CS 110
Computer Architecture

Superscalar CPUs

Instructor:
Sören Schwertfeger

https://robotics.shanghaitech.edu.cn/courses/ca

School of Information Science and Technology SIST

ShanghaiTech University

Slides based on UC Berkley's CS61C

Agenda

- Control Hazards
- Processor Performance
- Complex Pipelines
  - Static Multiple Issues (VLIW)
  - Dynamic Multiple Issues (Superscalar)

Pipelined RISC-V RV32I Datapath

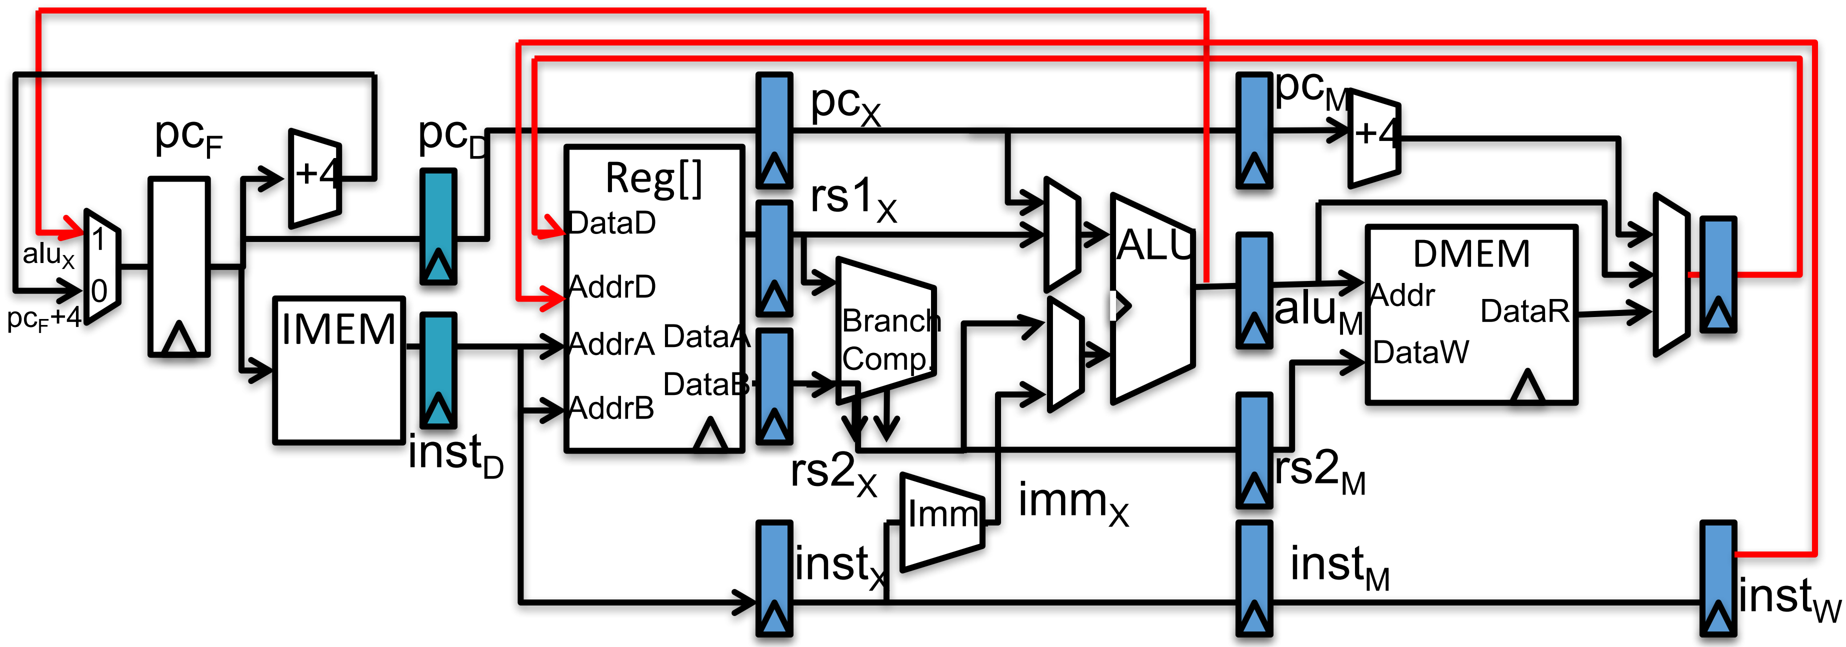

Pipelining Hazards

A *hazard* is a situation that prevents starting the next instruction in the next clock cycle

