

18-600 Foundations of Computer Systems

Lecture 9: "Pipelined Processor Design"

John P. Shen & Zhiyi Yu
September 28, 2016

18-600

CS: AAP

CS: APP

➤ Required Reading Assignment:

- Chapter 4 of CS:APP (3rd edition) by Randy Bryant & Dave O'Hallaron.

➤ Recommended Reference:

- ❖ Chapters 1 and 2 of Shen and Lipasti (SnL).



18-600 Foundations of Computer Systems

Lecture 9: "Pipelined Processor Design"

1. TYPICAL Pipelined Processor

- a. Instruction Pipeline Design
 - Balancing Pipeline Stages
 - Unifying Instruction Types
 - Resolving Pipeline Hazards

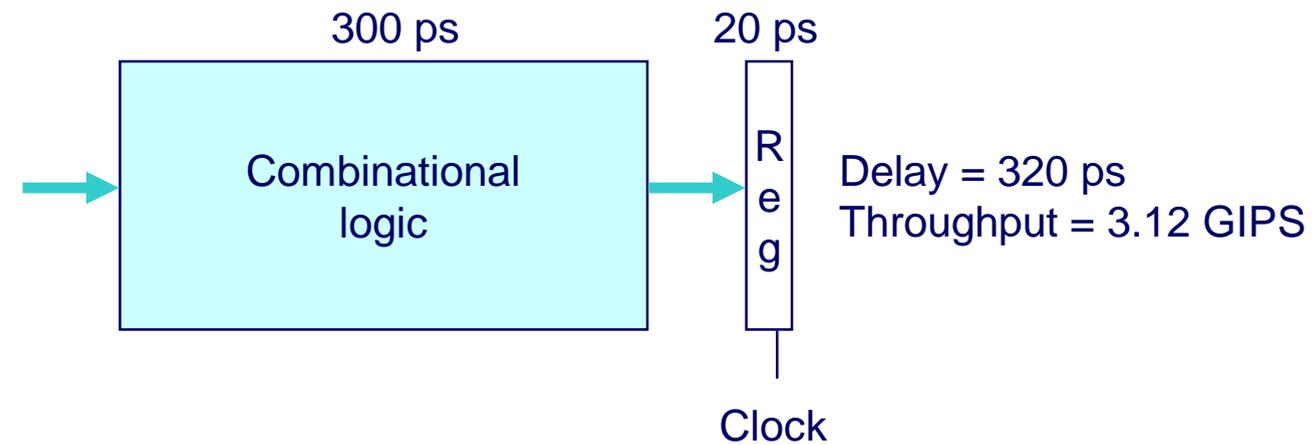
2. Y86-64 Pipelined Processor (PIPE)

- a. Pipelining of the SEQ Processor
- b. Dealing with Data Hazards
- c. Dealing with Control Hazards

3. Motivation for Superscalar



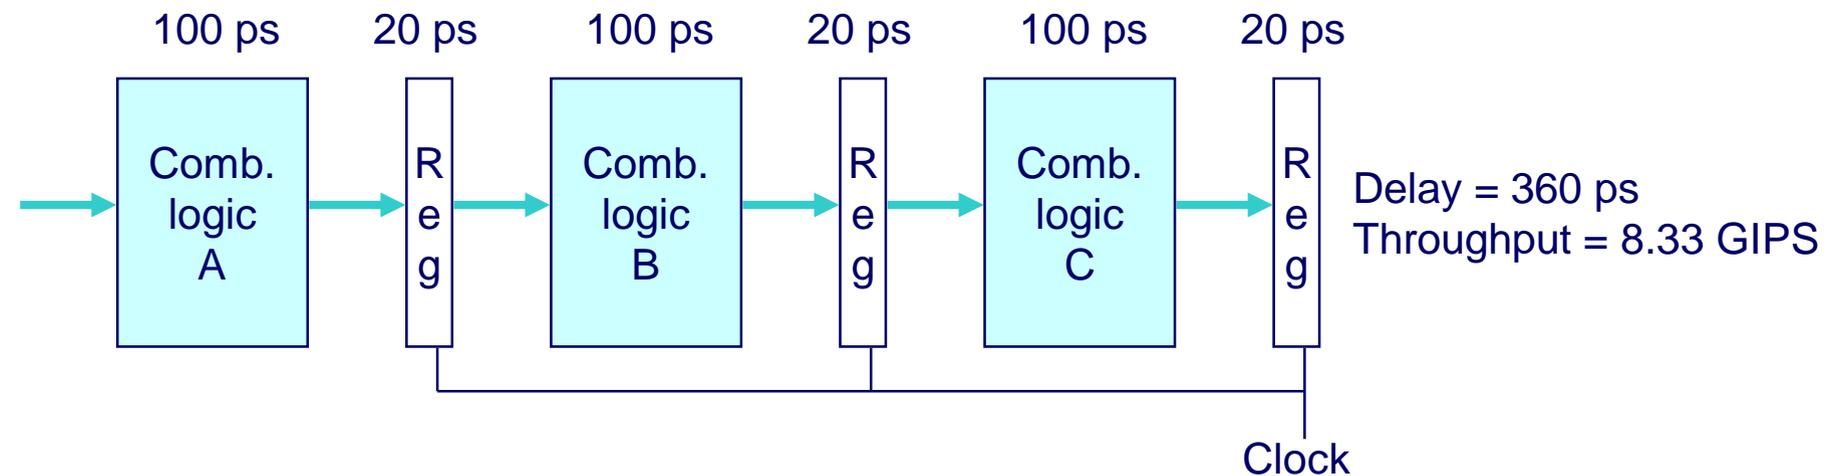
Computational Example



➤ System

- Computation requires total of 300 picoseconds
- Additional 20 picoseconds to save result in register
- Must have clock cycle of at least 320 ps

3-Way Pipelined Version

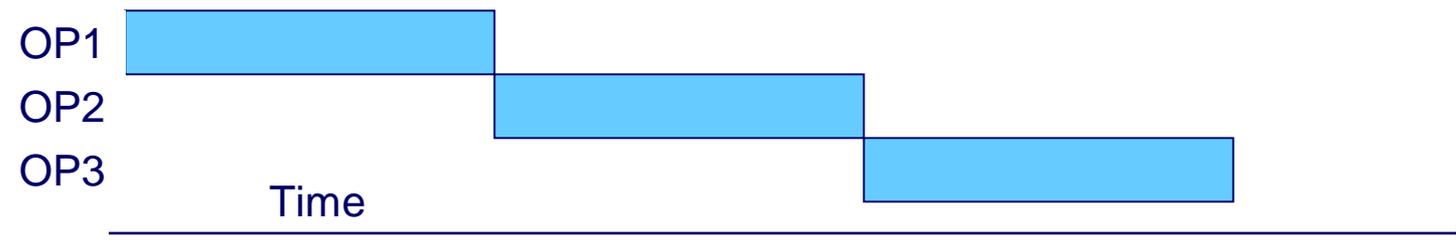


➤ System

- Divide combinational logic into 3 blocks of 100 ps each
- Can begin new operation as soon as previous one passes through stage A.
 - Begin new operation every 120 ps
- Overall latency increases
 - 360 ps from start to finish

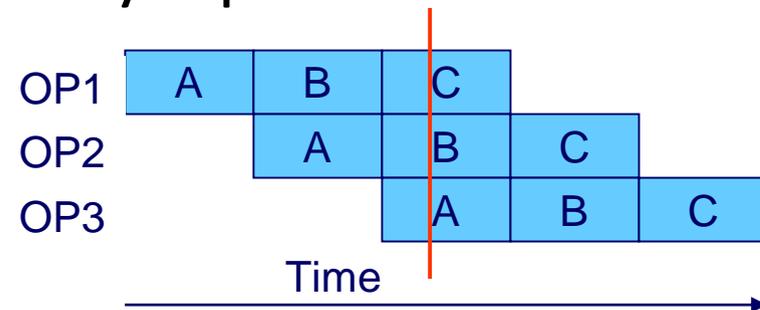
Pipeline Diagrams

➤ Unpipelined



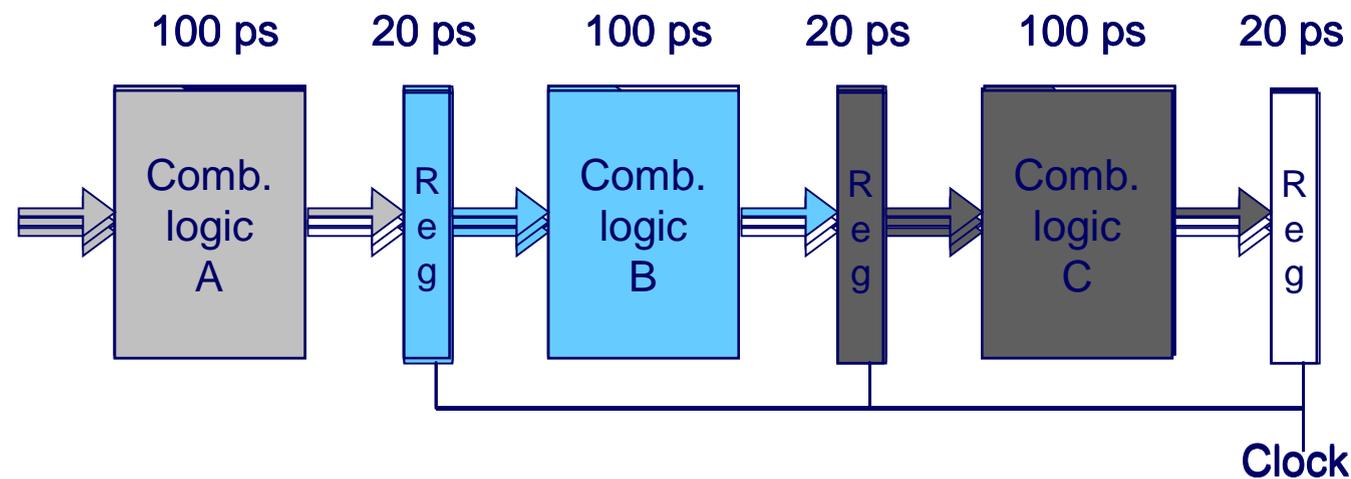
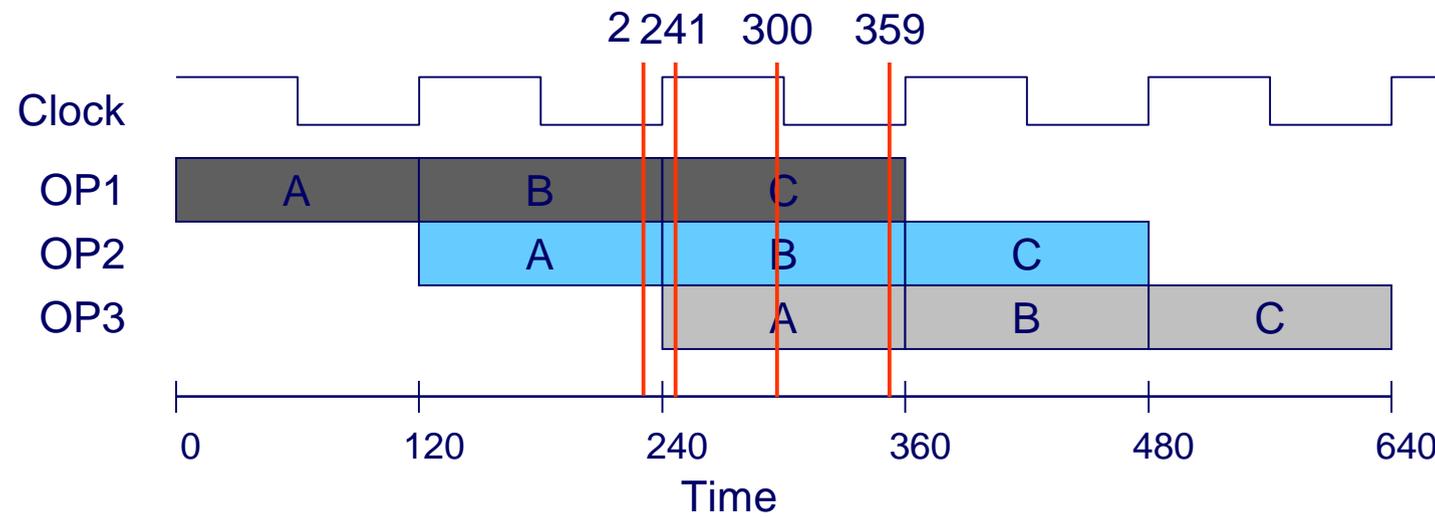
- Cannot start new operation until previous one completes

➤ 3-Way Pipelined



- Up to 3 operations in process simultaneously

Operating a Pipeline



Pipelining Fundamentals

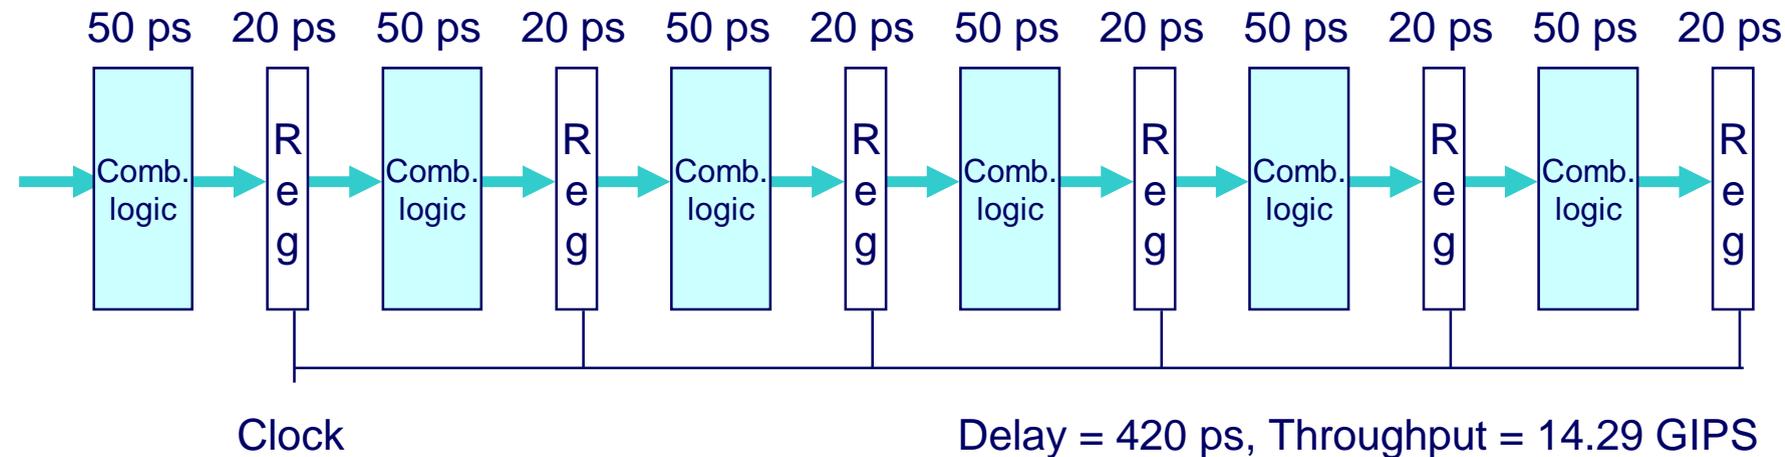
➤ Motivation:

- Increase throughput with little increase in hardware.

Bandwidth or Throughput = Performance

- Bandwidth (BW) = no. of tasks/unit time
- For a system that operates on one task at a time:
 - $BW = 1/\text{delay (latency)}$
- BW can be increased by pipelining if many operands exist which need the same operation, i.e. many repetitions of the same task are to be performed.
- Latency required for each task remains the same or may even increase slightly.

Limitations: Register Overhead



- As we try to deepen pipeline, overhead of loading registers becomes more significant
- Percentage of clock cycle spent loading register:
 - 1-stage pipeline: 6.25%
 - 3-stage pipeline: 16.67%
 - 6-stage pipeline: 28.57%
- High speeds of modern processor designs obtained through very deep pipelining

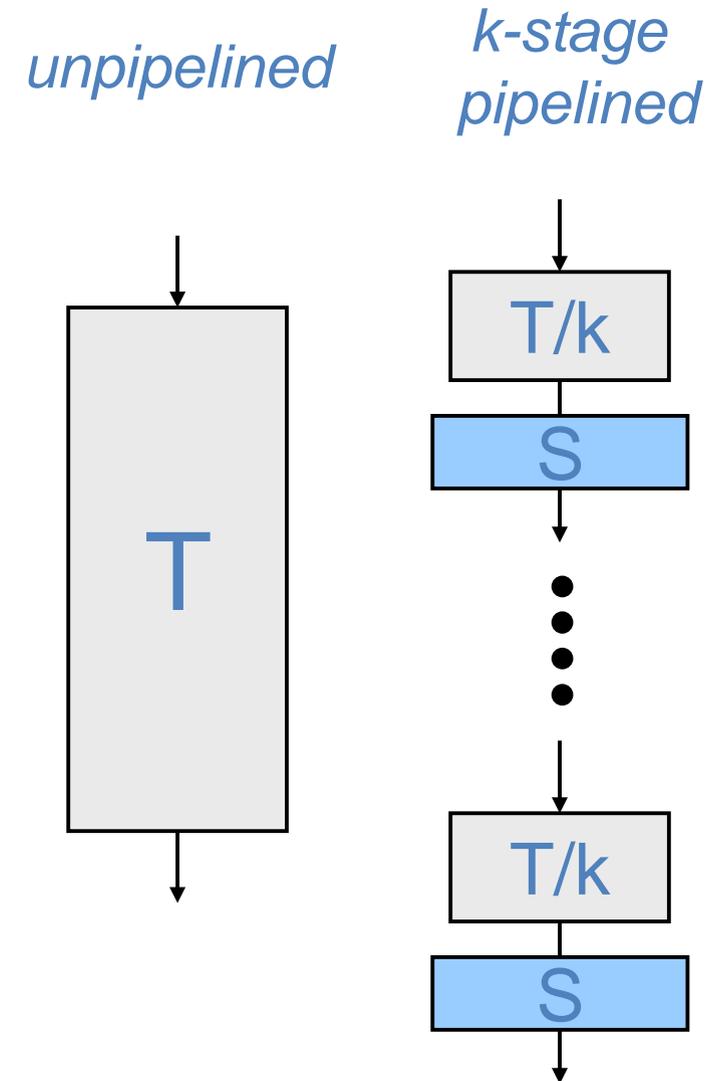
Pipelining Performance Model

- Starting from an un-pipelined version with propagation delay T and $BW = 1/T$

$$P_{\text{pipelined}} = BW_{\text{pipelined}} = 1 / (T/k + S)$$

where

S = delay through latch and overhead



Hardware Cost Model

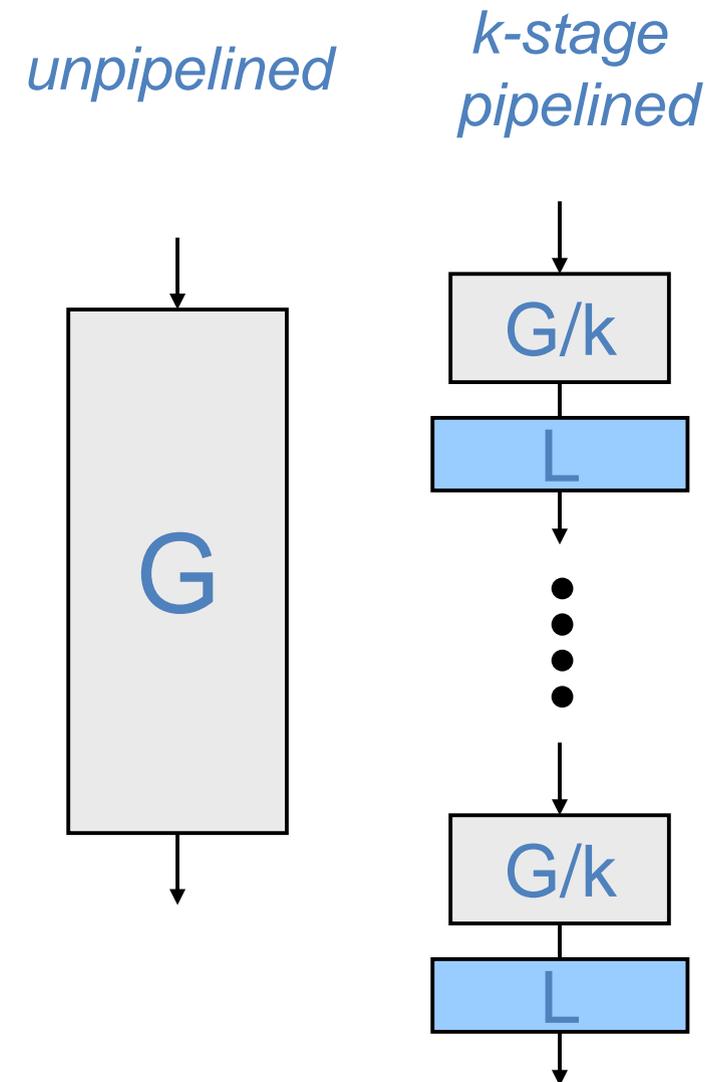
- Starting from an un-pipelined version with hardware cost G

$$\text{Cost}_{\text{pipelined}} = kL + G$$

where

L = cost of adding each latch, and

k = number of stages



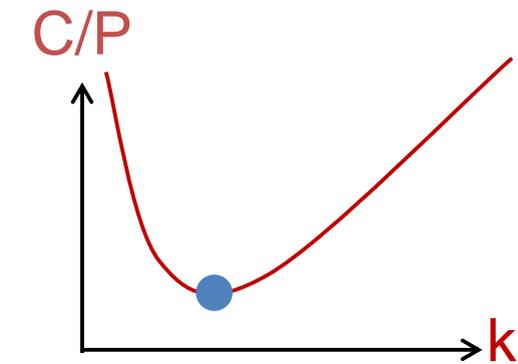
[Peter M. Kogge, 1981]

Cost/Performance Trade-off

Cost/Performance:

$$\begin{aligned} C/P &= [Lk + G] / [1/(T/k + S)] = (Lk + G) (T/k + S) \\ &= LT + GS + LS k + GT/k \end{aligned}$$

