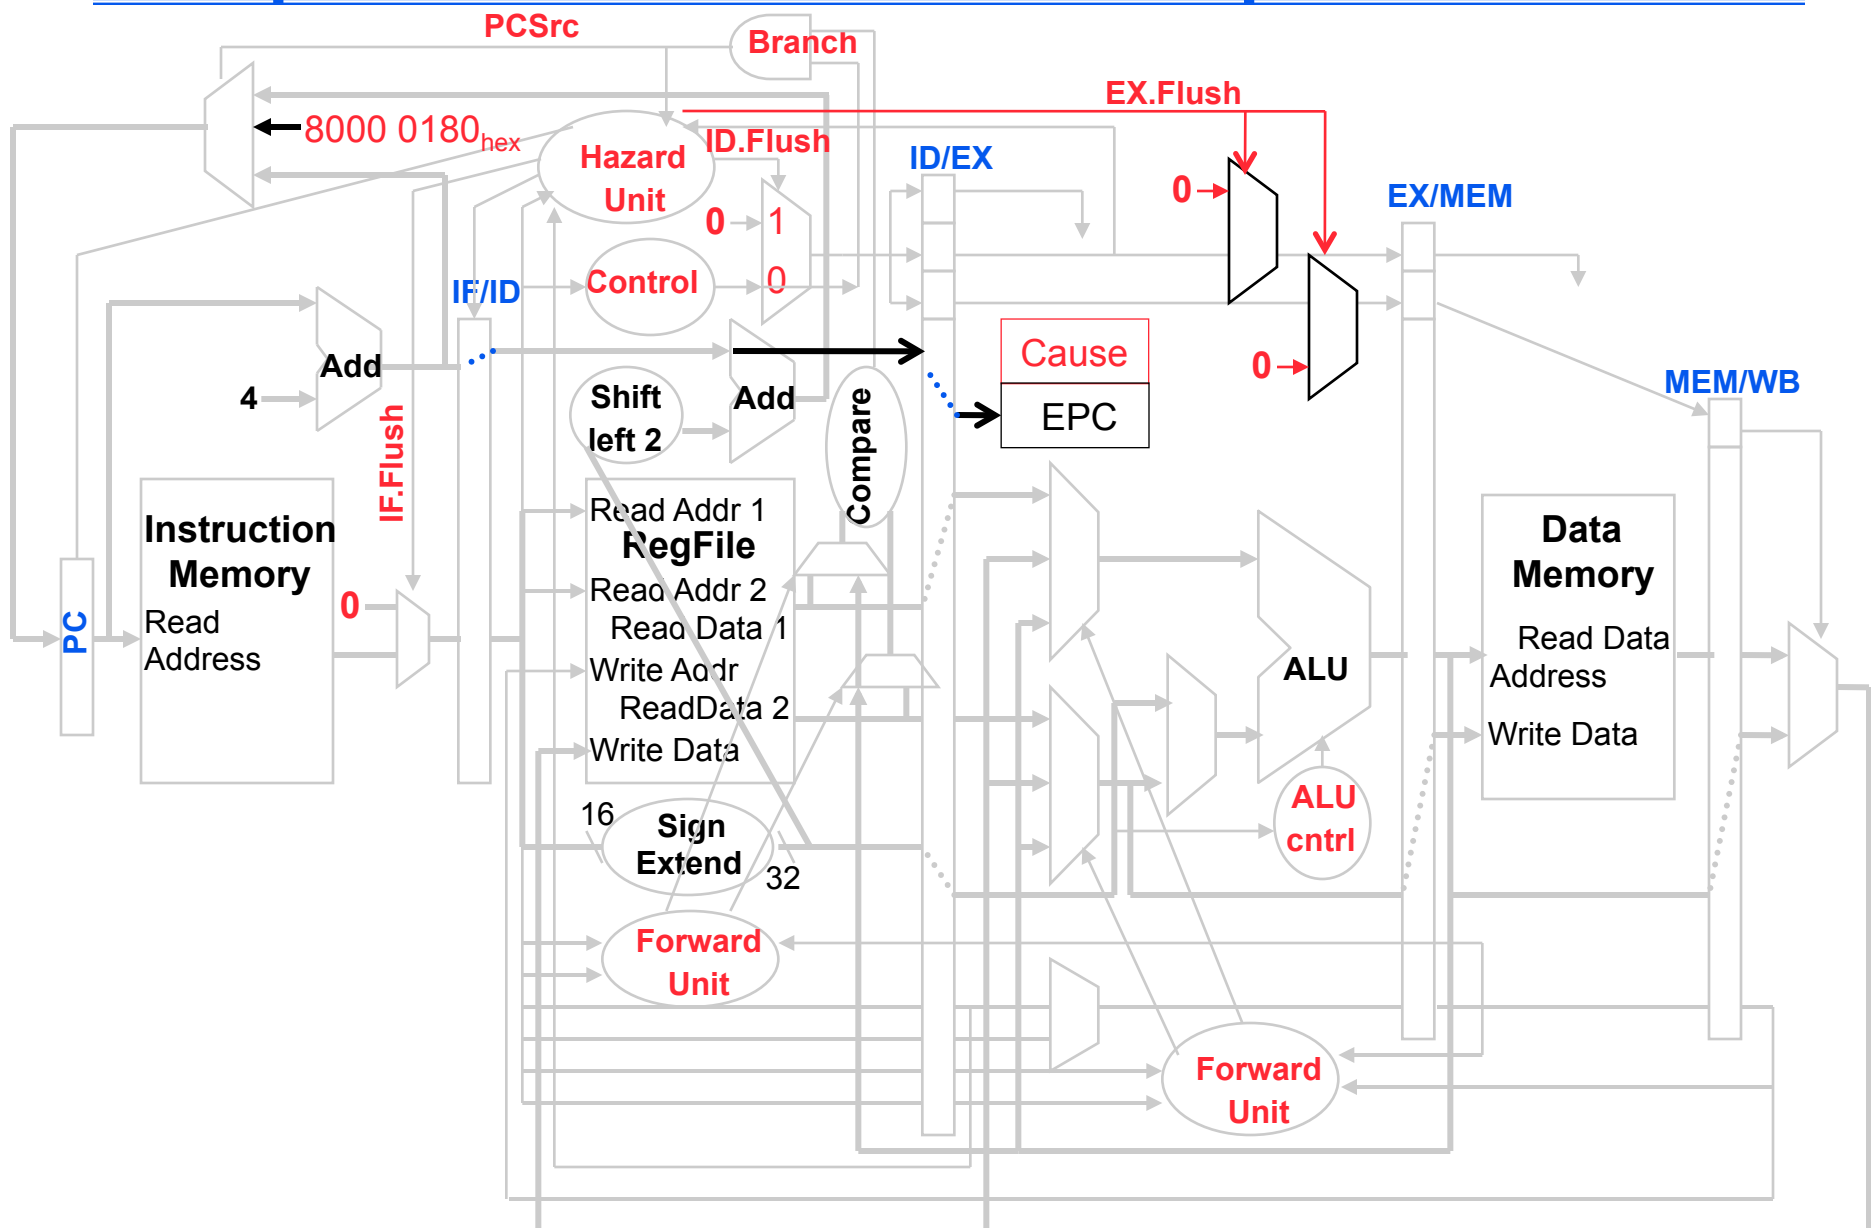


- ❑ Hardware sorts the exceptions so that the earliest instruction (Inst 1) is the one interrupted first

Additions to MIPS to Handle Exceptions (Fig 6.42)

- ❑ *Cause* register (records exceptions) – hardware to record in *Cause* the exceptions and a signal to control writes to it (*CauseWrite*)
- ❑ *EPC* register (records the addresses of the offending instructions) – hardware to record in *EPC* the address of the offending instruction and a signal to control writes to it (*EPCWrite*) (*e*xception *p*rogram *c*ounter)
 - Exception software must match the exception to the instruction
- ❑ A way to load the PC with the address of the exception handler
 - Expand the PC input mux where the new input is hardwired to the exception handler address - (e.g., 8000 0180_{hex} for arithmetic overflow)
- ❑ A way to flush offending instruction and the ones that follow it

Datapath with Controls for Exceptions



Summary

- ❑ All modern day processors use pipelining for performance (a CPI of 1 and a fast Clock Rate)
- ❑ Pipeline clock rate limited by **slowest** pipeline stage – so designing a balanced pipeline is important
- ❑ Must detect and resolve hazards
 - Structural hazards – resolved by designing the pipeline correctly
 - Data hazards
 - Stall (impacts CPI)
 - Forward (requires hardware support)
 - Control hazards – put the branch decision hardware in as early a stage in the pipeline as possible
 - Stall (impacts CPI)
 - Delay decision (requires compiler support)
 - Static and **dynamic prediction** (requires hardware support)
- ❑ Pipelining complicates exception handling

Next Lecture and Reminders

□ Next lecture

- Multiple issue processors
 - Reading assignment – PH, rest of Chapter 4, mainly 4.10

□ Reminders

- HW2 due Thursday, 10/29/09, in class or in 1308.
- HW3 will come out early in November.
- Midterm exam will be in class early in November, probably Thursday, 11/5/09.