1) Structural hazard

- A required resource is busy
  (e.g. needed in multiple stages)

2) Data hazard

- Data dependency between instructions
- Need to wait for previous instruction to complete its data read/write

3) Control hazard

- Flow of execution depends on previous instruction

Structural Hazards: More Hardware Instruction and Data Caches

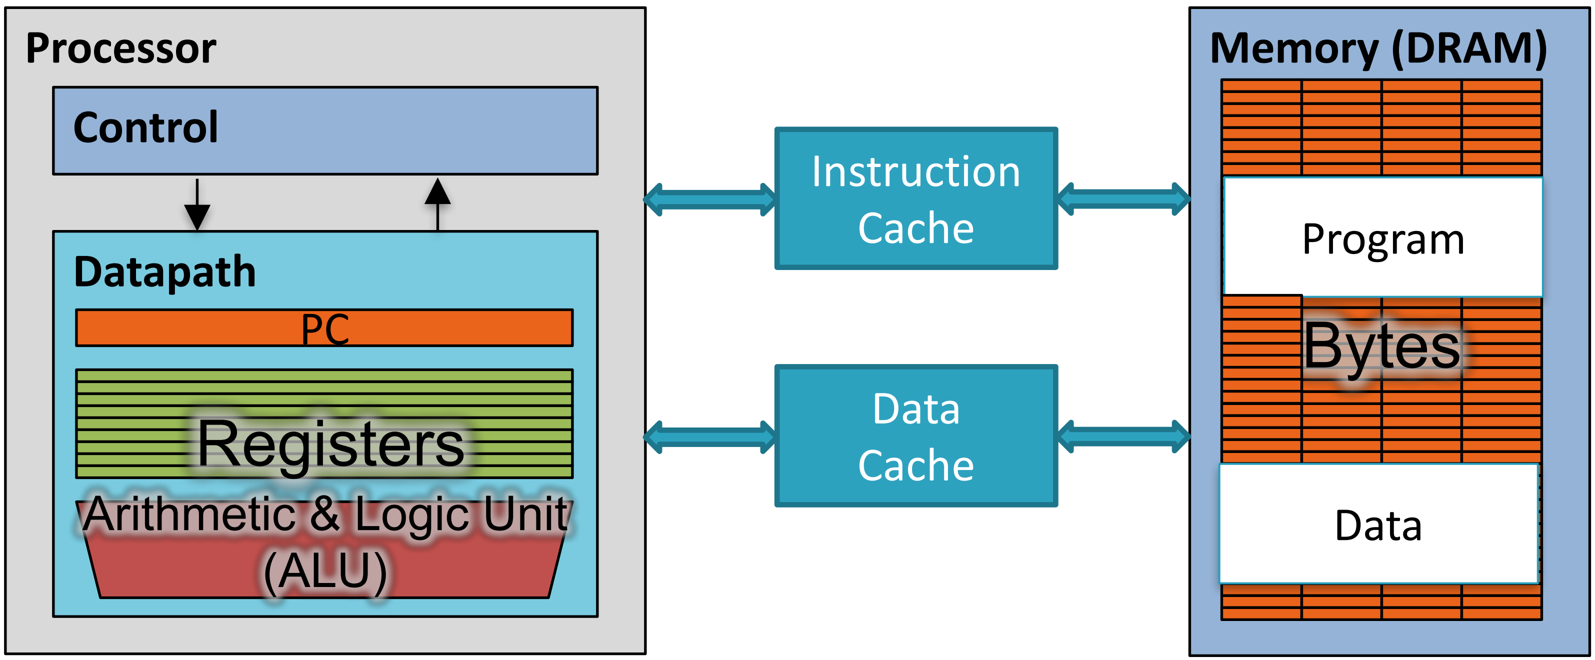

Data Hazards: Forwarding

Value of t0									
	5	5	5	5	5/9	9	9	9	9

instruction sequence

- add t0, t1, t2
- or t3, t0, t5
- sub t6, t0, t3
- xor t5, t1, t0
- sw t0, 8(t3)