Optimal Cost/Performance: find min. C/P w.r.t. choice of k

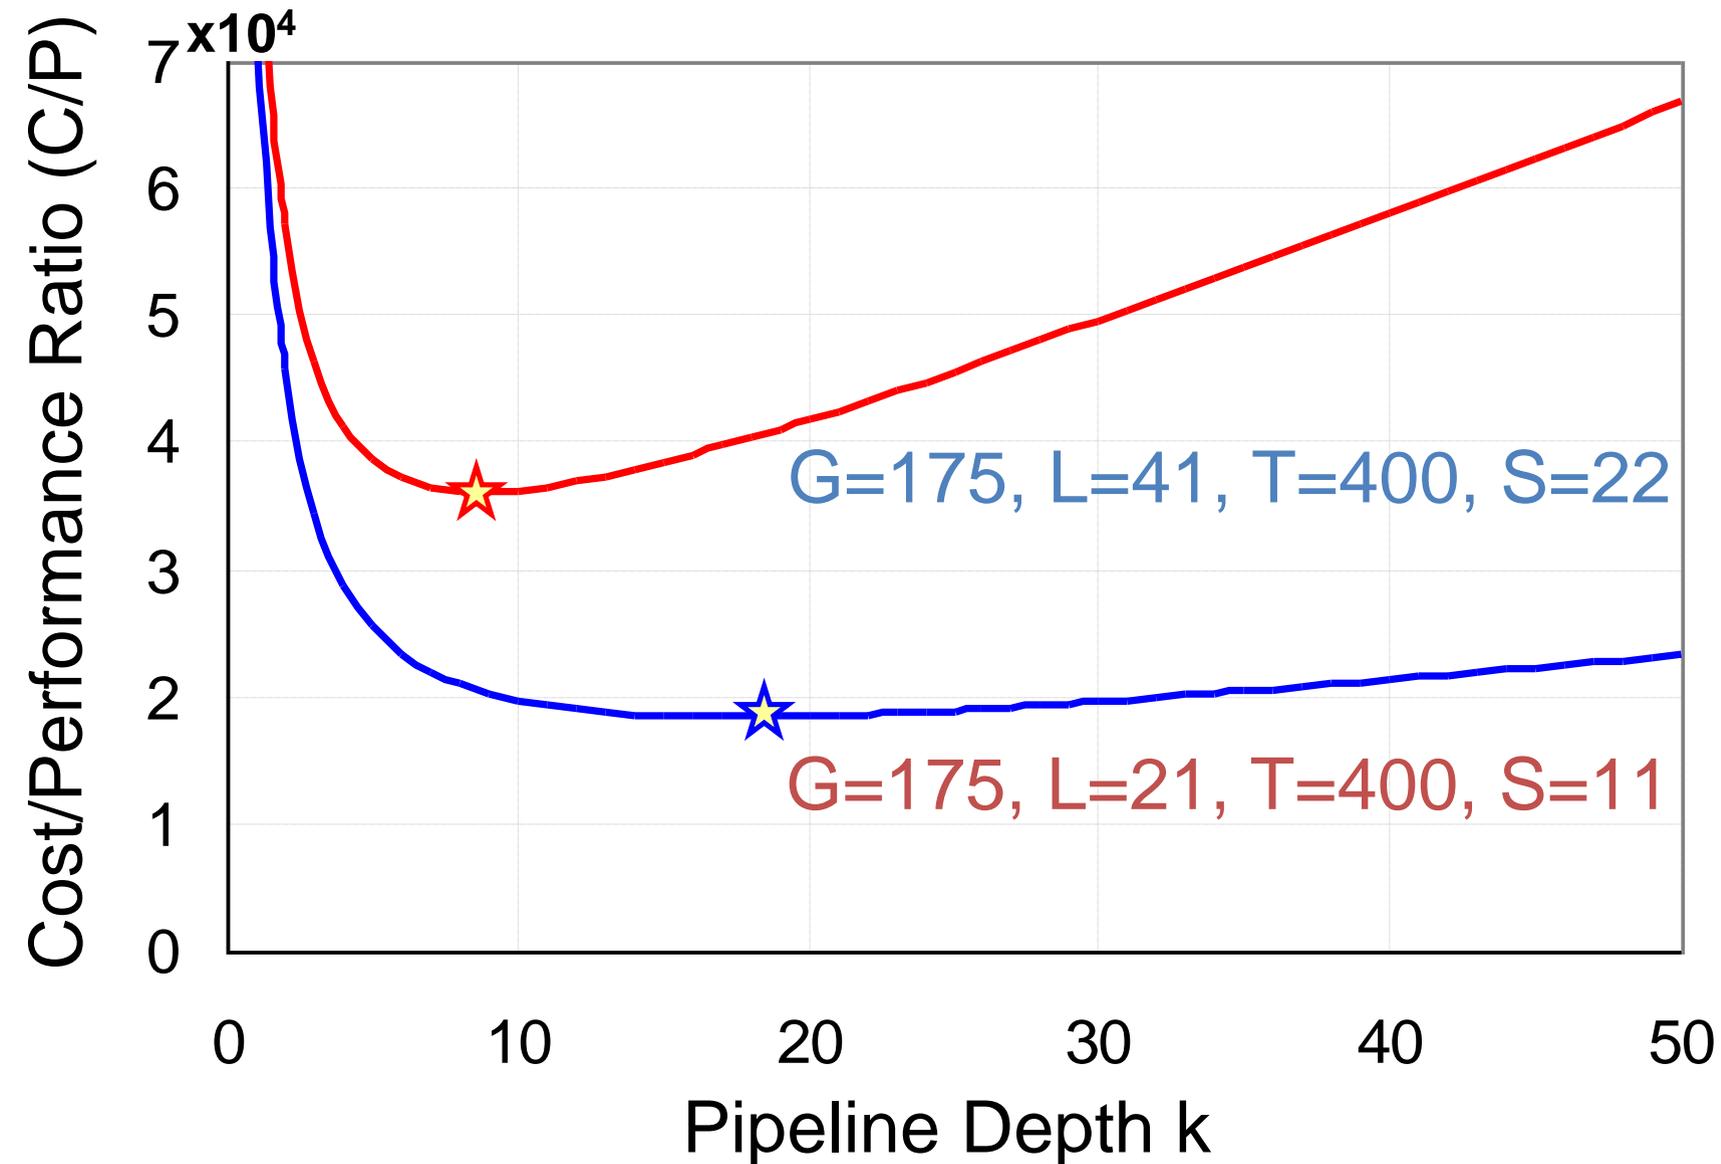


$$\frac{d}{dk} \left(\frac{Lk + G}{\frac{T}{k} + S} \right) = 0 + 0 + LS - \frac{GT}{k^2}$$

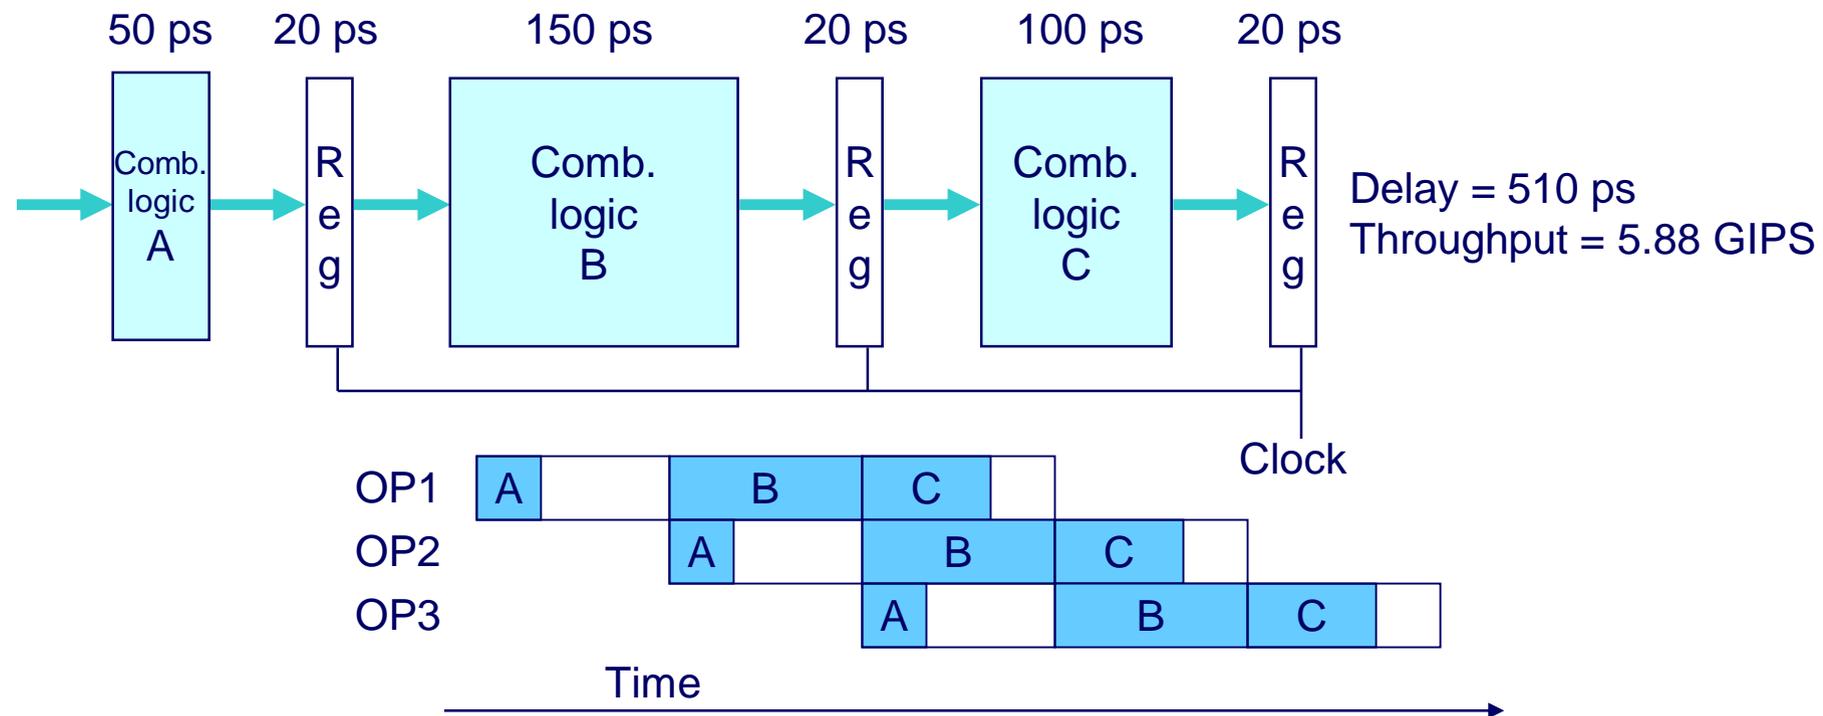
$$LS - \frac{GT}{k^2} = 0$$

$$k_{opt} = \sqrt{\frac{GT}{LS}}$$

"Optimal" Pipeline Depth (k_{opt}) Examples



Limitations: Nonuniform Delays



- Throughput limited by slowest stage
- Other stages sit idle for much of the time
- Challenging to partition system into balanced stages

Pipelining Idealistic Assumptions

➤ Uniform Sub-computations

The computation to be pipelined can be evenly partitioned into uniform-latency sub-operations

➤ Repetition of Identical Computations

The same computations are to be performed repeatedly on a large number of different inputs

➤ Repetition of Independent Computations

All the repetitions of the same computation are mutually independent, i.e. no data dependency and no resource conflicts

Instruction Pipeline Design

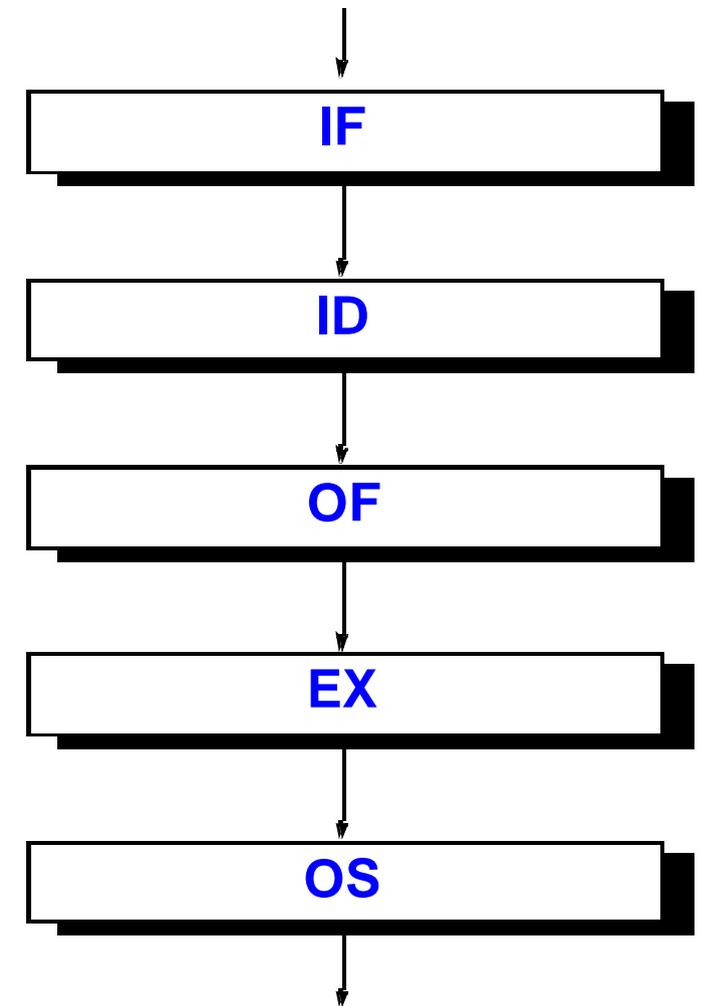
- **Uniform Sub-computations ... NOT!**
 - ⇒ **balancing pipeline stages**
 - stage quantization to yield balanced pipe stages
 - minimize internal fragmentation (some waiting stages)
- **Identical Computations ... NOT!**
 - ⇒ **unifying instruction types**
 - coalescing instruction types into one multi-function pipe
 - minimize external fragmentation (some idling stages)
- **Independent Computations ... NOT!**
 - ⇒ **resolving pipeline hazards**
 - inter-instruction dependency detection and resolution
 - minimize performance loss due to pipeline stalls

Instruction Pipelining: A Generic Pipeline

➤ “computation” to be pipelined.

- Instruction Fetch (IF)
- Instruction Decode (ID)
- Operand(s) Fetch (OF)
- Instruction Execution (EX)
- Operand Store (OS)
- Update Program Counter (PC)

1. Instruction Fetch
2. Instruction Decode
3. Operand Fetch
4. Instruction Execute
5. Operand Store



- Based on “obvious” subcomputations

Balancing Pipeline Stages

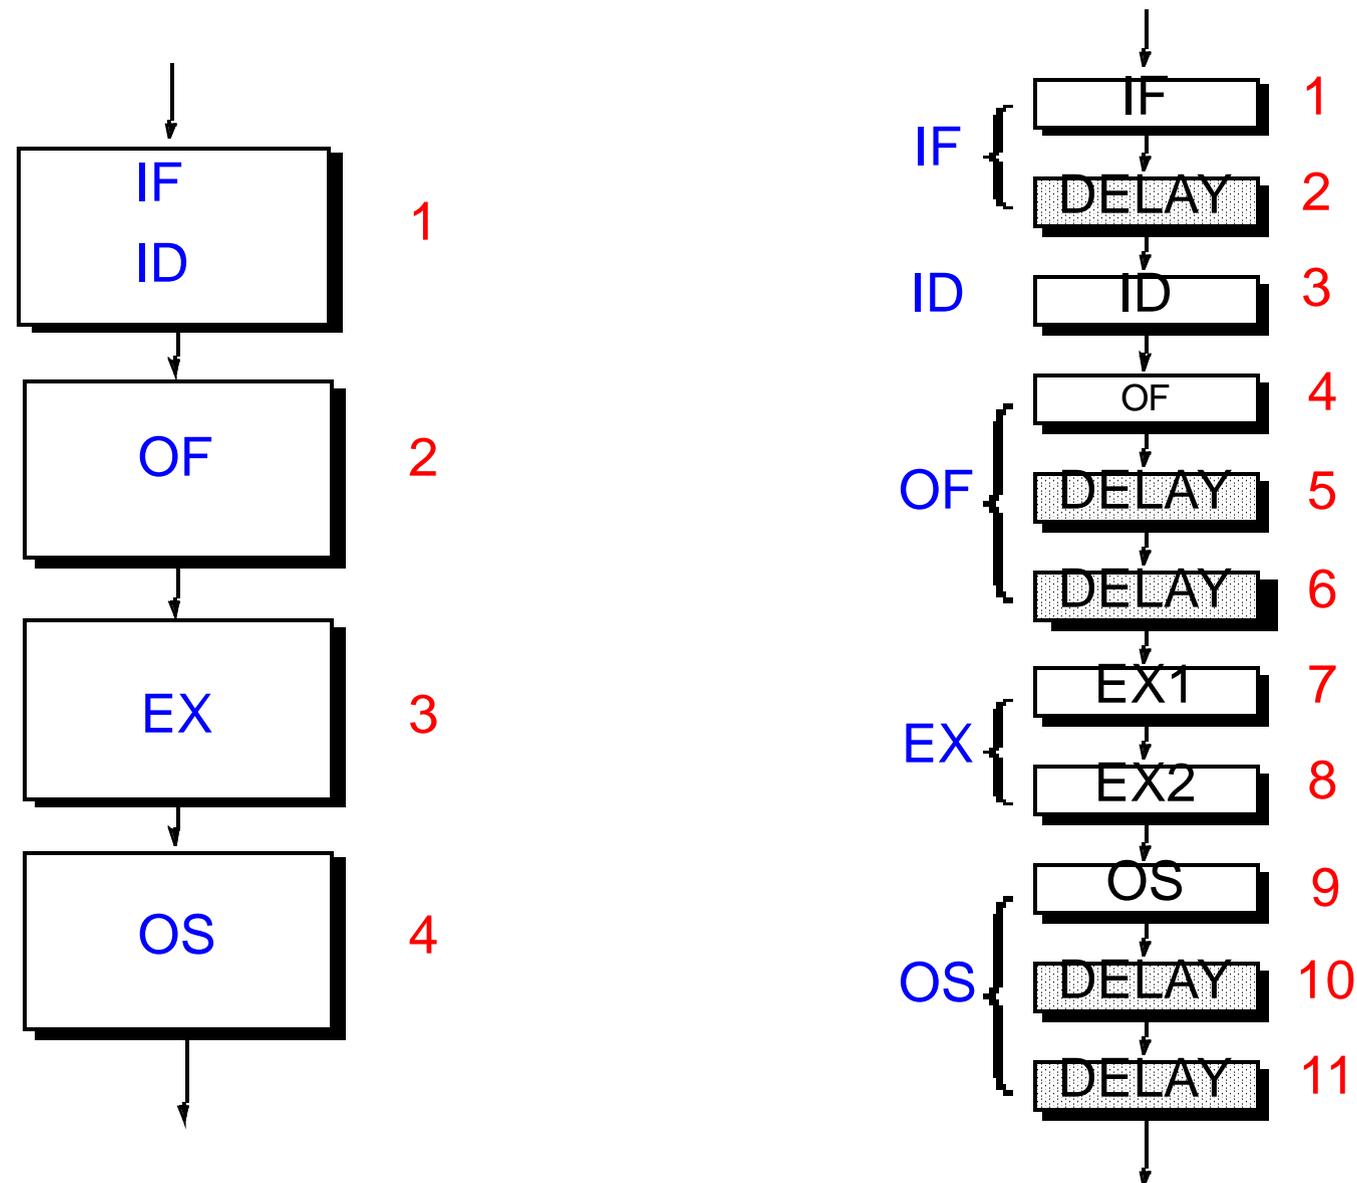
➤ Two Methods for Stage Quantization:

- Merging of multiple sub-computations into one.
- Subdividing a sub-computation into multiple sub-computations.

➤ Recent Trends:

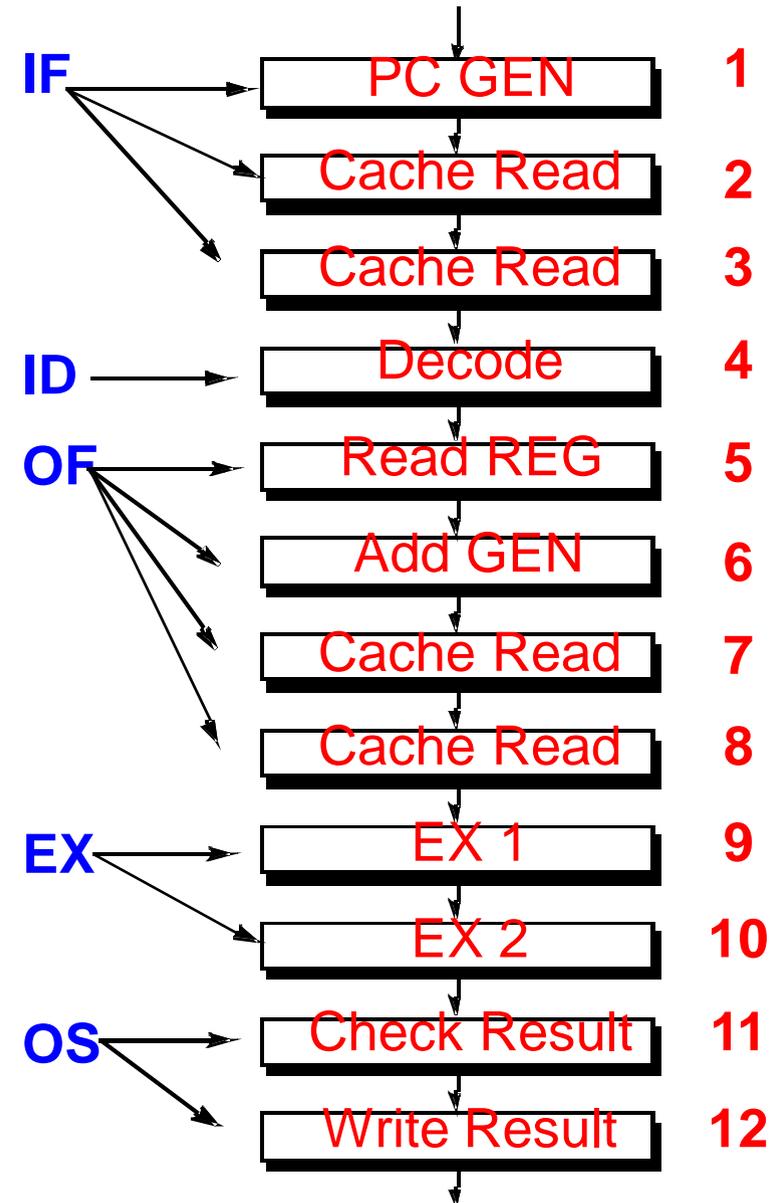
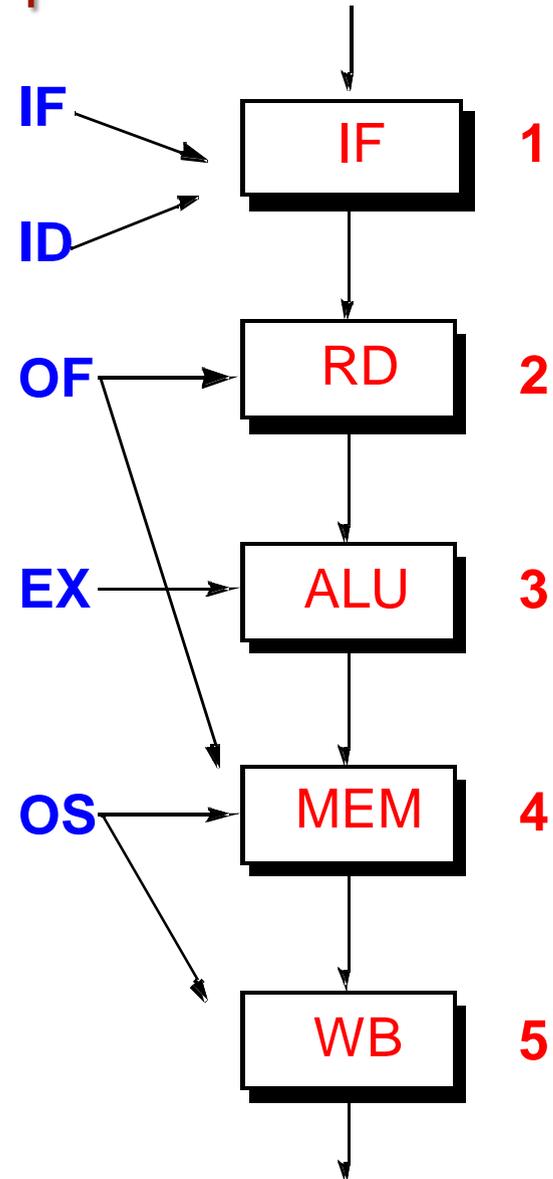
- Deeper pipelines with more and more stages.
- Multiplicity of different sub-pipelines.
- Pipelining of memory access becomes tricky.