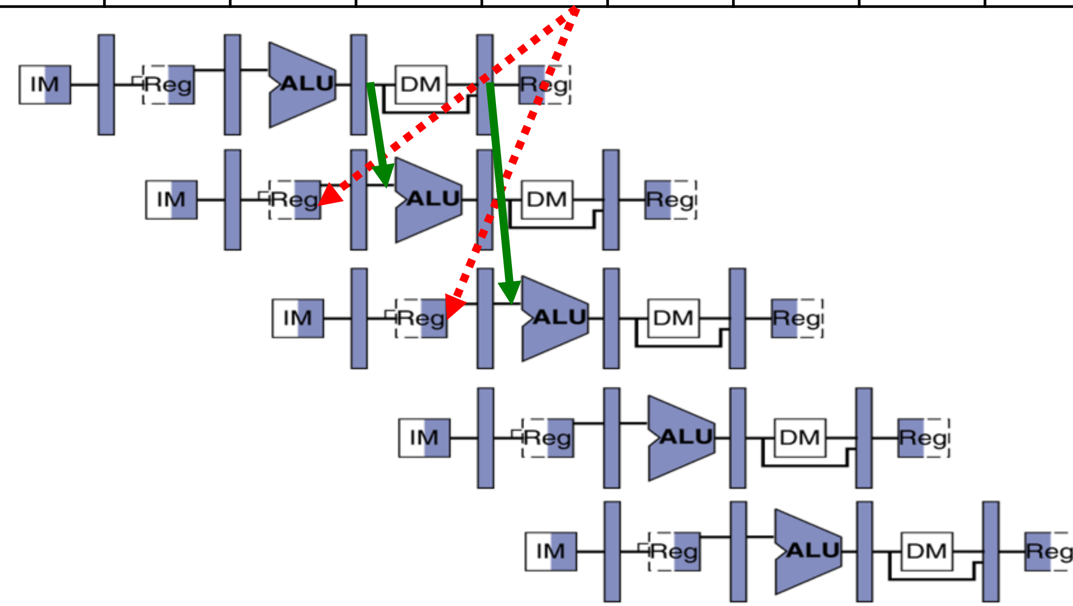

Forwarding: grab operand from pipeline stage, rather than register file

Forwarding Path

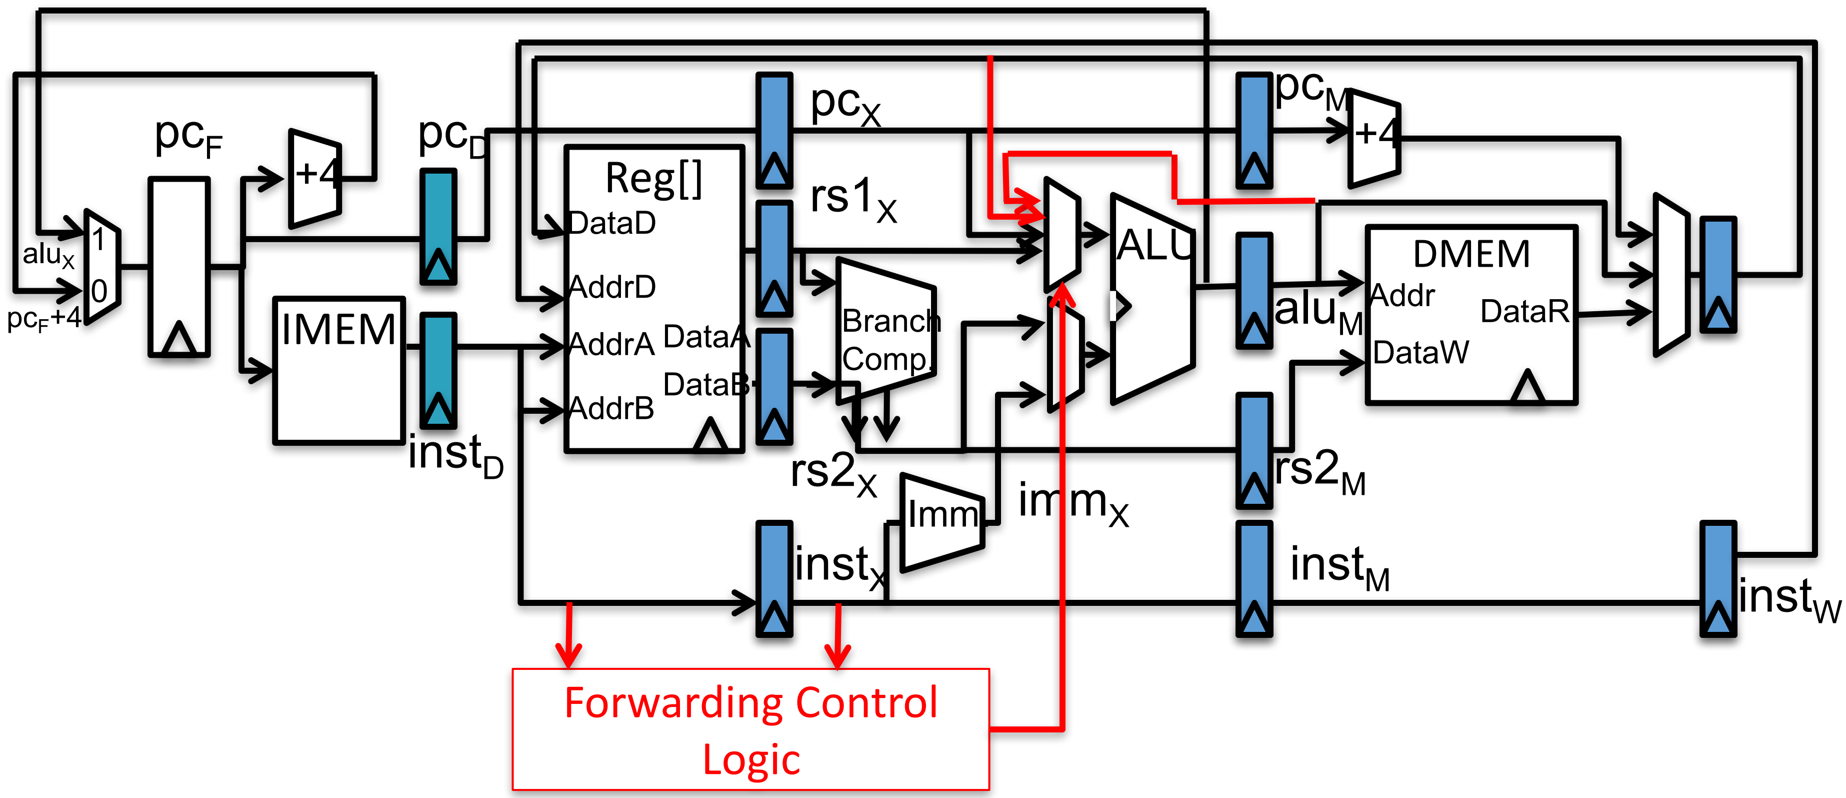

Load Data Hazard

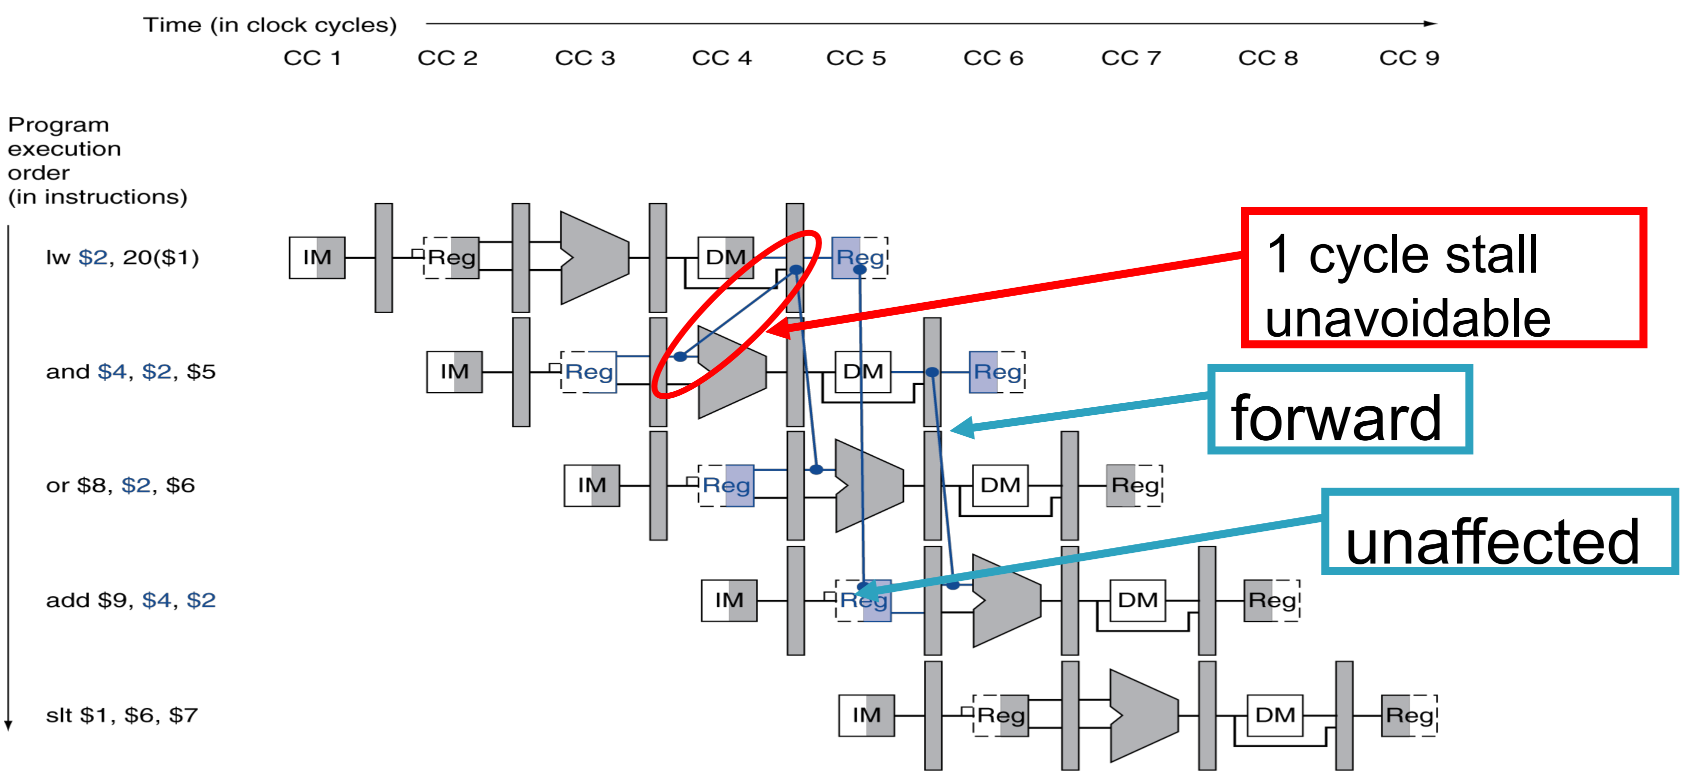

`lw` Data Hazard

- Slot after a load is called a *load delay slot*