Granularity of Pipeline Stages Can Vary



Actual Pipeline Examples

MIPS R2000/R3000



AMDAHL 470V/7

Unifying Instruction Types

➤ Procedure:

1. Classification of Instruction Types: Analyze the sequence of register transfers required by each instruction type.
2. Coalescing Resource Requirements: Find commonality across instruction types and merge them to share the same pipeline stage and hardware resource.
3. Instruction Pipeline Implementation: If there exists flexibility, shift or reorder some register transfers to facilitate further merging.

ALU Instruction Specification

Generic subcomputations	1. ALU Instruction Type:	
	Integer instruction	Floating-point instruction
IF	- Fetch instruction (access I-cache)	- Fetch instruction (access I-cache)
ID	- Decode instruction	- Decode instruction
OF	- Access register file	- Access FP register file
EX	- Perform ALU operation	- Perform FP operation
OS	- Write back to reg. file	- Write back to FP reg. file

Memory Instruction Specification

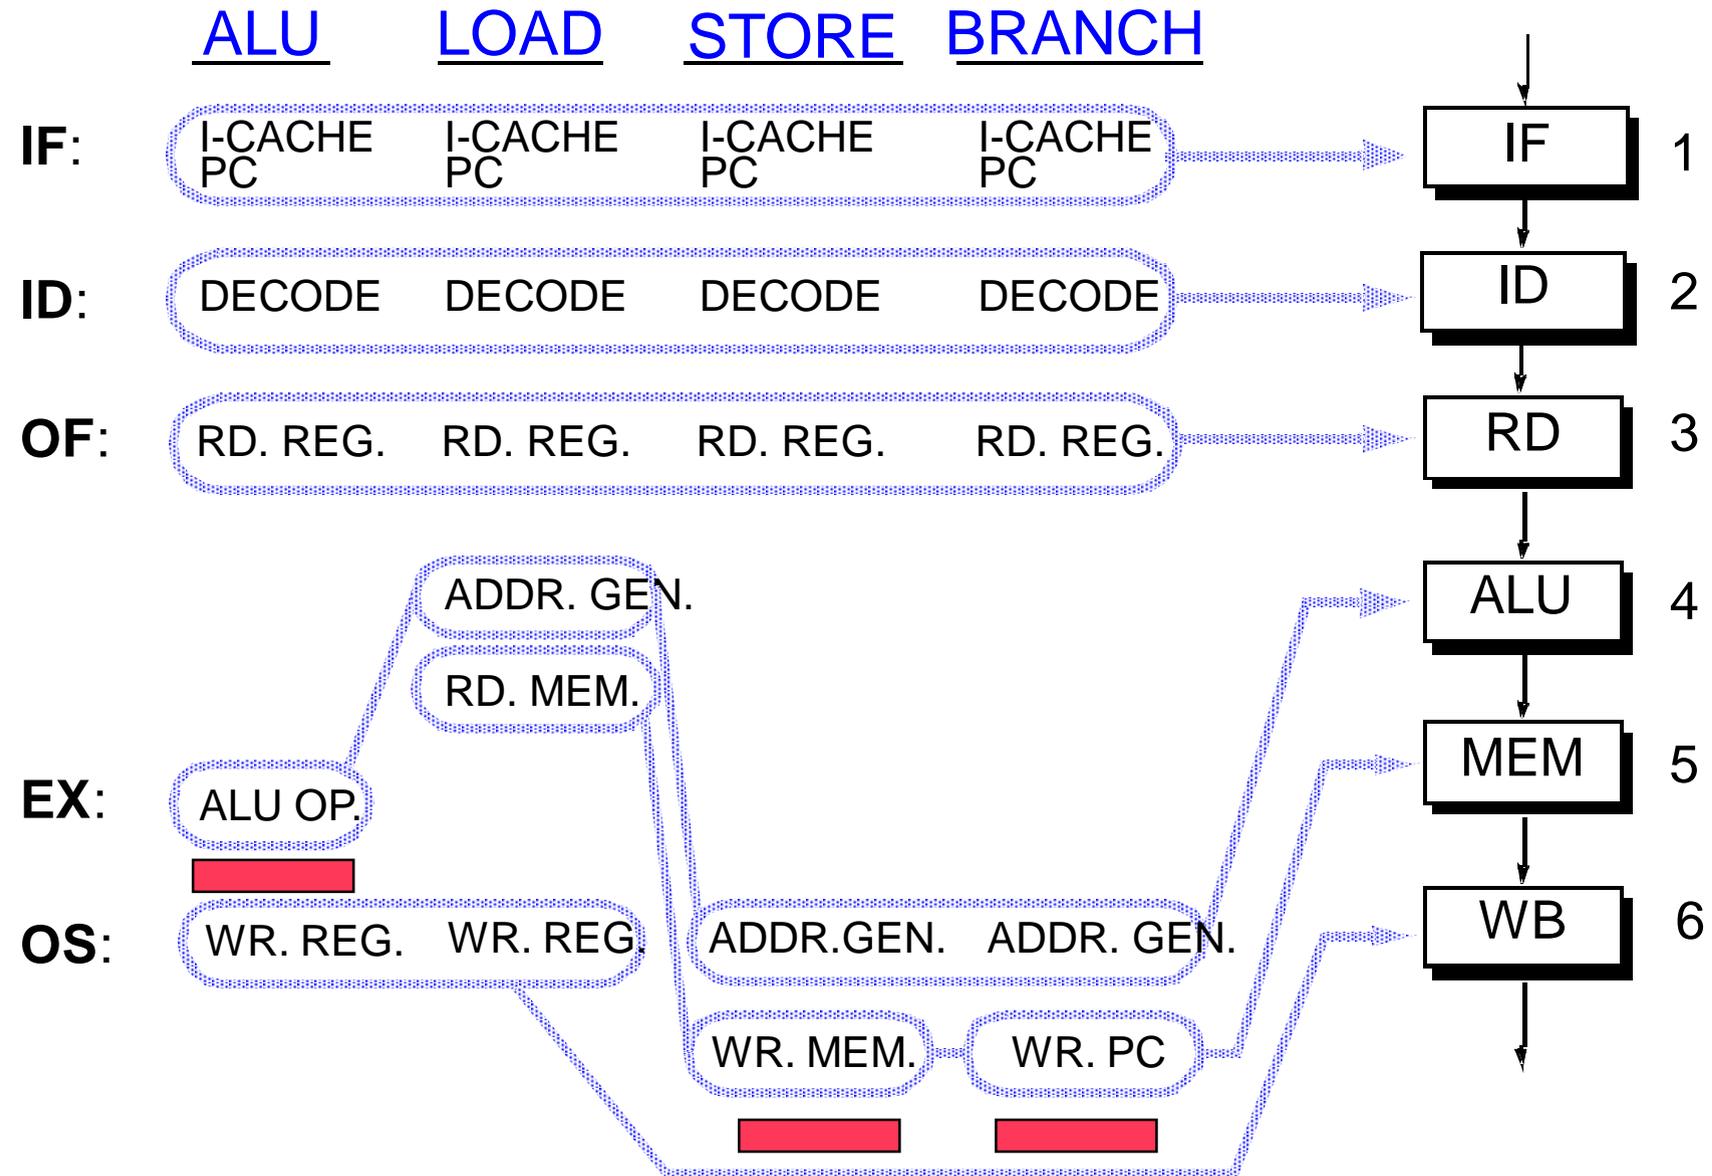
Generic subcomputations	2. Load/Store Instruction Type:	
	Load instruction	Store instruction
IF	- Fetch instruction (access I-cache)	- Fetch instruction (access I-cache)
ID	- Decode instruction	- Decode instruction
OF	- Access register file (base address) - Generate effective address (base + offset) - Access (read) memory location (D-cache)	- Access register file (register operand, and base address)
EX	-	-
OS	- Write back to reg. file	- Generate effective address (base + offset) - Access (write) memory location (D-cache)

Branch Instruction Specification

Generic subcomputations	3. Branch Instruction Type:	
	Jump (uncond.) instruction	Conditional branch instr.
IF	- Fetch instruction (access I-cache)	- Fetch instruction (access I-cache)
ID	- Decode instruction	- Decode instruction
OF	- Access register file (base address) - Generate effective address (base + offset)	- Access register file (base address) - Generate effective address (base + offset)
EX	-	- Evaluate branch condition
OS	- Update program counter with target address	- If condition is true, update program counter with target address

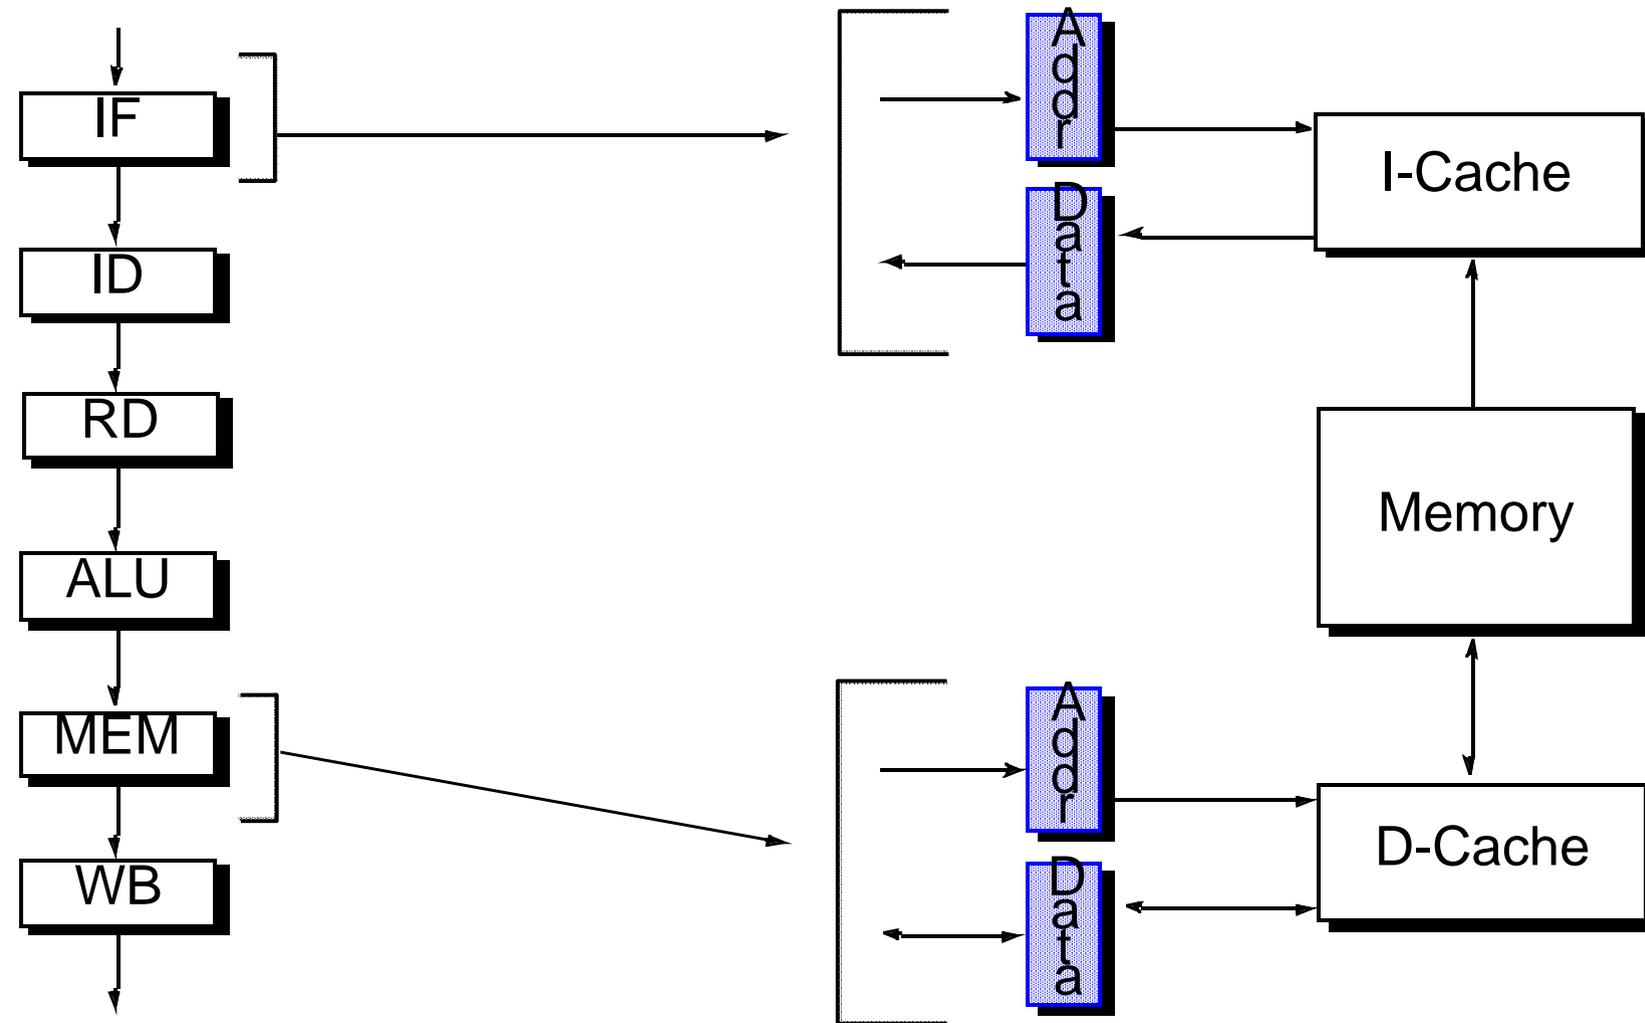
Unifying Instruction Types

The 6-stage TYPICAL pipeline:



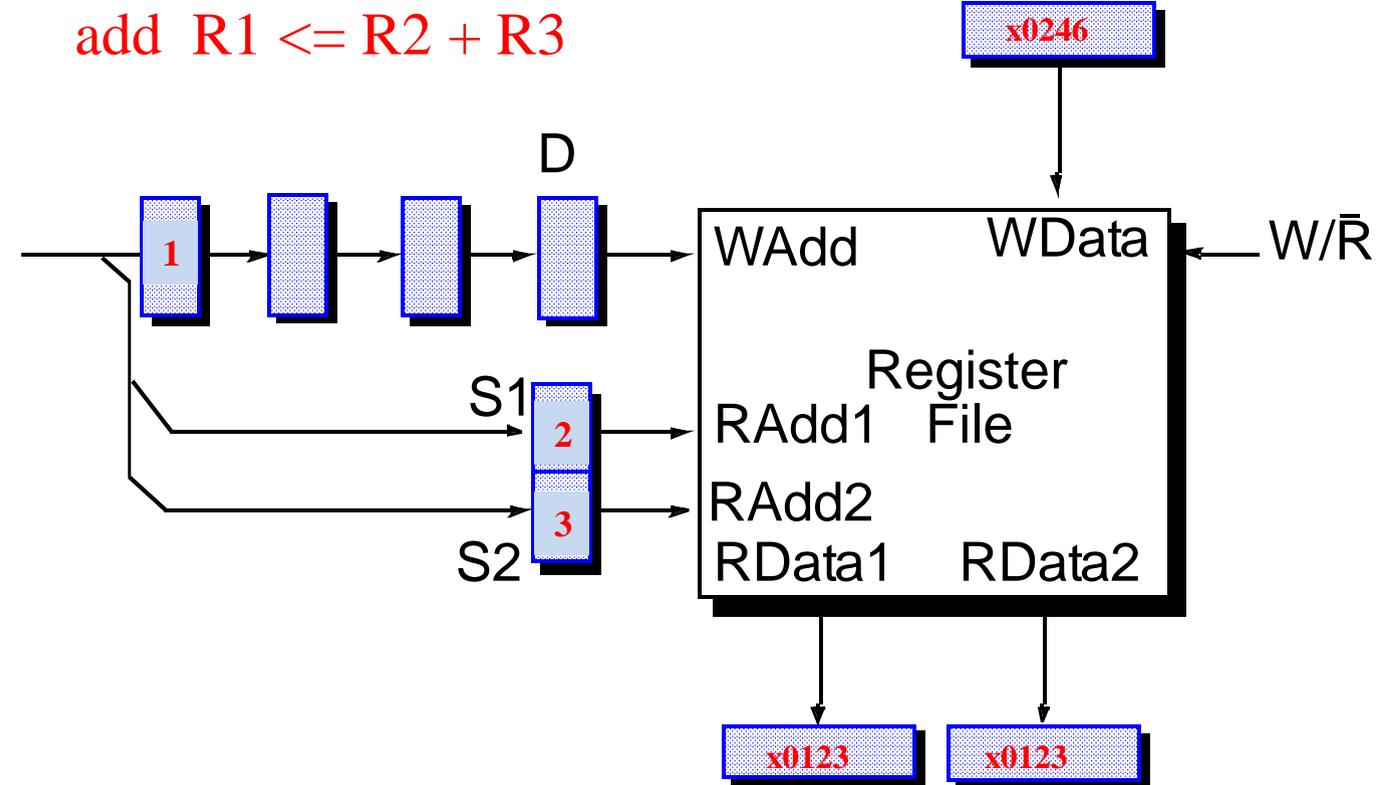
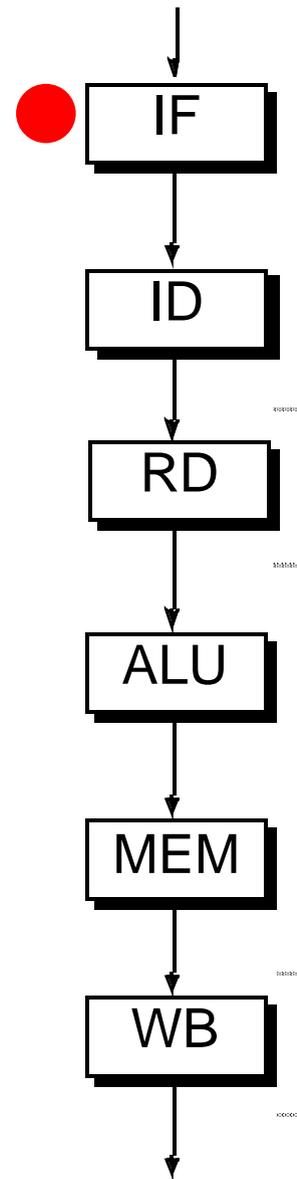
Pipeline Interface to Memory Subsystem

The 6-stage TYPICAL pipeline:



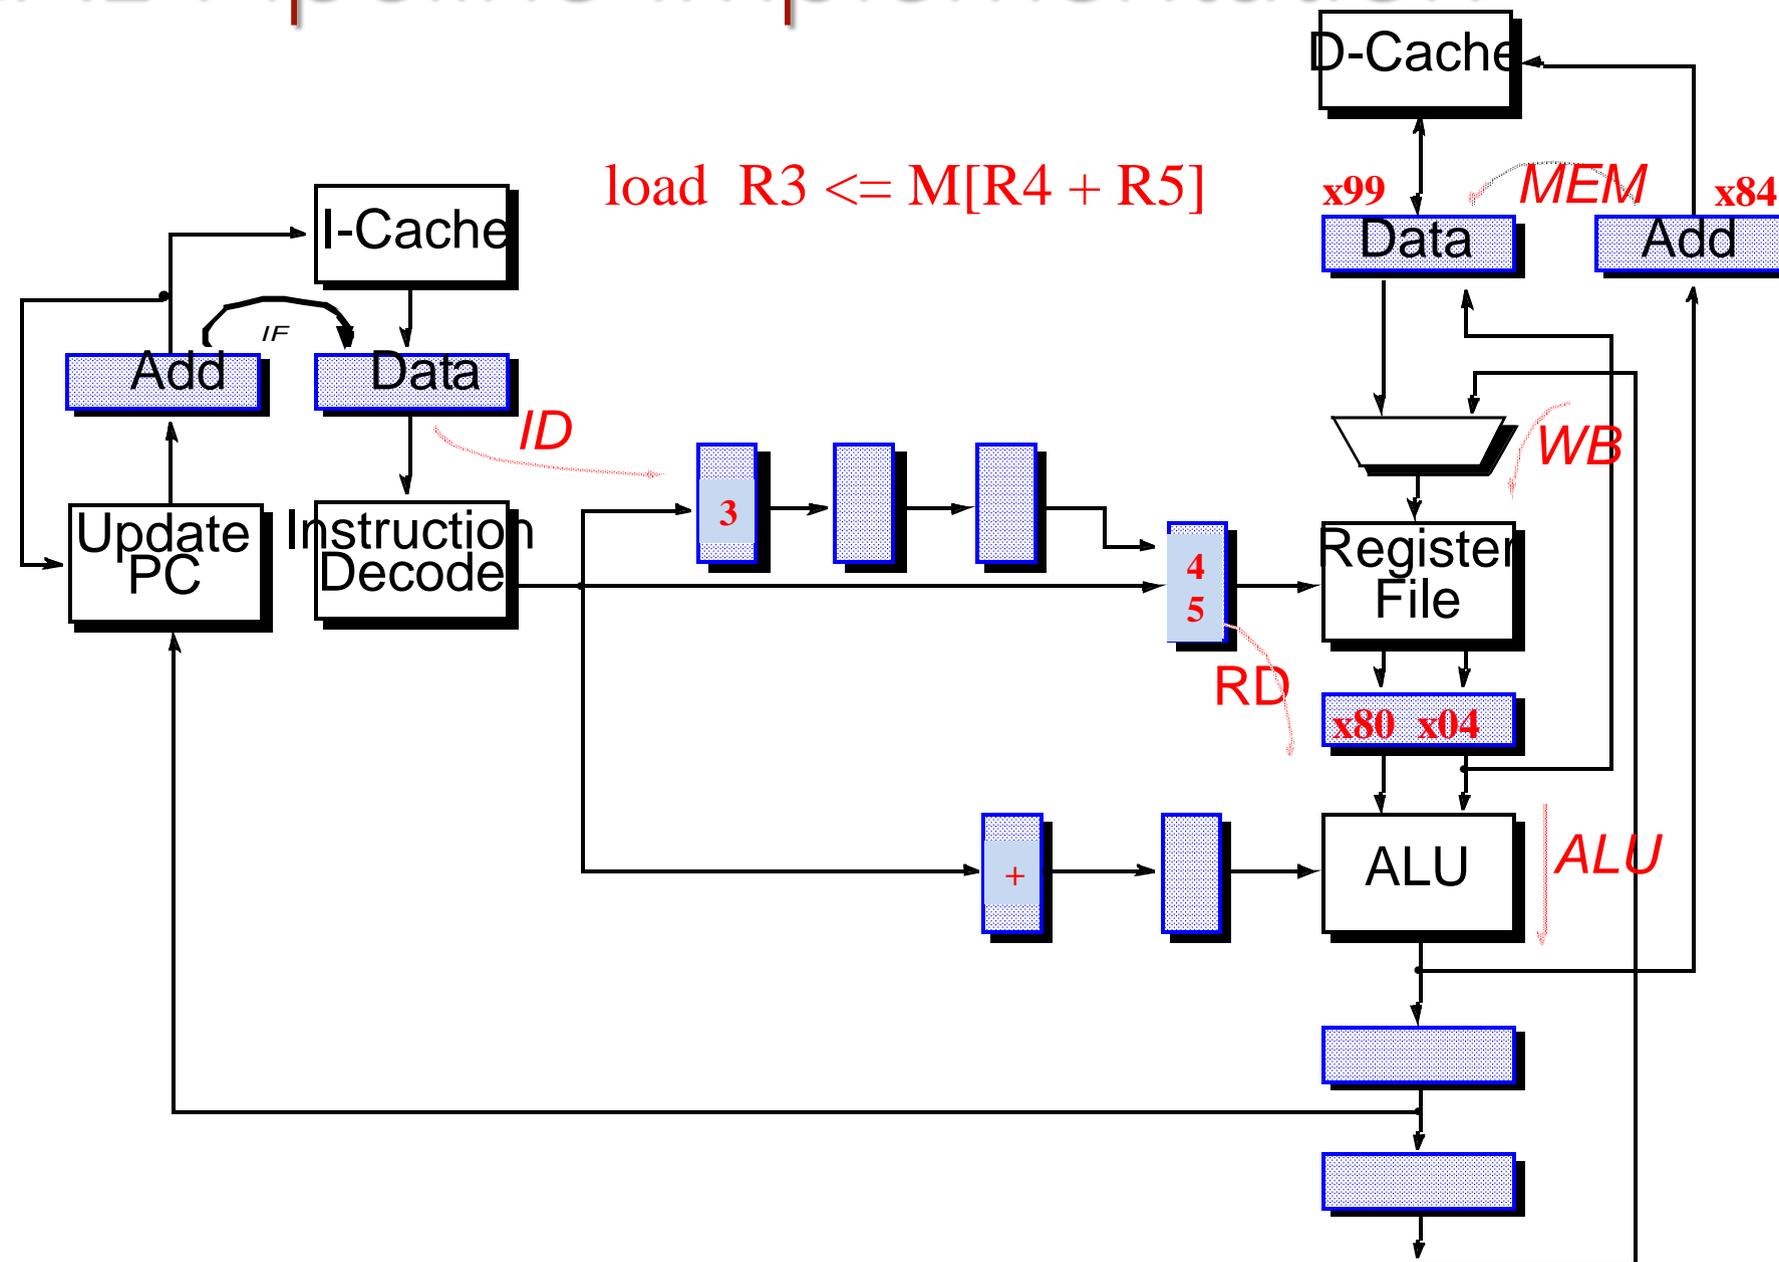
Pipeline Interface to Register File

The 6-stage TYPICAL pipeline:



TYPICAL Pipeline Implementation

The 6-stage TYPICAL pipeline:

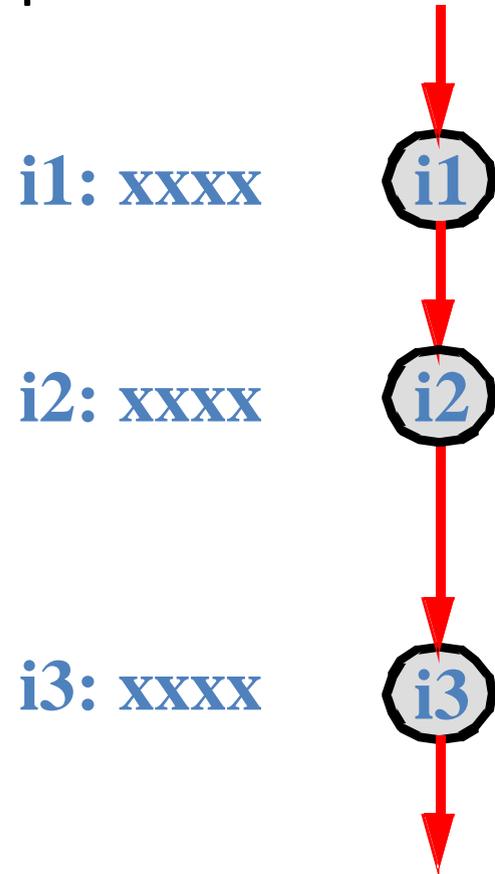


Pipelining Idealistic Assumptions

- ☑ Uniform subcomputations
 - Can pipeline into stages with equal delay
 - Balance pipeline stages
- ☑ Identical computations
 - Can fill pipeline with identical work
 - Unify instruction types
- Independent computations
 - No relationships between work units
 - Resolve Pipeline Hazards
 - Minimize pipeline stalls

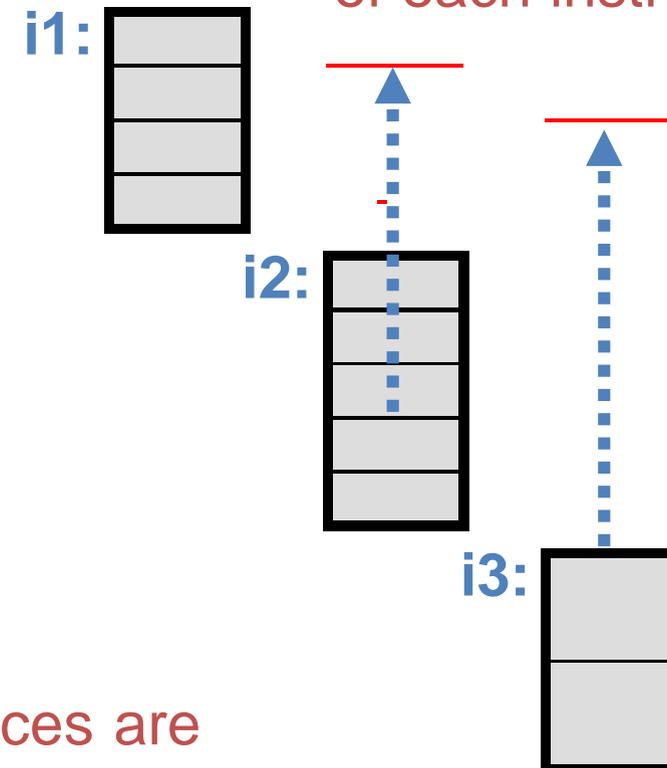
Instruction Dependencies & Pipeline Hazards

Sequential Code Semantics



The implied sequential precedences are over-specifications. It is sufficient but not necessary to ensure program correctness.

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



Inter-Instruction Dependencies

◆ Data dependency

$$\begin{array}{l} r_3 \leftarrow r_1 \text{ op } r_2 \\ r_5 \leftarrow r_3 \text{ op } r_4 \end{array}$$

Read-after-Write
(RAW)

◆ Anti-dependency

$$\begin{array}{l} r_3 \leftarrow r_1 \text{ op } r_2 \\ r_1 \leftarrow r_4 \text{ op } r_5 \end{array}$$

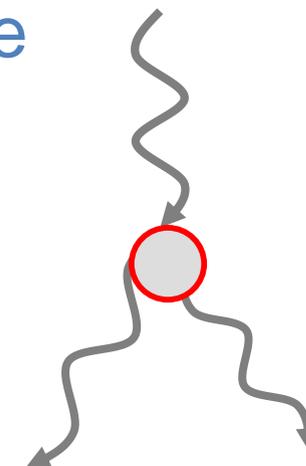
Write-after-Read
(WAR)

◆ Output dependency

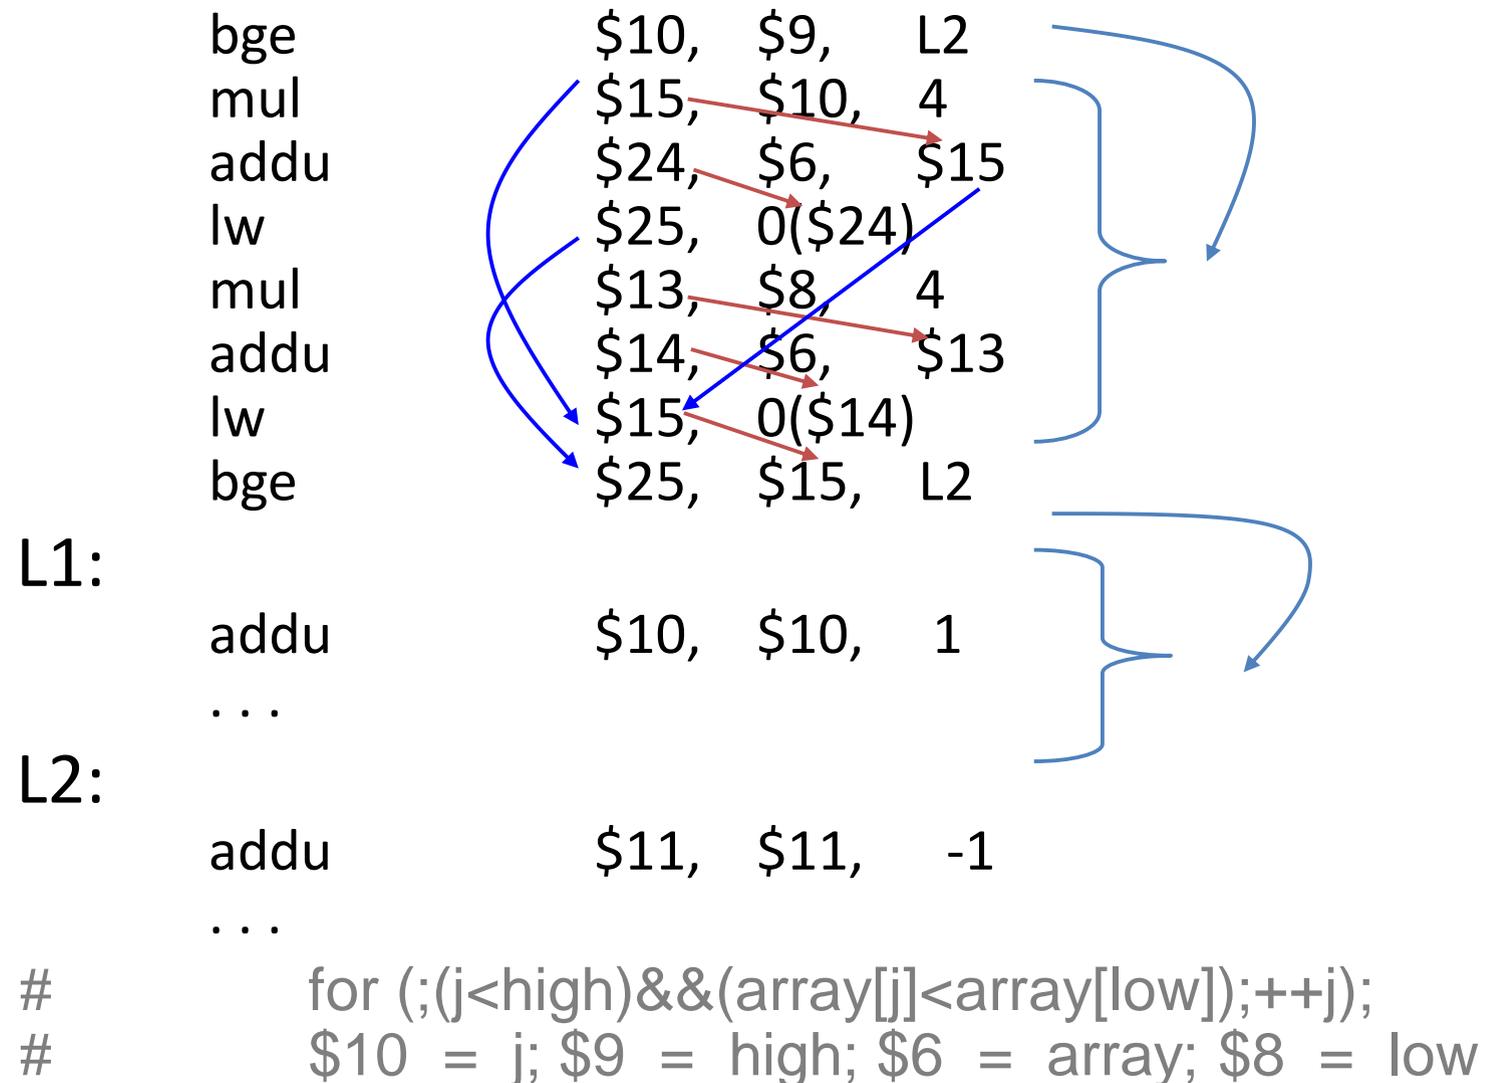
$$\begin{array}{l} r_3 \leftarrow r_1 \text{ op } r_2 \\ r_5 \leftarrow r_3 \text{ op } r_4 \\ r_3 \leftarrow r_6 \text{ op } r_7 \end{array}$$

Write-after-Write
(WAW)

◆ Control dependency



Example: Quick Sort for MIPS



Resolving Pipeline Hazards

➤ Pipeline Hazards:

- Potential violations of program dependencies
- Must ensure program dependencies are not violated

➤ Hazard Resolution:

- Static Method: Performed at compiled time in software
- Dynamic Method: Performed at run time using hardware

➤ Pipeline Interlock:

- Hardware mechanisms for dynamic hazard resolution
- Must detect and enforce dependencies at run time

Pipeline Hazards

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

Hazards due to Memory Data Dependences

<u>Pipe Stage</u>	<u>ALU Inst.</u>	<u>Load</u> inst.	<u>Store</u> inst.	<u>Branch</u> inst.
1. IF	I-cache PC<PC+4	I-cache PC<PC+4	I-cache PC<PC+4	I-cache PC<PC+4
2. ID	decode	decode	decode	decode
3. RD	read reg.	read reg.	read reg.	read reg.
4. ALU	ALU op.	addr. gen.	addr. gen.	addr. gen. cond. gen.
5. MEM	-----	read mem.	write mem.	PC<-br. addr.
6. WB	write reg.	write reg.	-----	-----

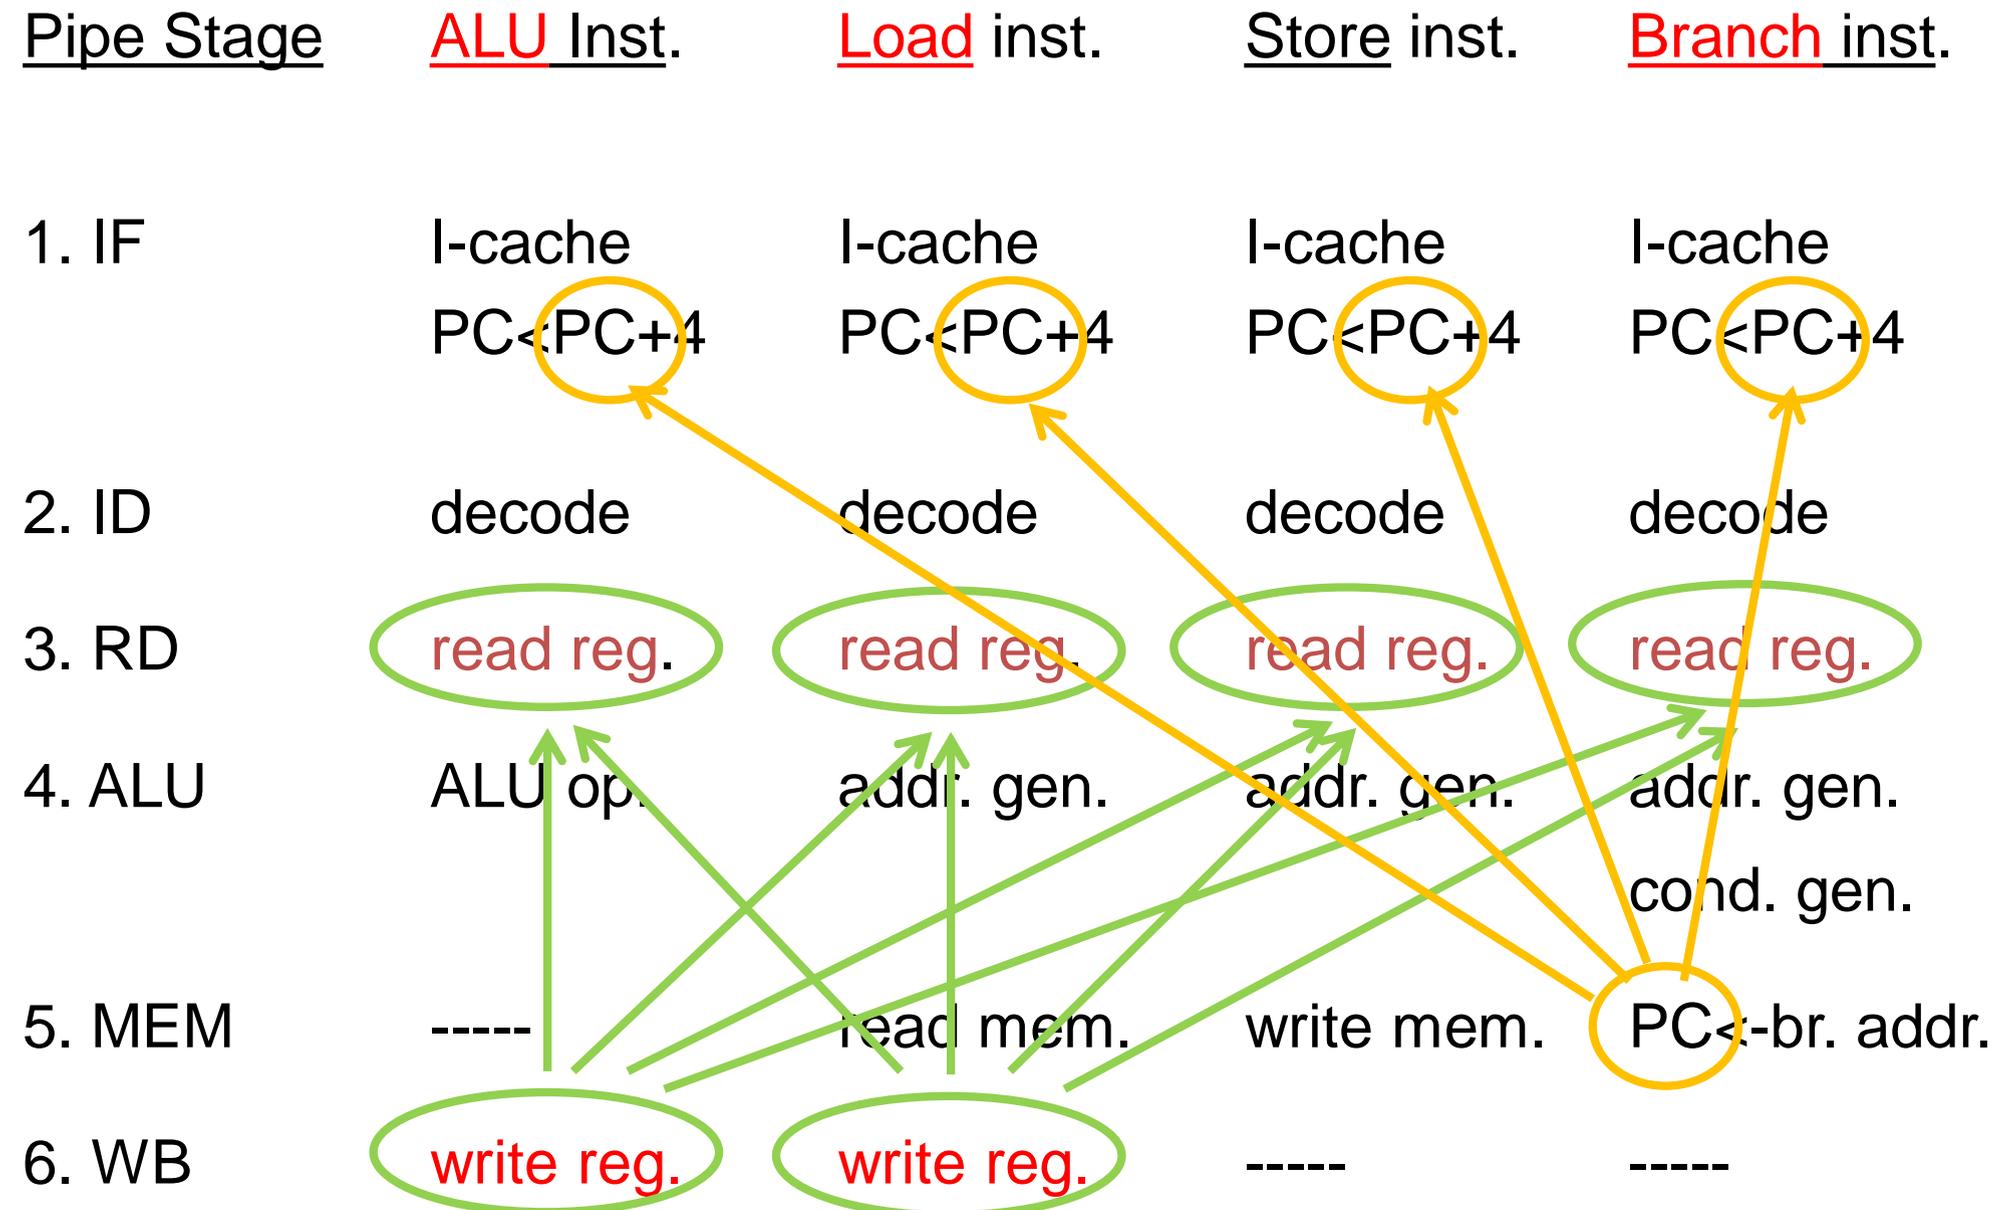
Hazards due to Register Data Dependences

<u>Pipe Stage</u>	<u>ALU</u> Inst.	<u>Load</u> inst.	<u>Store</u> inst.	<u>Branch</u> inst.
1. IF	I-cache PC<PC+4	I-cache PC<PC+4	I-cache PC<PC+4	I-cache PC<PC+4
2. ID	decode	decode	decode	decode
3. RD	read reg.	read reg.	read reg.	read reg.
4. ALU	ALU op.	addr. gen.	addr. gen.	addr. gen. cond. gen.
5. MEM	-----	read mem.	write mem.	PC<-br. addr.
6. WB	write reg.	write reg.	-----	-----

Hazards due to Control Dependences

<u>Pipe Stage</u>	<u>ALU</u> Inst.	<u>Load</u> inst.	<u>Store</u> inst.	<u>Branch</u> inst.
1. IF	I-cache PC<PC+4	I-cache PC<PC+4	I-cache PC<PC+4	I-cache PC<PC+4
2. ID	decode	decode	decode	decode
3. RD	read reg.	read reg.	read reg.	read reg.
4. ALU	ALU op.	addr. gen.	addr. gen.	addr. gen. cond. gen.
5. MEM	-----	read mem.	write mem.	PC<-br. addr.
6. WB	write reg.	write reg.	-----	-----

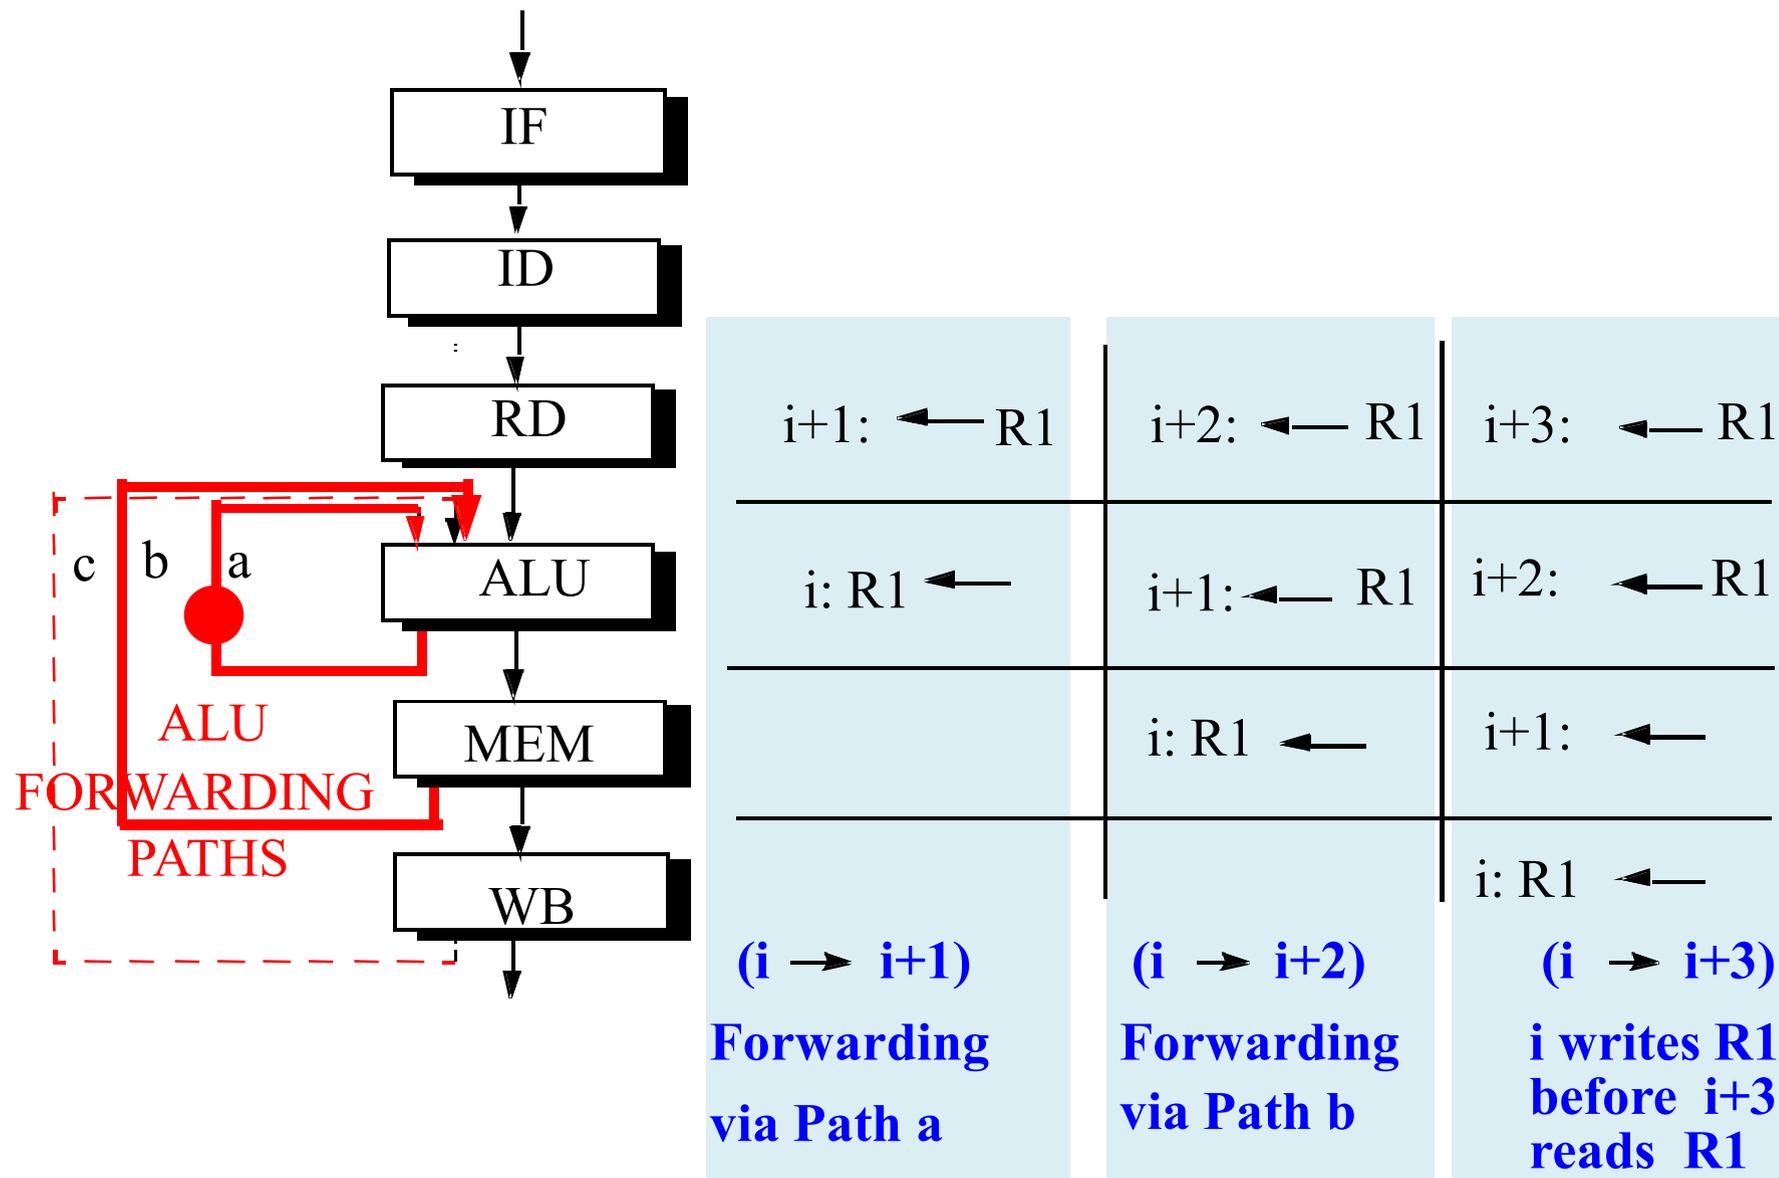
Inter-Instruction Hazards



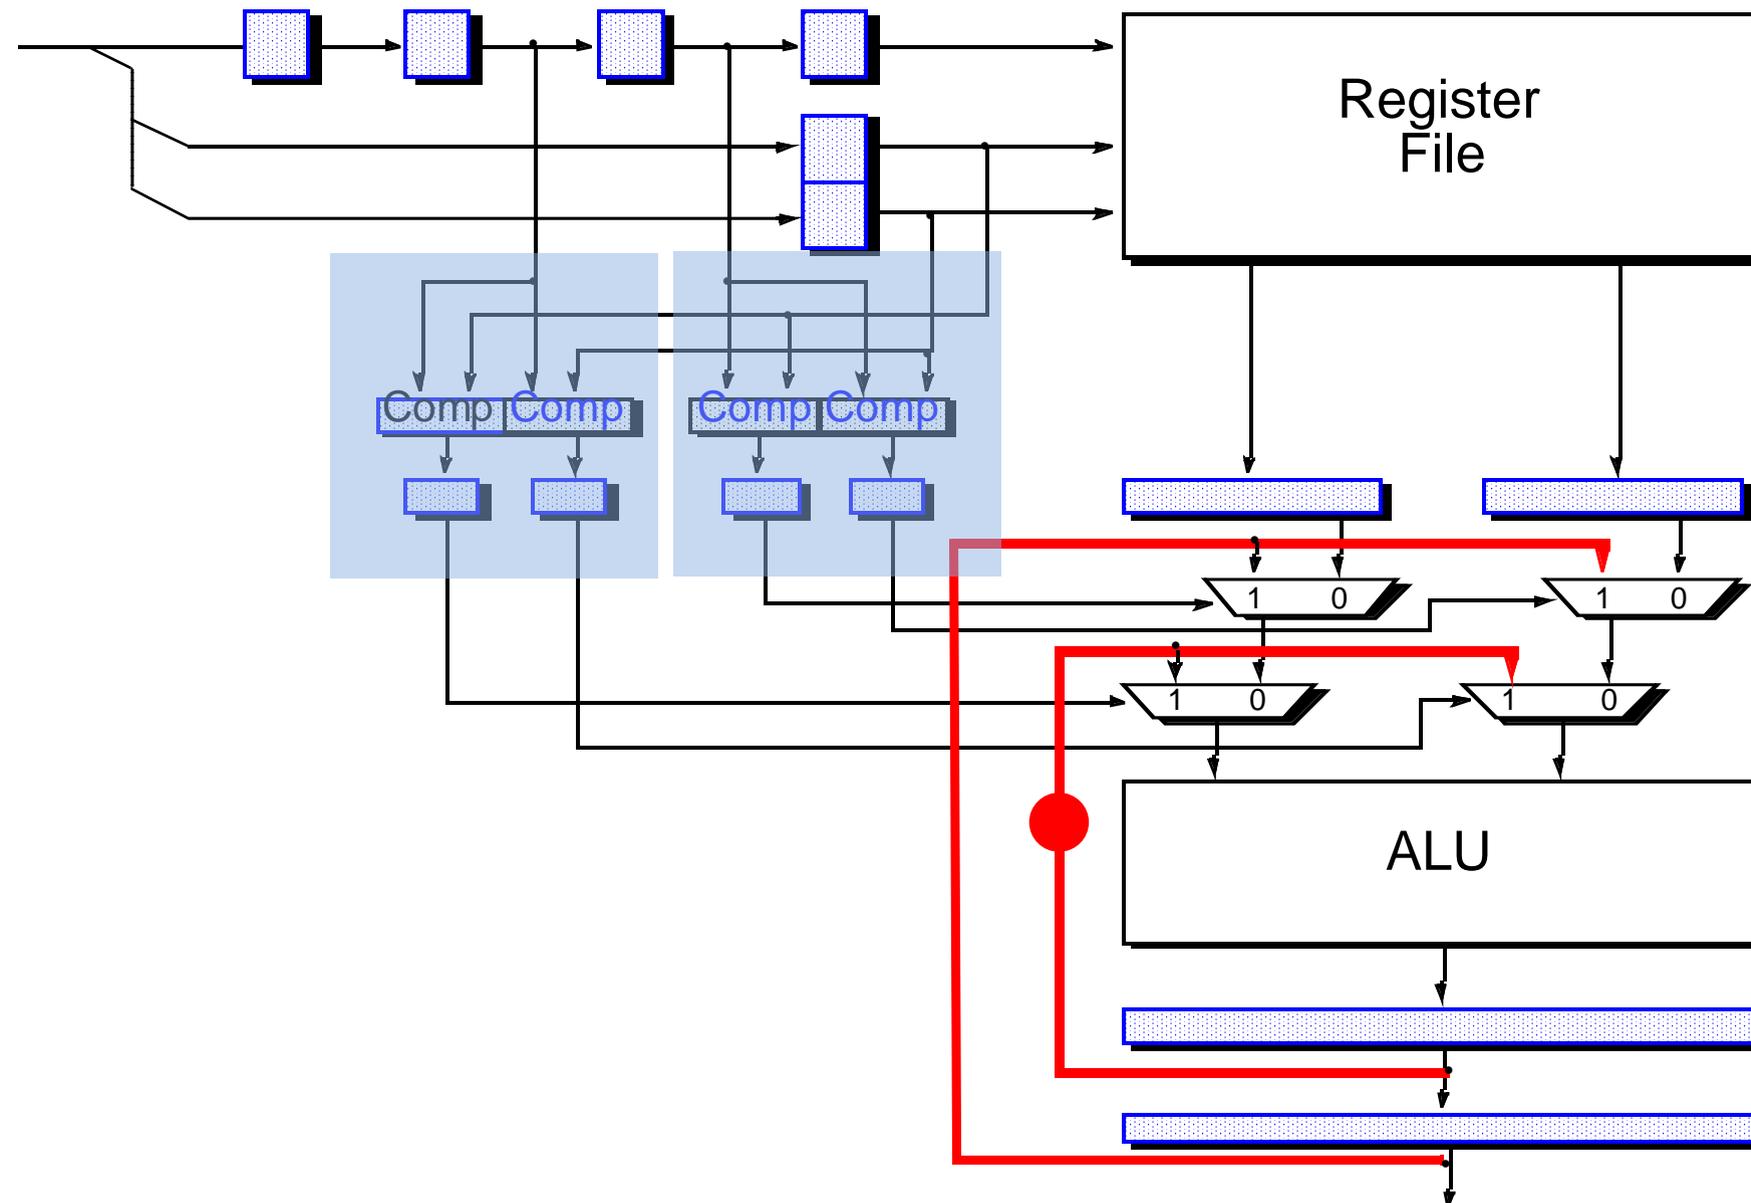
Dealing with Data Hazards

- Must first detect **RAW hazards**
 - Compare read register specifiers for newer instructions with write register specifiers for older instructions
 - Newer instruction in ID; older instructions in EX, MEM
- Resolve hazard dynamically
 - Stall or forward
- Not all hazards because
 - No register written (store or branch)
 - No register is read (e.g. addi, jump)
 - Do something only if necessary
 - Use special encodings for these cases to prevent spurious detection

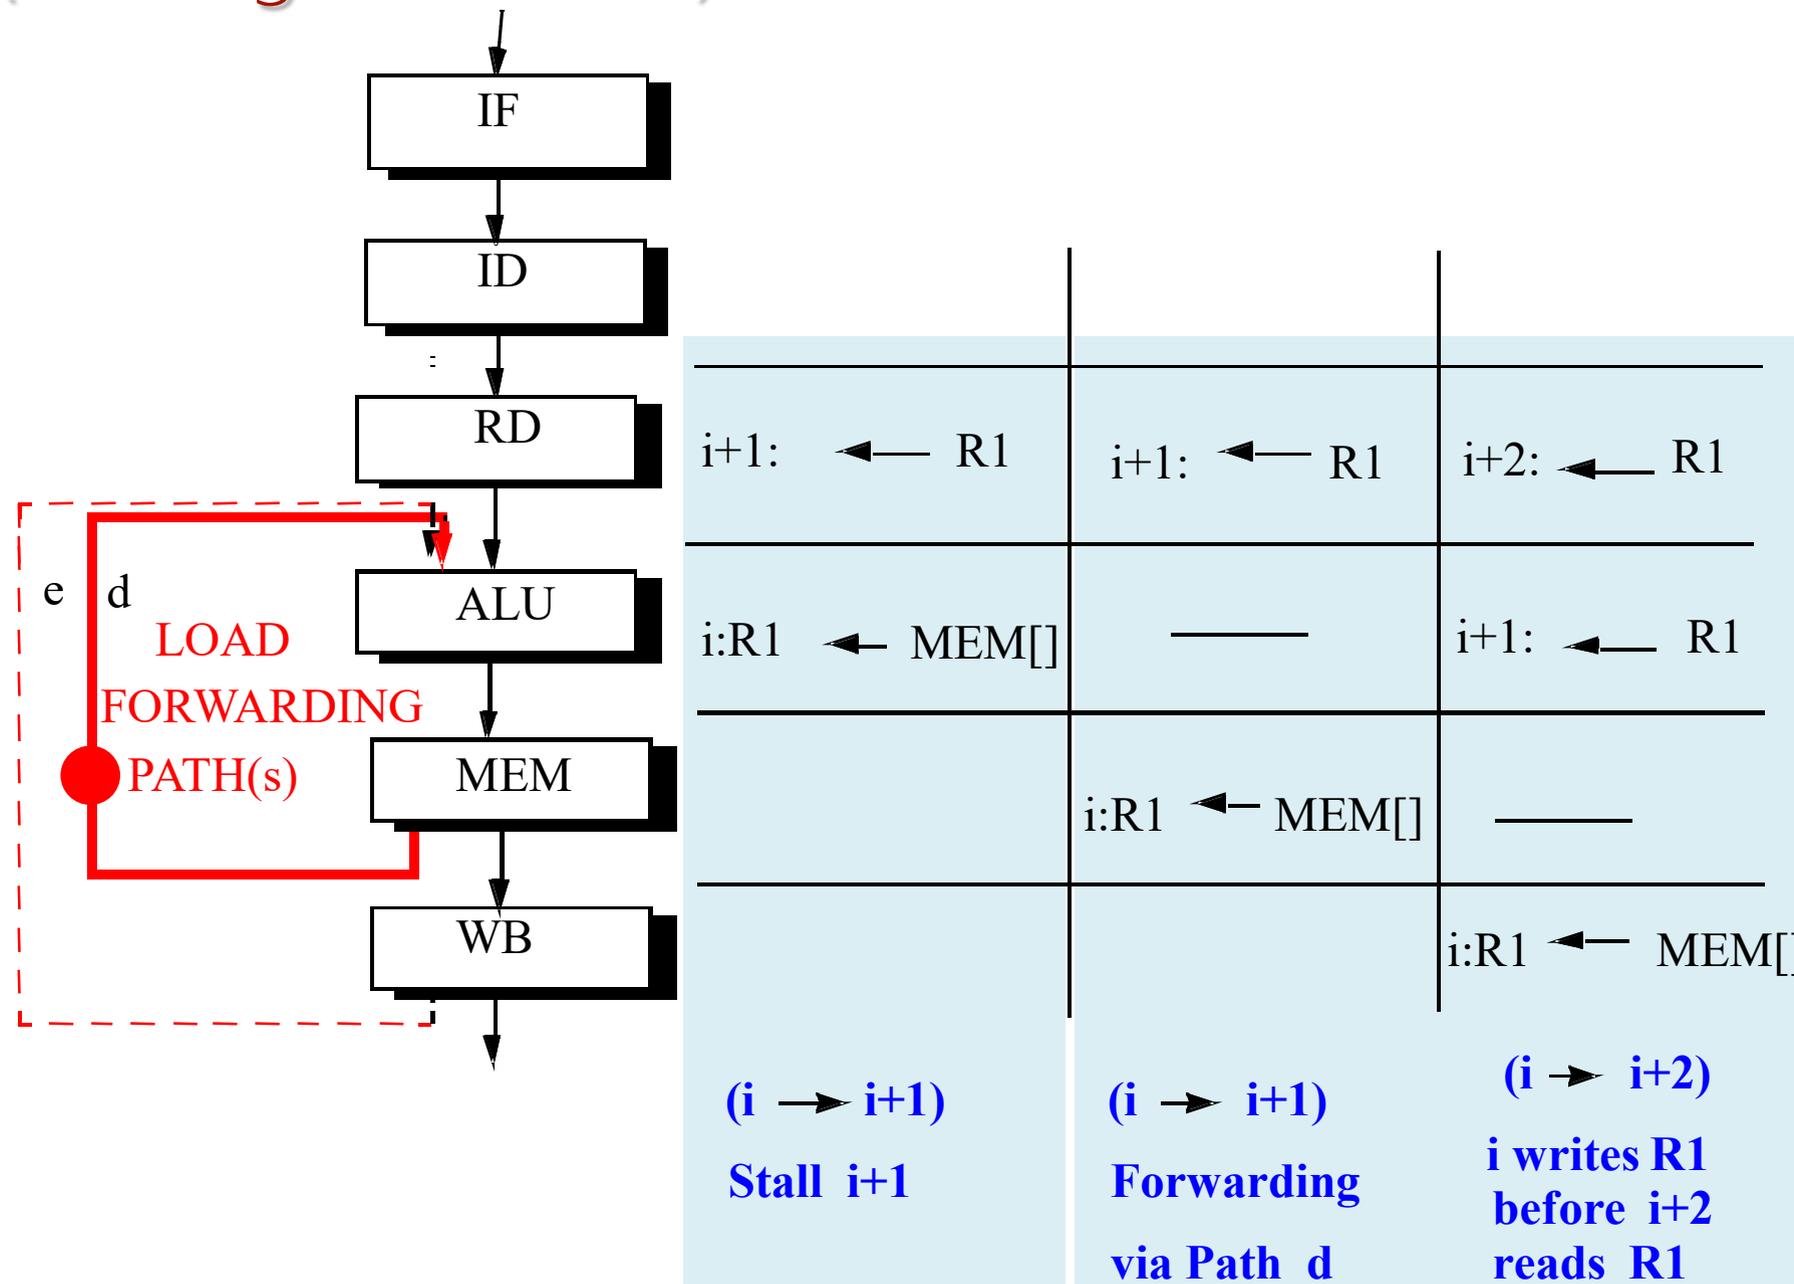
ALU (Leading Instruction) Interlock and Penalty