  - If that instruction uses the result of the load, then the hardware will stall for one cycle
  - Equivalent to inserting an explicit **`nop`** in the slot
    - except the latter uses more code space
  - Performance loss
- Idea:
  - Put unrelated instruction into load delay slot
  - No performance loss!

Control Hazards

beq t0, t1, label

sub t2, s0, t5

or t6, s0, t3

xor t5, t1, s0

sw s0, 8(t3)

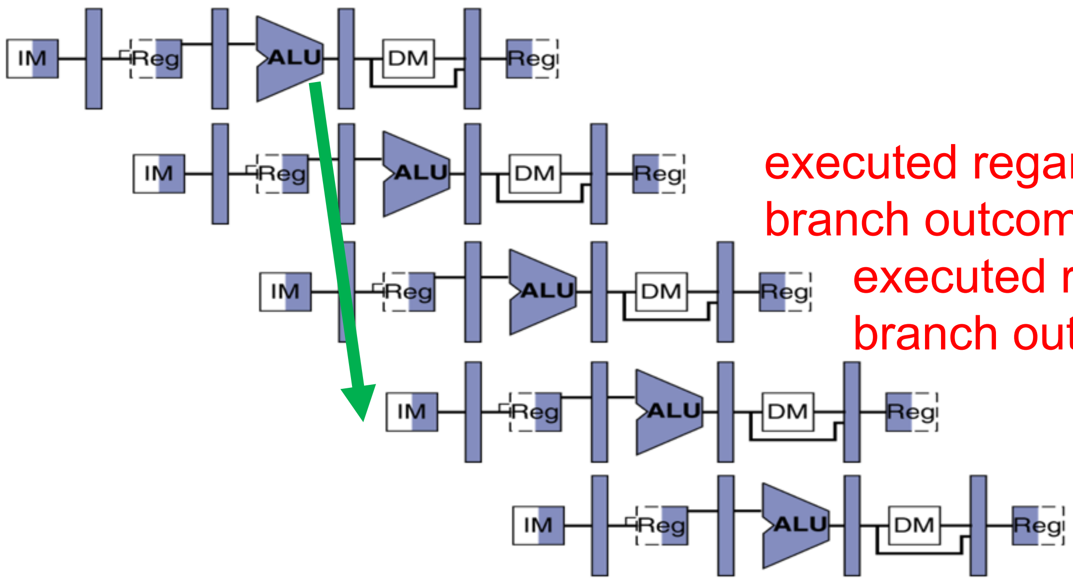

executed regardless of branch outcome!

executed regardless of branch outcome!!!