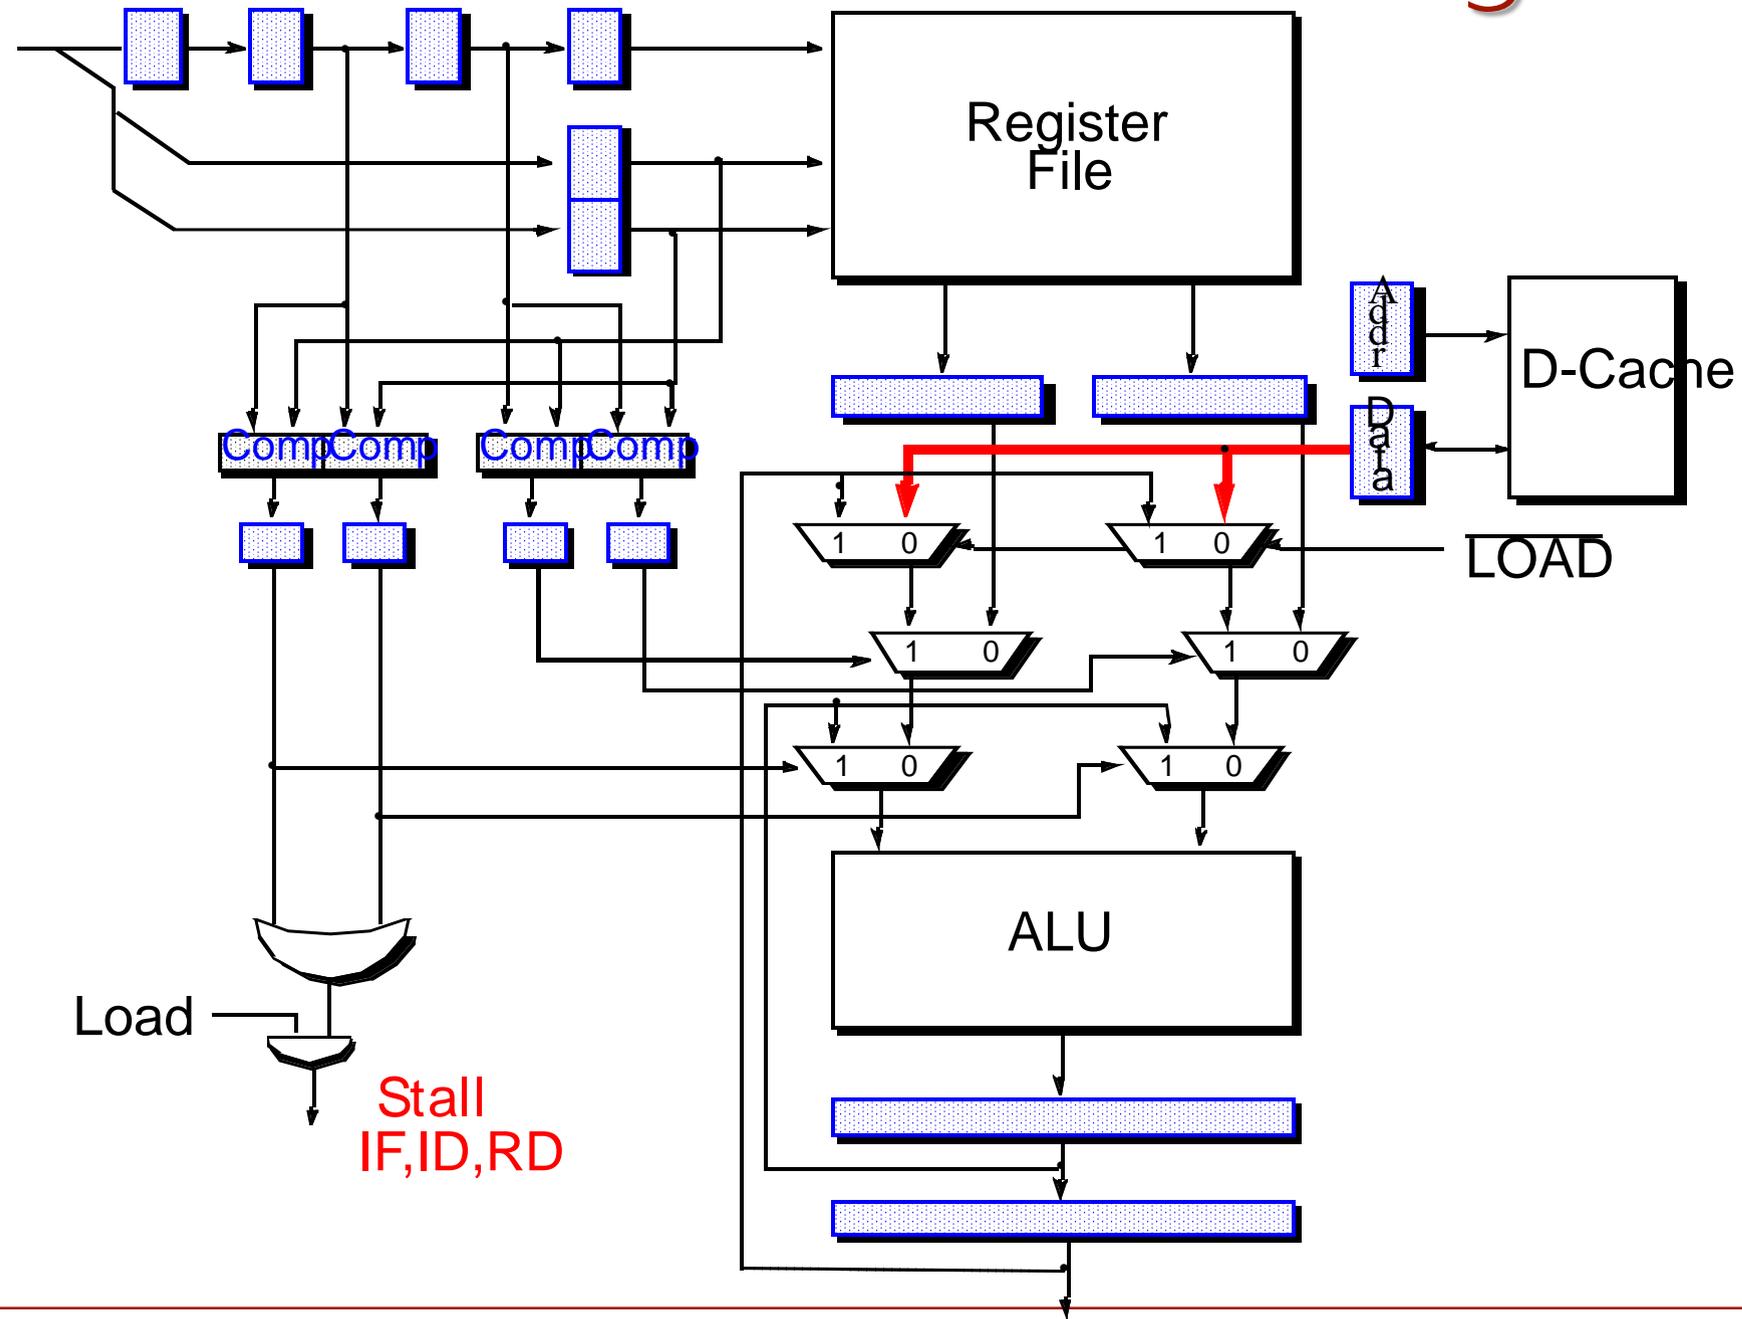
Implementation of ALU Forwarding



Load (Leading Instruction) Interlock and Penalty



Implementation of Load Forwarding

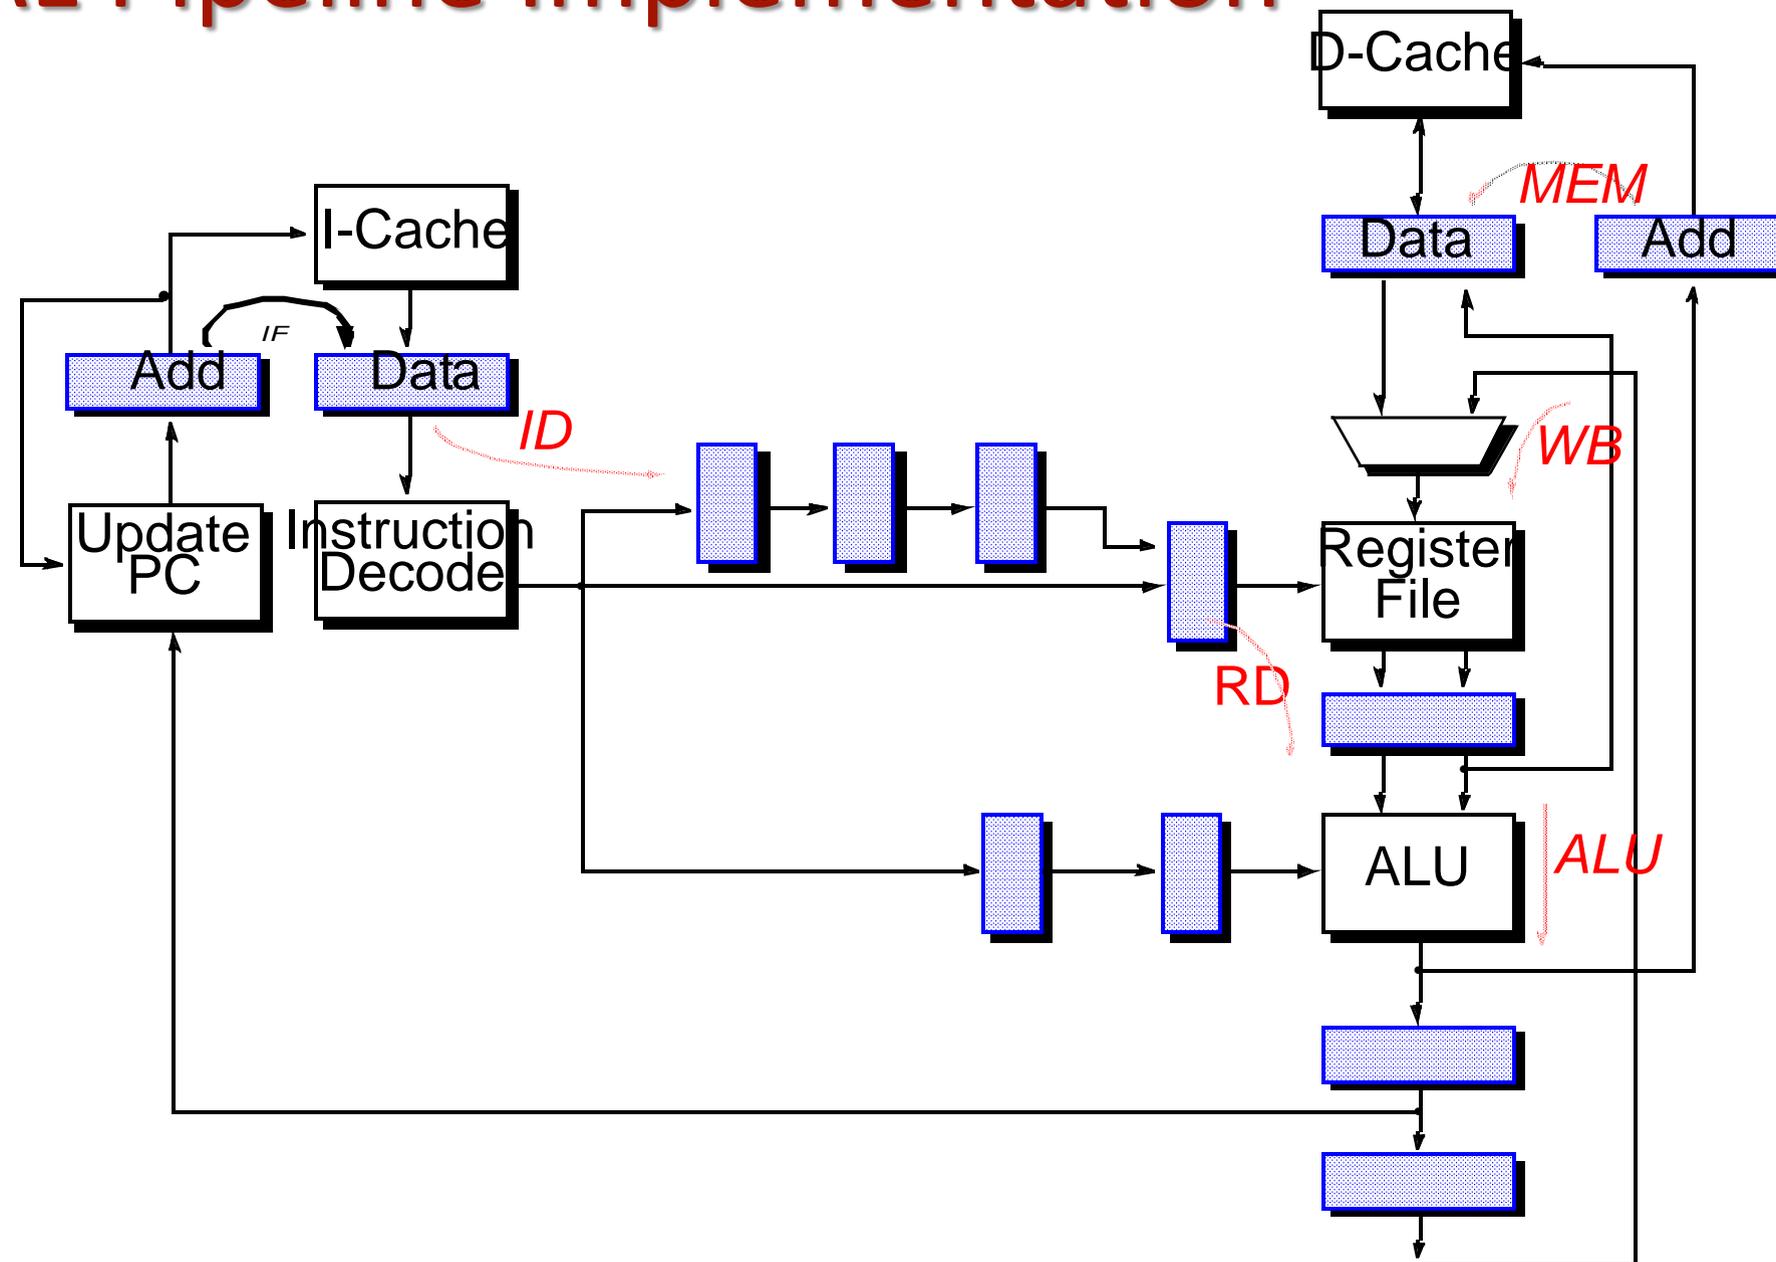


Branch Instruction Hazards

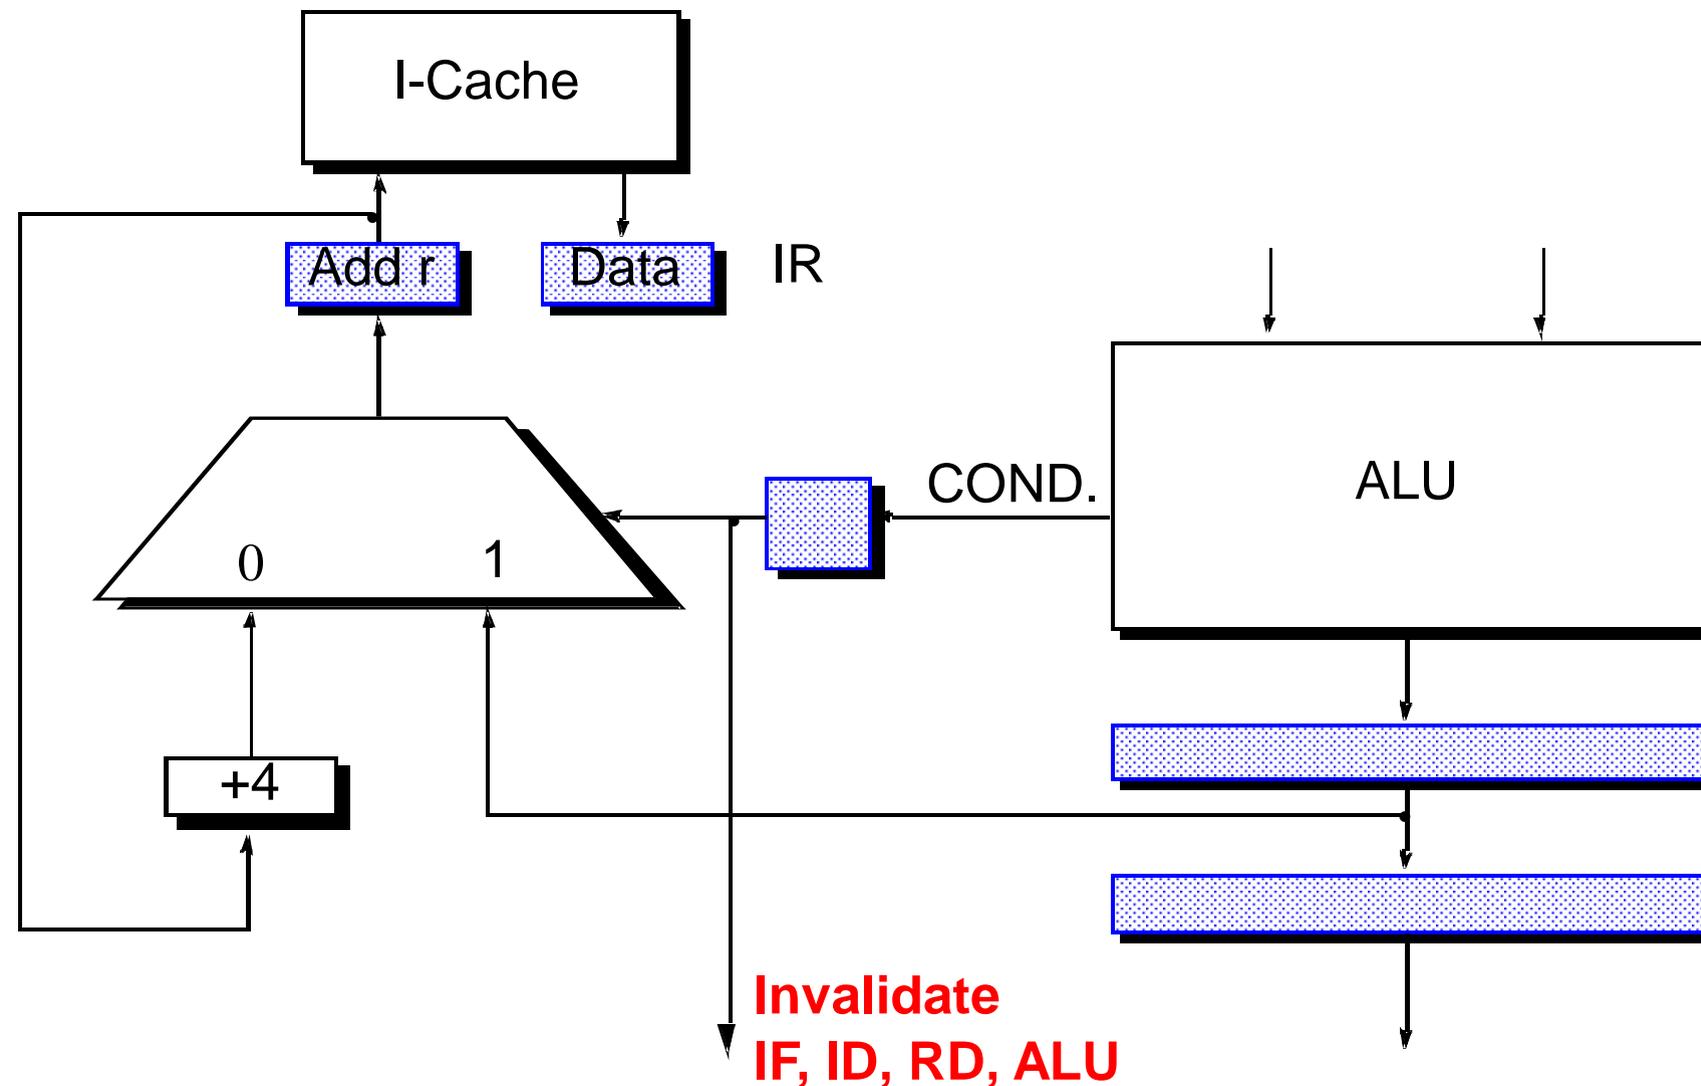
<u>Pipe Stage</u>	<u>ALU</u> Inst.	<u>Load</u> inst.	<u>Store</u> inst.	<u>Branch</u> inst.
1. IF	I-cache PC \leftarrow PC+4	I-cache PC \leftarrow PC+4	I-cache PC \leftarrow PC+4	I-cache PC \leftarrow PC+4
2. ID	decode	decode	decode	decode
3. RD	read reg.	read reg.	read reg.	read reg.
4. ALU	ALU op.	addr. gen.	addr. gen.	addr. gen. cond. gen.
5. MEM	-----	read mem.	write mem.	PC \leftarrow -br. addr.
6. WB	write reg.	write reg.	-----	-----

TYPICAL Pipeline Implementation

The 6-stage TYPICAL pipeline:



Implementation of Branch Instruction Interlock



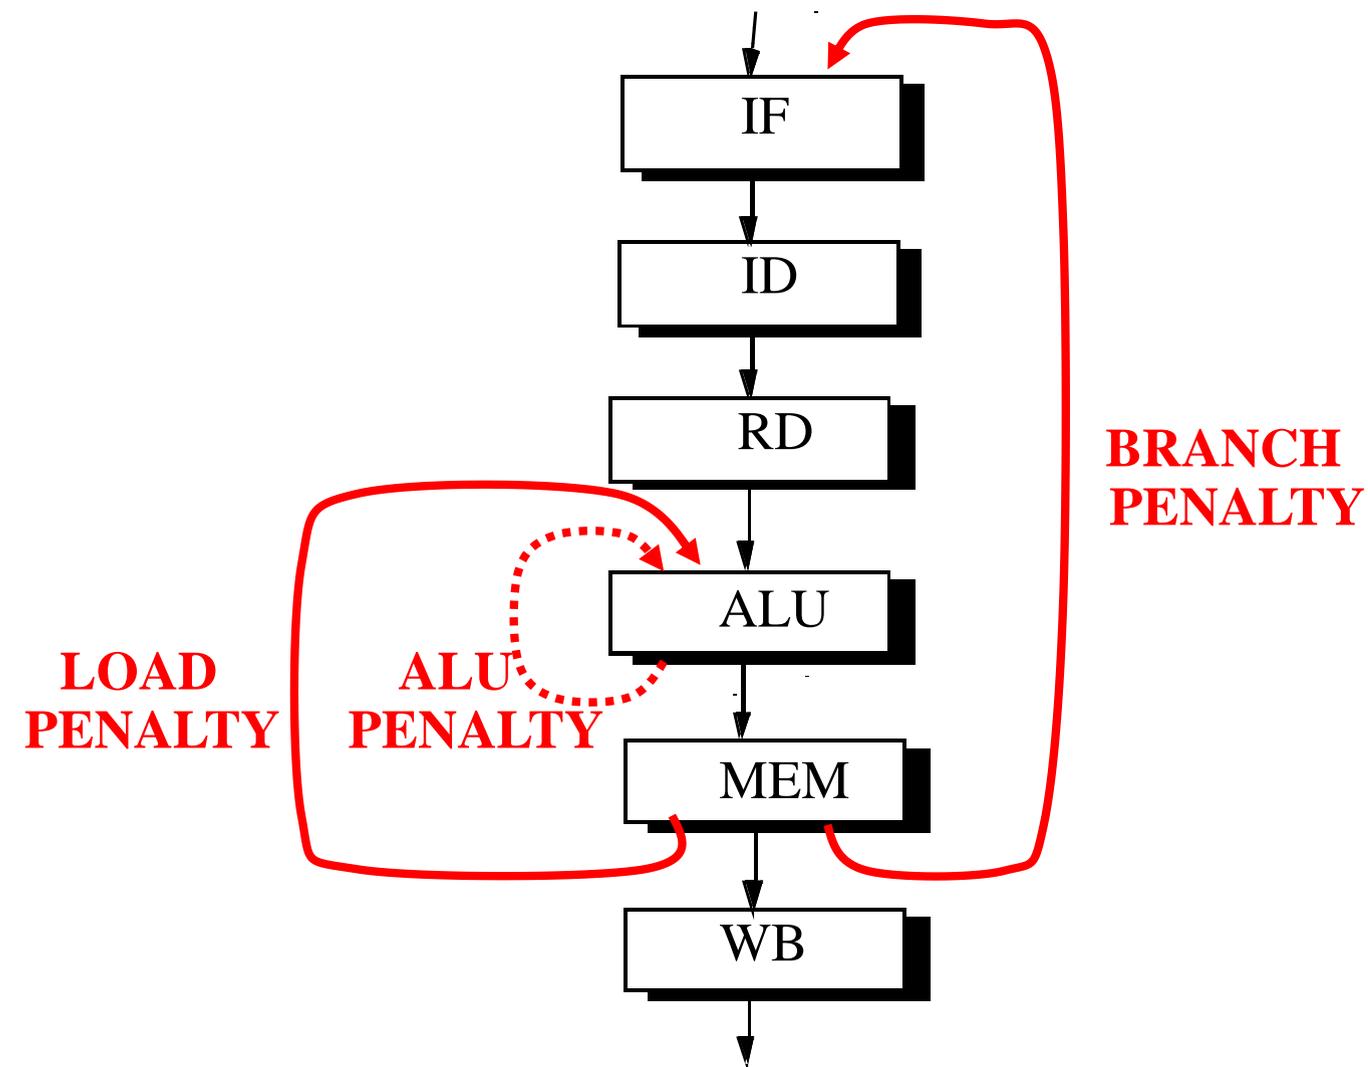
Penalties Due to Stalling for RAW

Leading Inst _i	ALU	Load	Branch
Trailing Inst _j	ALU, L/S, Br.	ALU, L/S, Br.	ALU, L/S, Br.
Hazard register	Int. Reg. (R _i)	Int. Reg. (R _i)	PC
Register WRITE stage (i)	WB (stage 6)	WB (stage 6)	MEM (stage 5)
Register READ stage (j)	RD (stage 3)	RD (stage 3)	IF (stage 1)
RAW distance or penalty:	3 cycles	3 cycles	4 cycles

Penalties with Forwarding Paths

Leading Inst _i	ALU	Load	Branch
Trailing Inst _j	ALU, L/S, Br.	ALU, L/S, Br.	ALU, L/S, Br.
Hazard register	Int. Reg. (R _i)	Int. Reg. (R _i)	PC
Register WRITE stage (i)	WB (stage 6)	WB (stage 6)	MEM (stage 5)
Register READ stage (j)	RD (stage 3)	RD (stage 3)	IF (stage 1)
Forward from outputs of:	ALU, MEM, WB	MEM, WB	MEM
Forward to input of:	ALU	ALU	IF
RAW distance or penalty:	0 cycles	1 cycles	4 cycles

3 Major Penalty Loops of (Scalar) Pipelining



Performance Objective: Reduce CPI as close to 1 as possible.

18-600 Foundations of Computer Systems

Lecture 9: "Pipelined Processor Design"

1. TYPICAL Pipelined Processor

- a. Instruction Pipeline Design
 - Balancing Pipeline Stages
 - Unifying Instruction Types
 - Resolving Pipeline Hazards

2. Y86-64 Pipelined Processor (PIPE) ← COVERED IN RECITATION

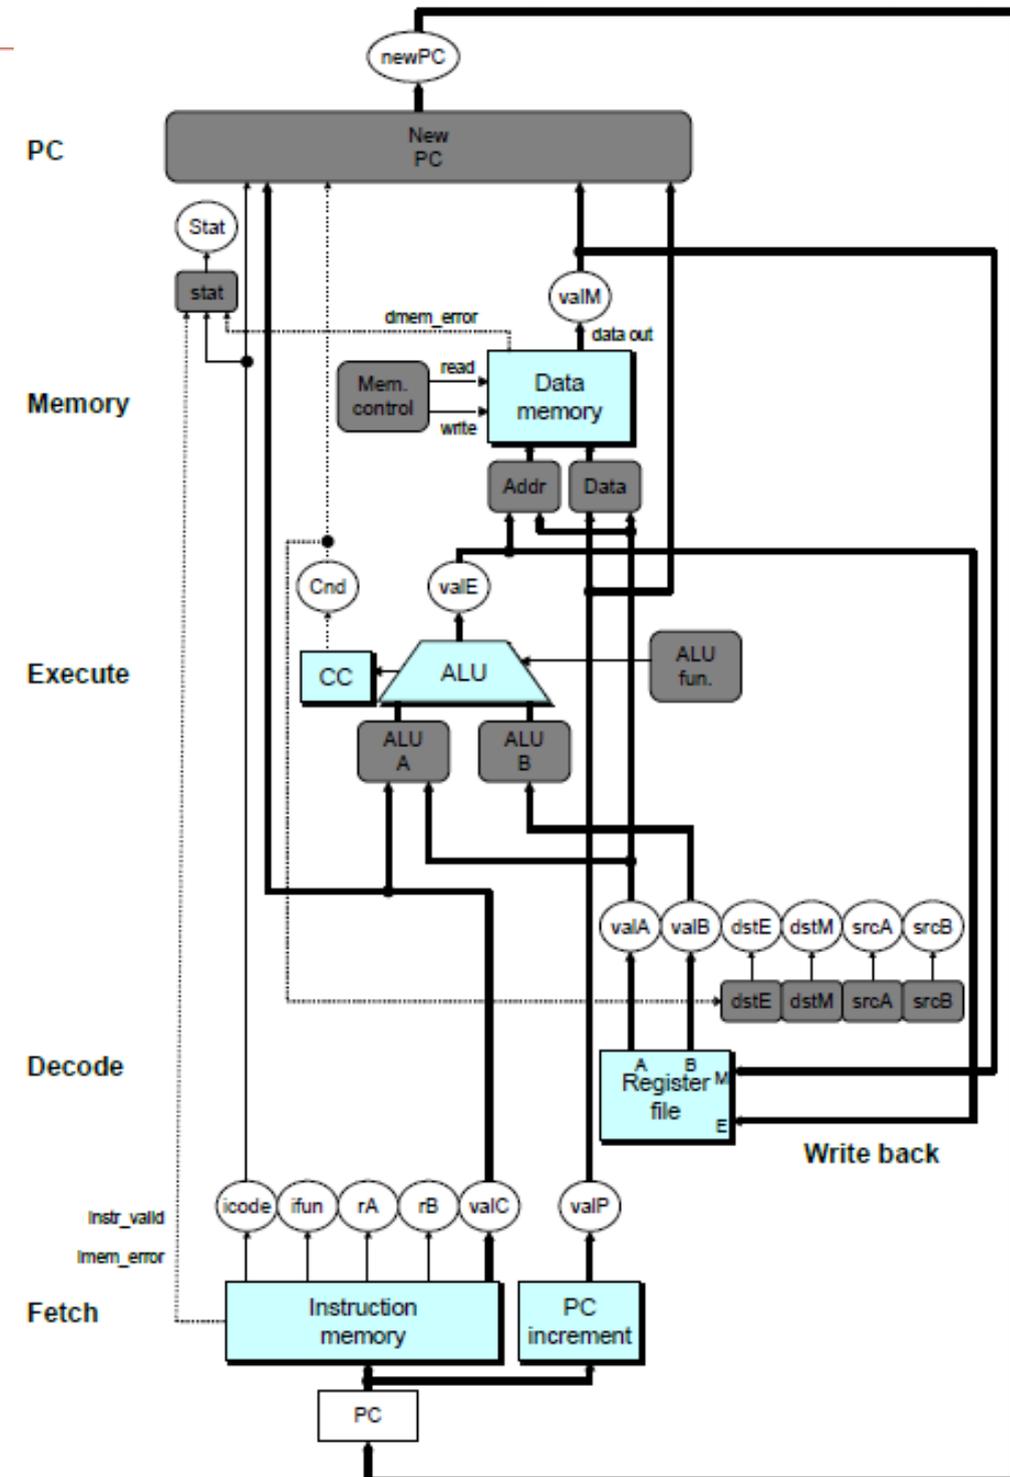
- a. Pipelining of the SEQ Processor
- b. Dealing with Data Hazards
- c. Dealing with Control Hazards

3. Motivation for Superscalar

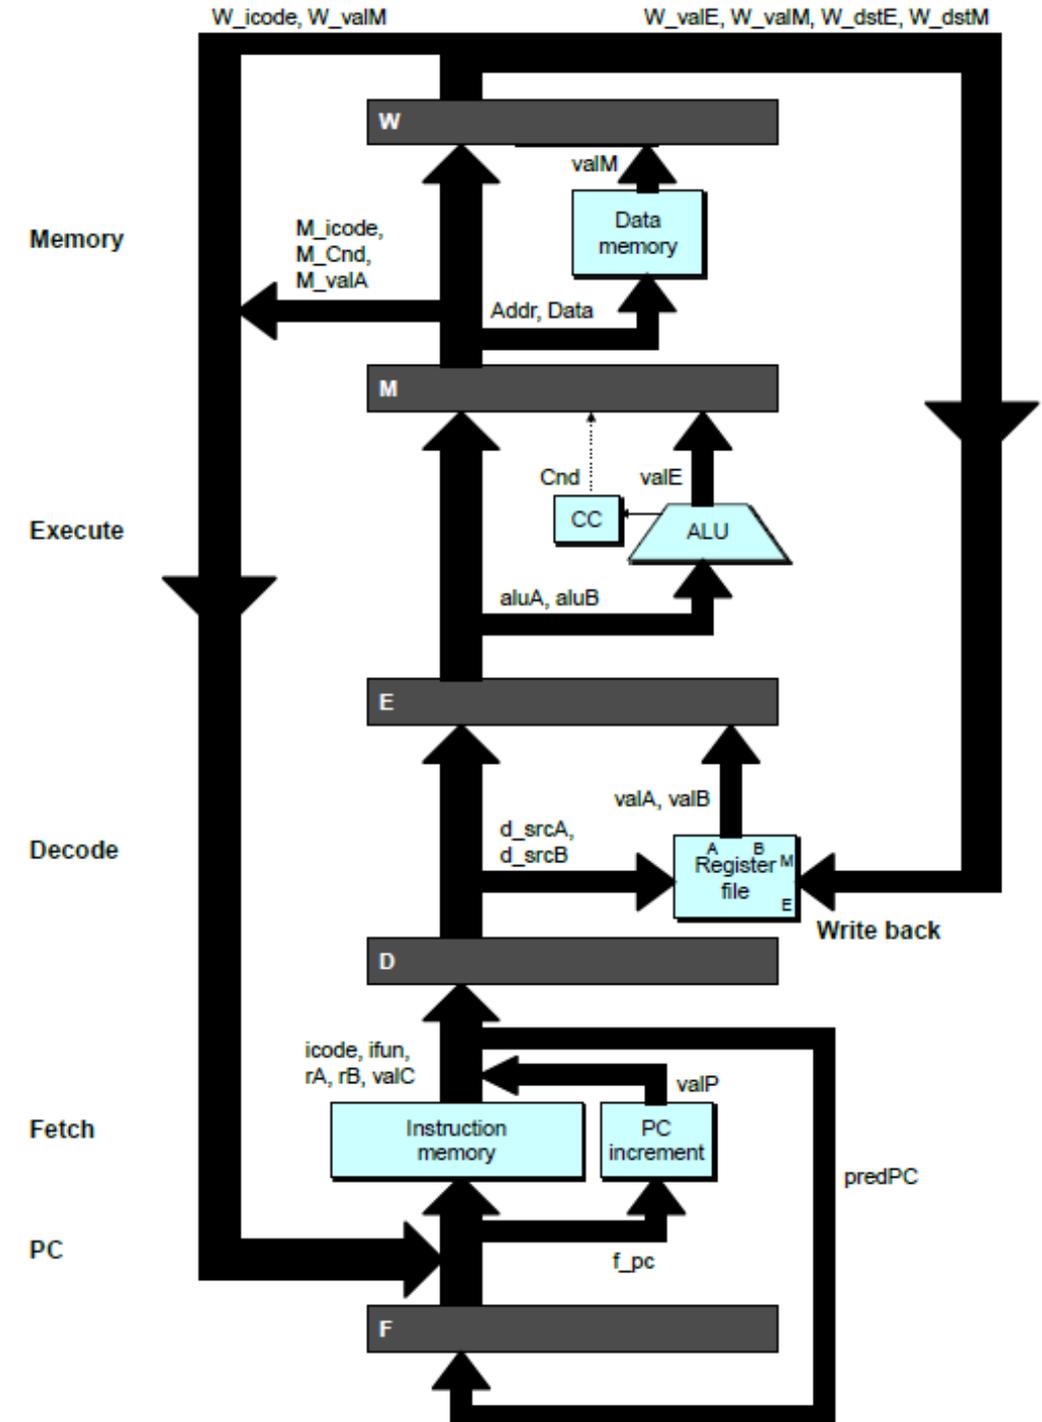
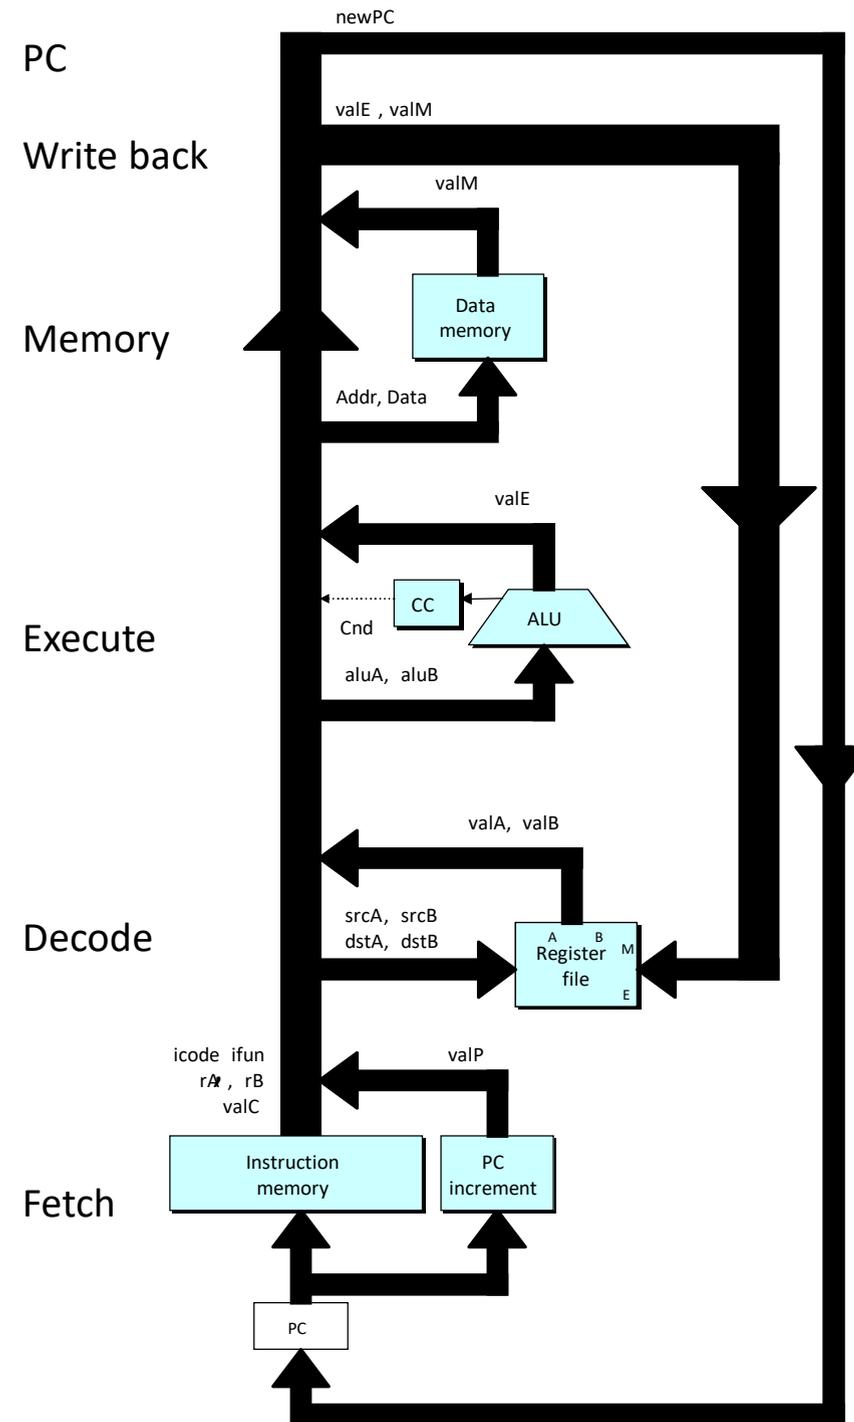


SEQ Processor Hardware

- Stages occur in sequence
- One operation in process at a time

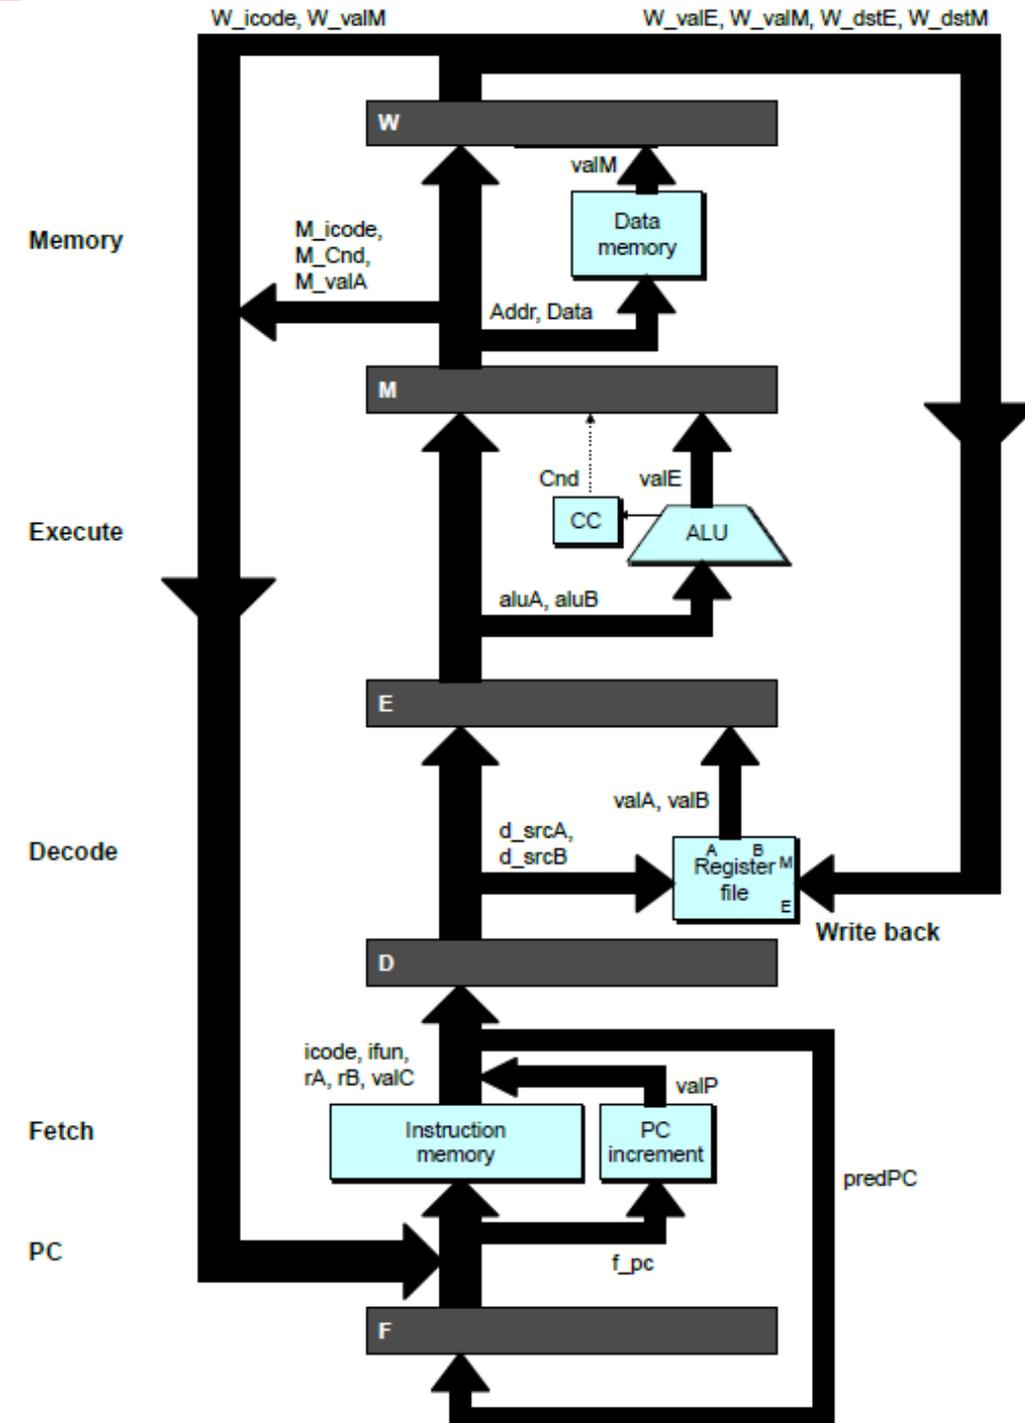


Adding Pipeline Registers



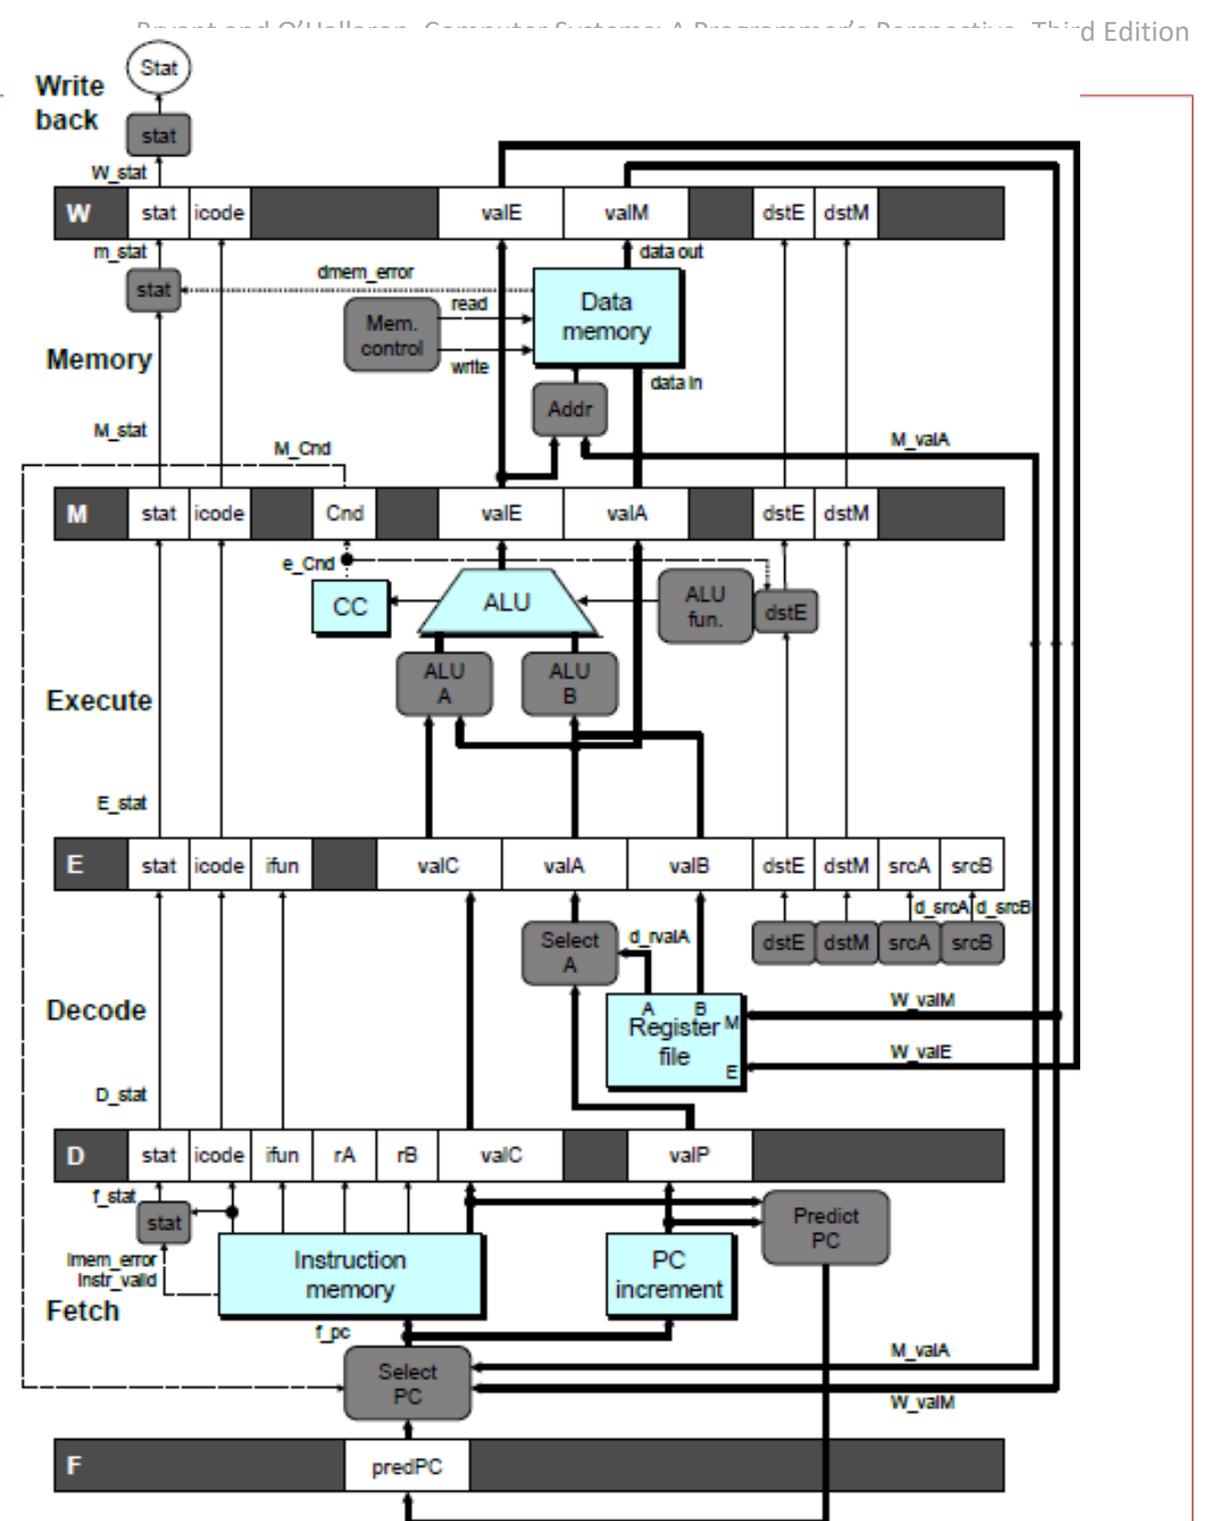
Pipeline Stages

- Fetch (F)
 - Select current PC
 - Read instruction
 - Compute incremented PC
- Decode (D)
 - Read program registers
- Execute (E)
 - Operate ALU
- Memory (M)
 - Read or write data memory
- Write Back (W)
 - Update register file



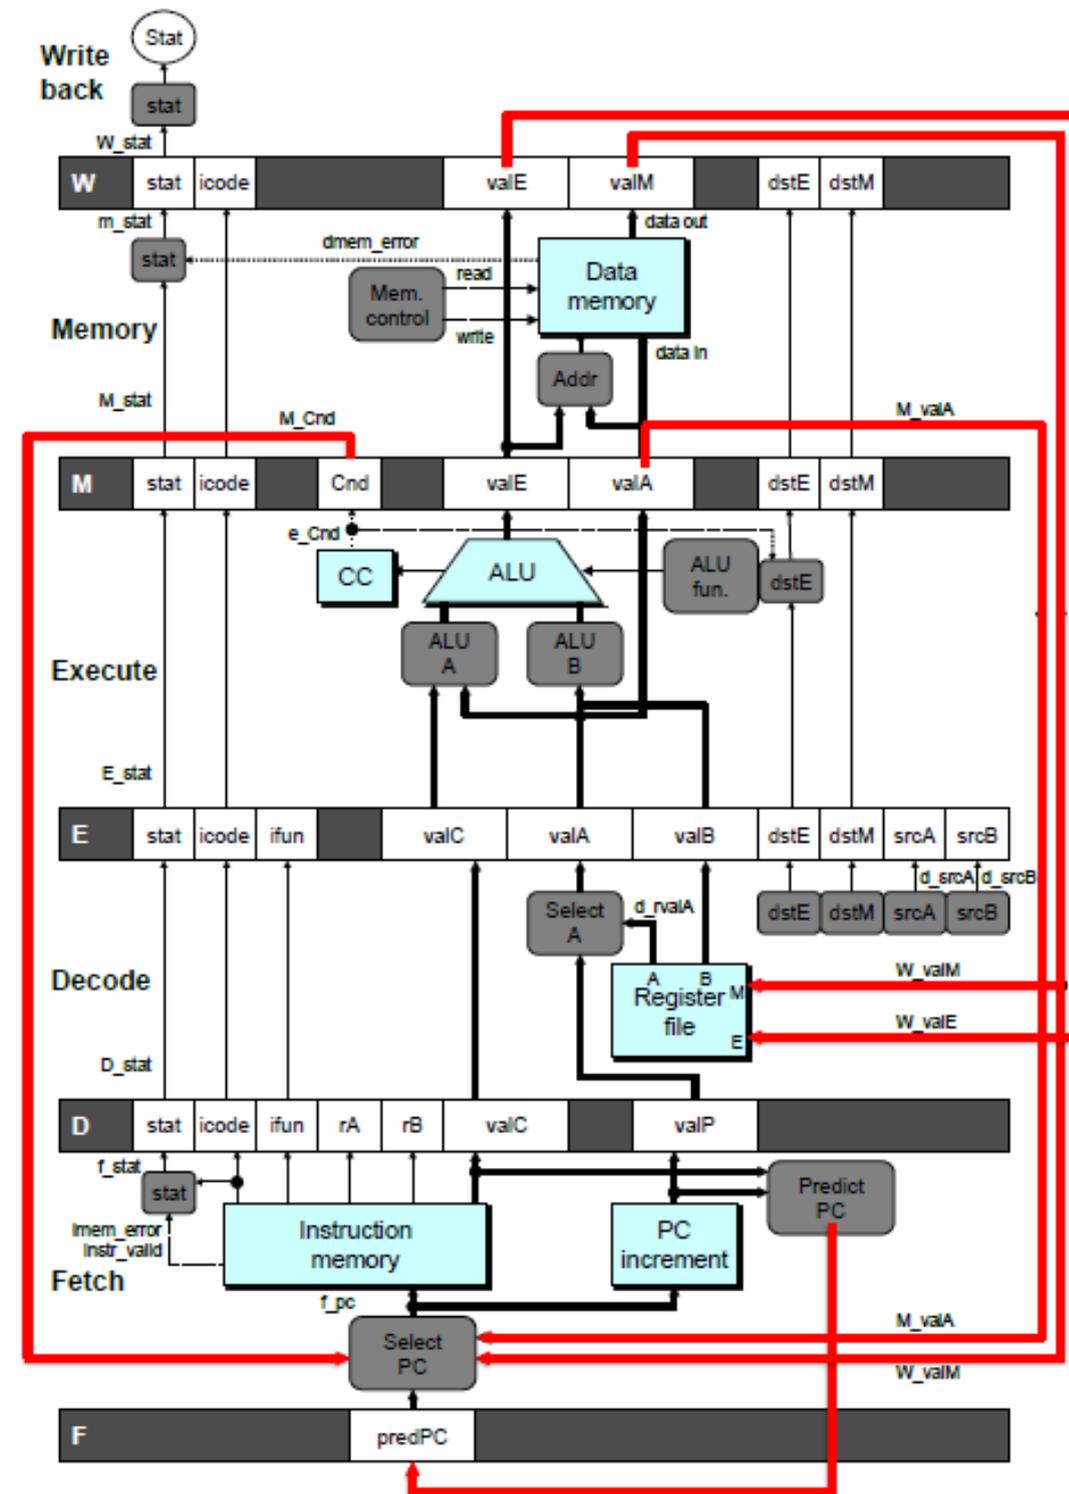
PIPE Processor Hardware

- Pipeline registers hold intermediate values from instruction execution
- Forward (Upward) Paths
 - Values passed from one stage to next
 - Cannot jump past stages
 - e.g., valC passes through decode



Feedback Paths

- Predicted PC
 - Guess value of next PC
- Branch information
 - Jump taken/not-taken
 - Fall-through or target address
- Return point
 - Read from memory
- Register updates
 - To register file write ports



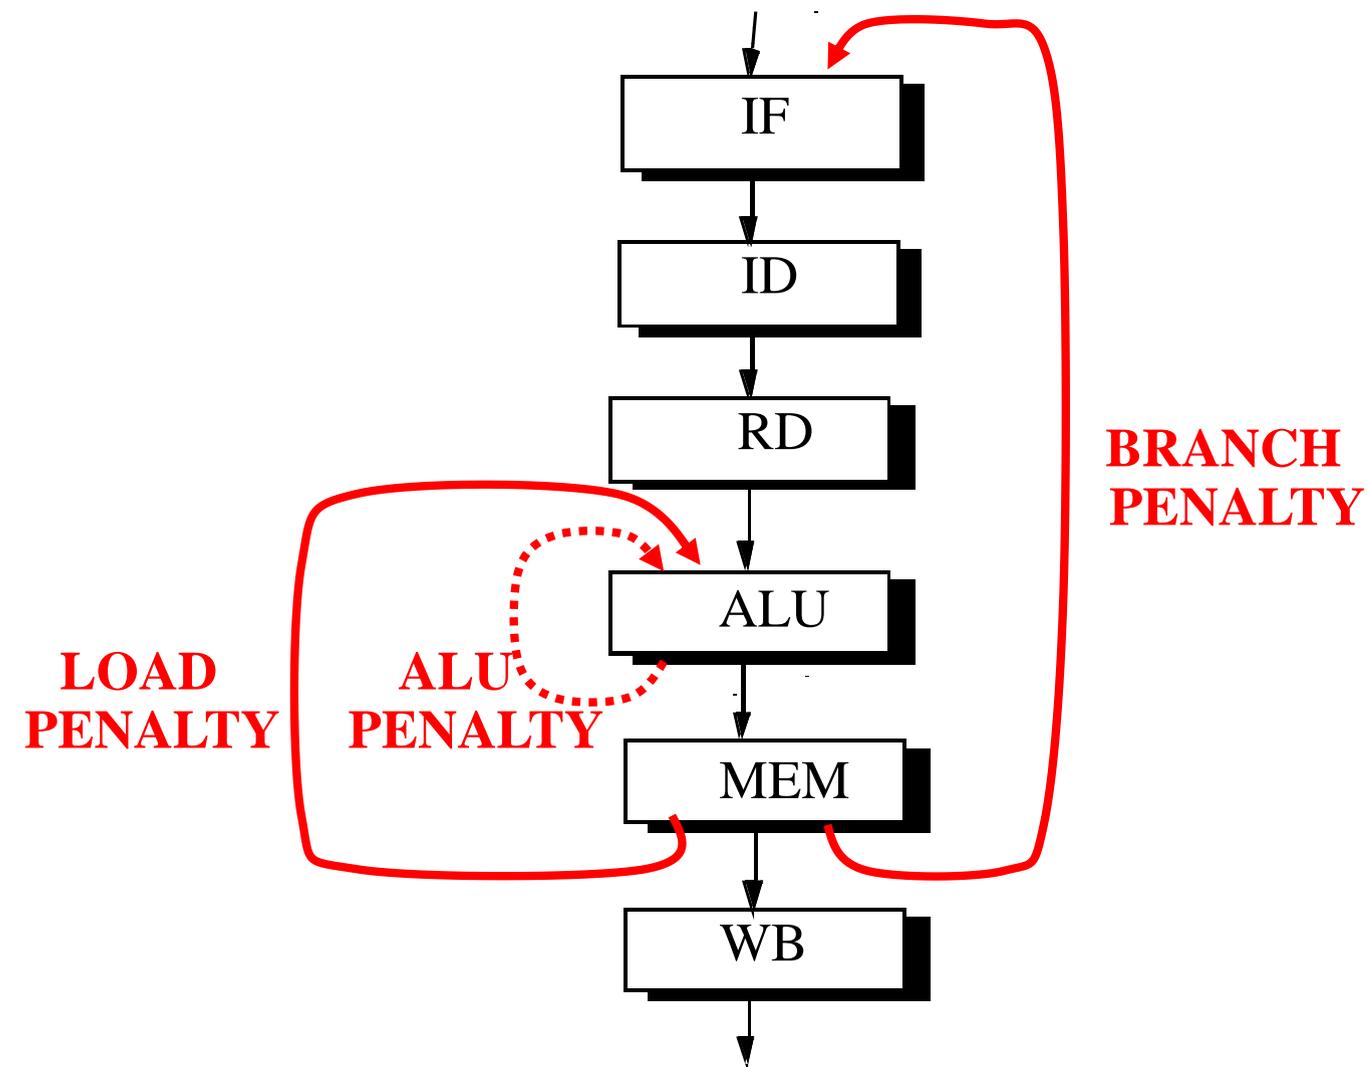
18-600 Foundations of Computer Systems

Lecture 9: "Pipelined Processor Design"

1. **TYPICAL Pipelined Processor**
 - a. Instruction Pipeline Design
 - Balancing Pipeline Stages
 - Unifying Instruction Types
 - Resolving Pipeline Hazards
2. **Y86-64 Pipelined Processor (PIPE)**
 - a. Pipelining of the SEQ Processor
 - b. Dealing with Data Hazards
 - c. Dealing with Control Hazards
3. **Motivation for Superscalar**



3 Major Penalty Loops of (Scalar) Pipelining



Performance Objective: Reduce CPI as close to 1 as possible.

MIPS R2000/R3000 Pipeline

Stage	Phase	Function performed
IF	ϕ_1	Translate virtual instr. addr. using TLB
	ϕ_2	Access I-cache
RD	ϕ_1	Return instruction from I-cache, check tags & parity
	ϕ_2	Read RF; if branch, generate target
ALU	ϕ_1	Start ALU op; if branch, check condition
	ϕ_2	Finish ALU op; if ld/st, translate addr
MEM	ϕ_1	Access D-cache
	ϕ_2	Return data from D-cache, check tags & parity
WB	ϕ_1	Write RF
	ϕ_2	

Separate
Adder

IBM RISC Experience [Agerwala and Cocke 1987]

- Internal IBM study: Limits of a scalar pipeline?
- Memory Bandwidth
 - Fetch 1 instr/cycle from I-cache
 - 40% of instructions are load/store (D-cache)
- Code characteristics (dynamic)
 - Loads – 25%
 - Stores 15%
 - ALU/RR – 40%
 - Branches – 20%
 - 1/3 unconditional (always taken)
 - 1/3 conditional taken, 1/3 conditional not taken

IBM Experience

➤ Cache Performance

- Assume 100% hit ratio (upper bound)
- Cache latency: $I = D = 1$ cycle default

➤ Load and branch scheduling

- Loads
 - 25% cannot be scheduled (delay slot empty)
 - 65% can be moved back 1 or 2 instructions
 - 10% can be moved back 1 instruction
- Branches
 - Unconditional – 100% schedulable (fill one delay slot)
 - Conditional – 50% schedulable (fill one delay slot)

CPI Optimizations

- Goal and impediments
 - CPI = 1, prevented by pipeline stalls
- No cache bypass of RF, no load/branch scheduling
 - Load penalty: 2 cycles: $0.25 \times 2 = 0.5$ CPI
 - Branch penalty: 2 cycles: $0.2 \times 2/3 \times 2 = 0.27$ CPI
 - Total CPI: $1 + 0.5 + 0.27 = 1.77$ CPI
- Bypass, no load/branch scheduling
 - Load penalty: 1 cycle: $0.25 \times 1 = 0.25$ CPI
 - Total CPI: $1 + 0.25 + 0.27 = 1.52$ CPI

More CPI Optimizations

- Bypass, scheduling of loads/branches
 - Load penalty:
 - $65\% + 10\% = 75\%$ moved back, no penalty
 - $25\% \Rightarrow 1$ cycle penalty
 - $0.25 \times 0.25 \times 1 = 0.0625$ CPI
 - Branch Penalty
 - $1/3$ unconditional 100% schedulable $\Rightarrow 1$ cycle
 - $1/3$ cond. not-taken, \Rightarrow no penalty (predict not-taken)
 - $1/3$ cond. Taken, 50% schedulable $\Rightarrow 1$ cycle
 - $1/3$ cond. Taken, 50% unschedulable $\Rightarrow 2$ cycles
 - $0.25 \times [1/3 \times 1 + 1/3 \times 0.5 \times 1 + 1/3 \times 0.5 \times 2] = 0.167$
- Total CPI: $1 + 0.063 + 0.167 = 1.23$ CPI

Simplify Branches

- Assume 90% can be PC-relative
 - No register indirect, no register access
 - Separate adder (like MIPS R3000)
 - Branch penalty reduced

15% Overhead
from program
dependencies

- Total CPI: $1 + 0.063 + 0.085 = 1.15$ CPI = 0.87 IPC

PC-relative	Schedulable	Penalty
Yes (90%)	Yes (50%)	0 cycle
Yes (90%)	No (50%)	1 cycle
No (10%)	Yes (50%)	1 cycle
No (10%)	No (50%)	2 cycles

Limits of Pipelining

➤ IBM RISC Experience

- Control and data dependencies add 15%
- Best case CPI of 1.15, IPC of 0.87
- Deeper pipelines (higher frequency) magnify dependency penalties

➤ This analysis assumes 100% cache hit rates

- Hit rates approach 100% for some programs
- Many important programs have much worse hit rates

18-600 Foundations of Computer Systems

Lecture 10: "From Pipelined to Superscalar Processors"

John P. Shen & Zhiyi Yu
October 3, 2016

Next Time ...

- Required Reading Assignment:
 - Chapter 4 of CS:APP (3rd edition) by Randy Bryant & Dave O'Hallaron.
- Recommended Reading Assignment:
 - ❖ Chapter 4 of Shen and Lipasti (SnL).