PC updated reflecting branch outcome

Observation

- If branch not taken, then instructions fetched sequentially after branch are correct
- If branch or jump taken, then need to flush incorrect instructions from pipeline by converting to NOPs

Kill Instructions after Branch if Taken

beq t0, t1, label

sub t2, s0, t5

or t6, s0, t3

label: xxxxxx

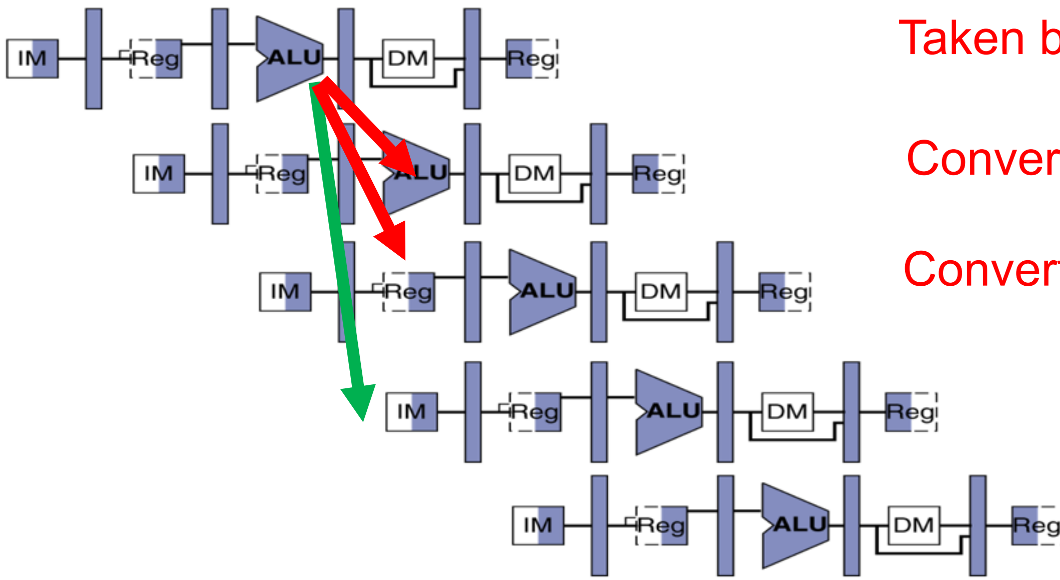

Taken branch

Convert to NOP

Convert to NOP

PC updated reflecting branch outcome

Reducing Branch Penalties

- Every taken branch in simple pipeline costs 2 dead cycles
- To improve performance, use “branch prediction” to guess which way branch will go earlier in pipeline
- Only flush pipeline if branch prediction was incorrect

Branch Prediction

beq t0, t1, label

label: …..

…..

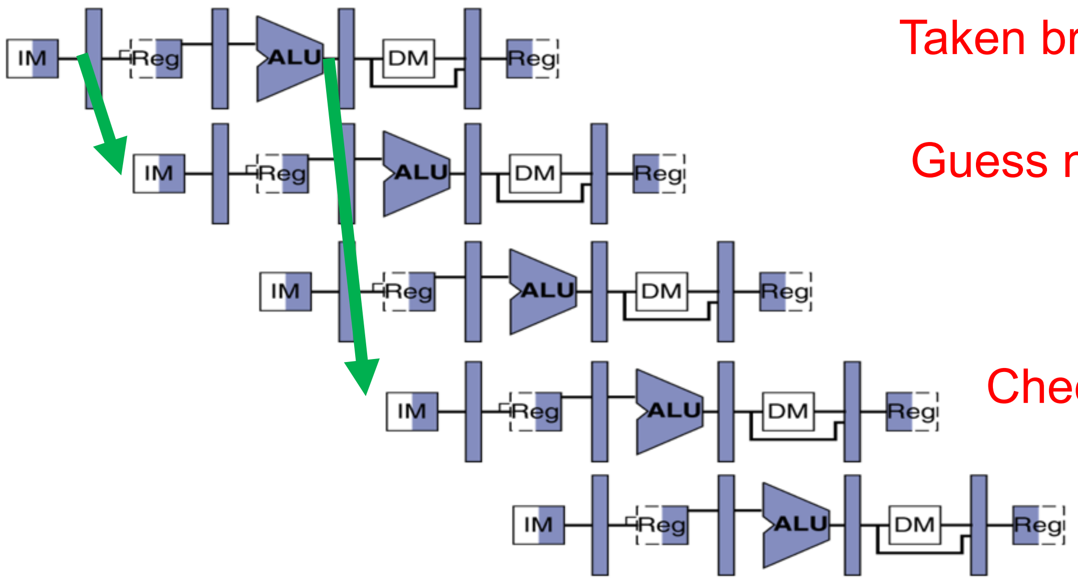

Taken branch

Guess next PC!

Check guess correct

Agenda

- Control Hazards
- Processor Performance
- Complex Pipelines
  - Static Multiple Issues (VLIW)
  - Dynamic Multiple Issues (Superscalar)

Increasing Processor Performance

1. Clock rate
   – Limited by technology and power dissipation
2. Pipelining
   – “Overlap” instruction execution
   – Deeper pipeline: 5 => 10 => 15 stages
      • Less work per stage → shorter clock cycle
      • But more potential for hazards
      • Multi-issue “superscalar” processor

Greater Instruction-Level Parallelism (ILP)

- Multiple issue "superscalar"
  - Replicate pipeline stages => multiple pipelines
  - Start multiple instructions per clock cycle
  - CPI < 1, so use Instructions Per Cycle (IPC)
  - E.g., 4GHz 4-way multiple-issue
    - 16 BIPS, peak CPI = 0.25, peak IPC = 4
  - But dependencies reduce this in practice
- "Out-of-Order" execution
  - Reorder instructions dynamically in hardware to reduce impact of hazards
- Hyper-threading

Hyper-threading (simplified)

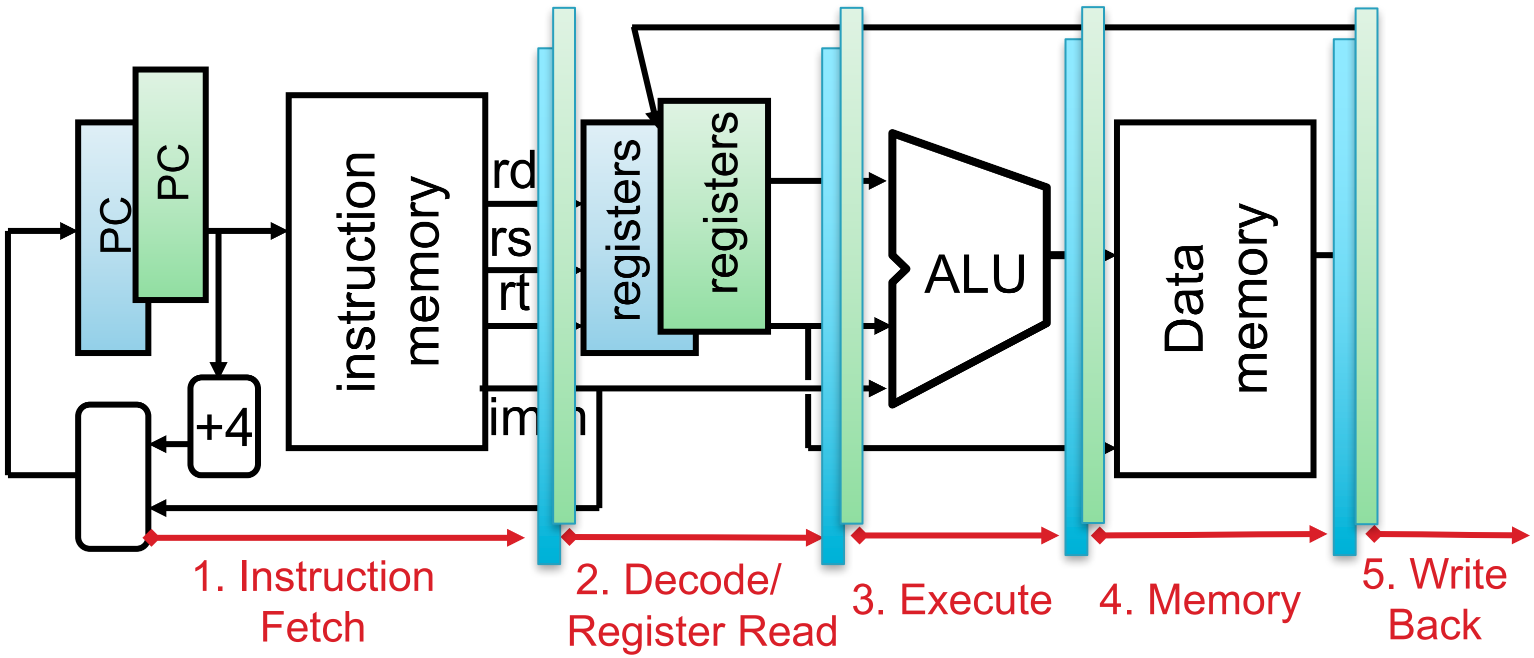

- Duplicate all elements that hold the state (registers)
- Use the same CL blocks
- Use muxes to select which state to use every clock cycle
- => run 2 independent processes
  - No Hazards: registers different; different control flow; memory different; Threads: memory hazard should be solved by software (locking, mutex, …)
- Speedup?
  - No obvious speedup; Complex pipeline: make use of CL blocks in case of unavailable resources (e.g. wait for memory)

Intel Nehalem i7

- Hyperthreading:
  - About 5% die area
  - Up to 30% speed gain (BUT also < 0% possible)
- Pipeline: 20-24 stages!
- Out-of-order execution
  1. Instruction fetch.
  2. Instruction dispatch to an instruction queue
  3. Instruction: Wait in queue until input operands are available => instruction can **leave queue before earlier**, older instructions.
  4. The instruction is issued to the appropriate functional unit and executed by that unit.
  5. The results are queued.
  6. Write to register only after all older instructions have their results written.

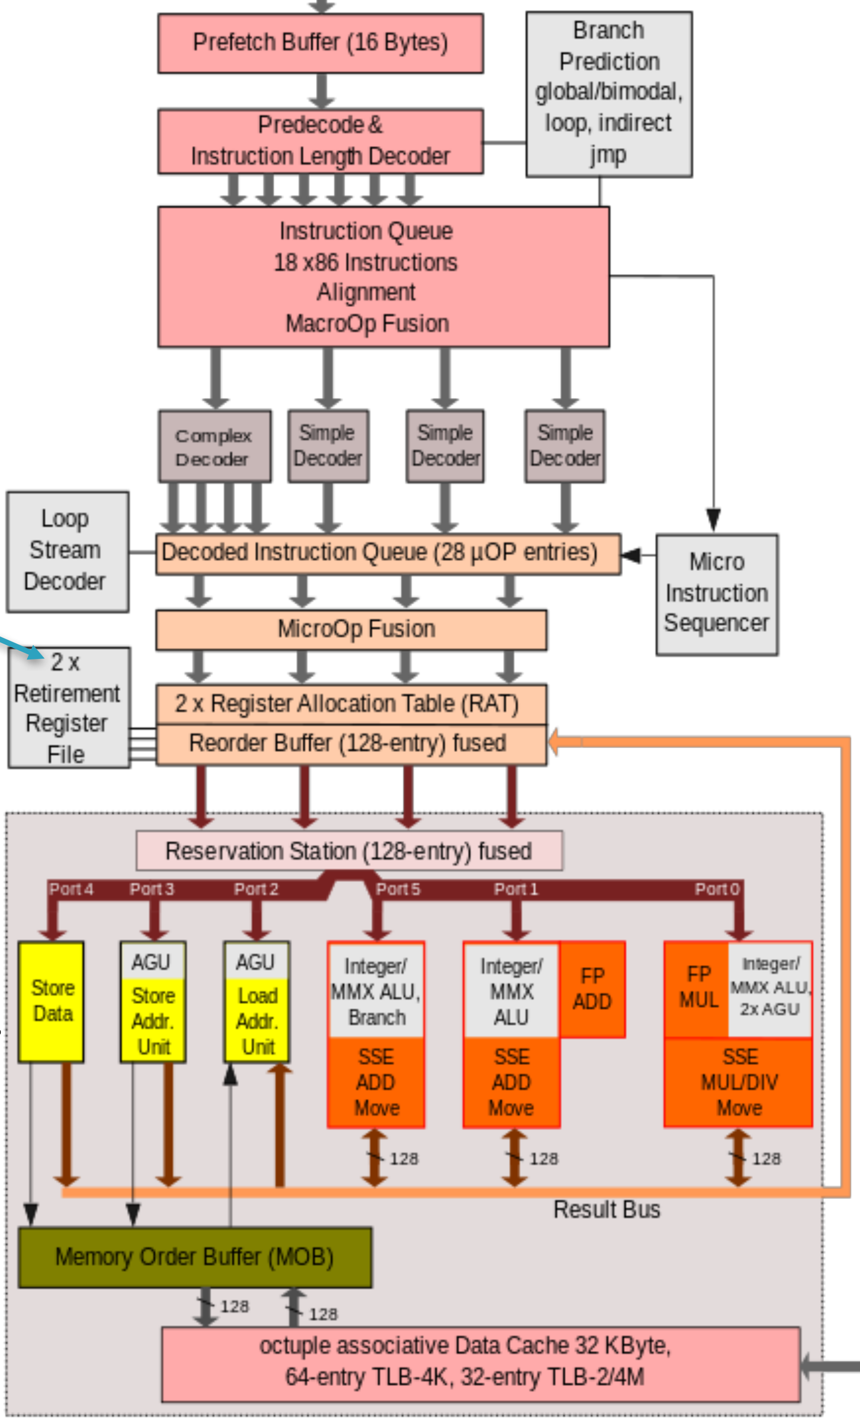

Superscalar Processor

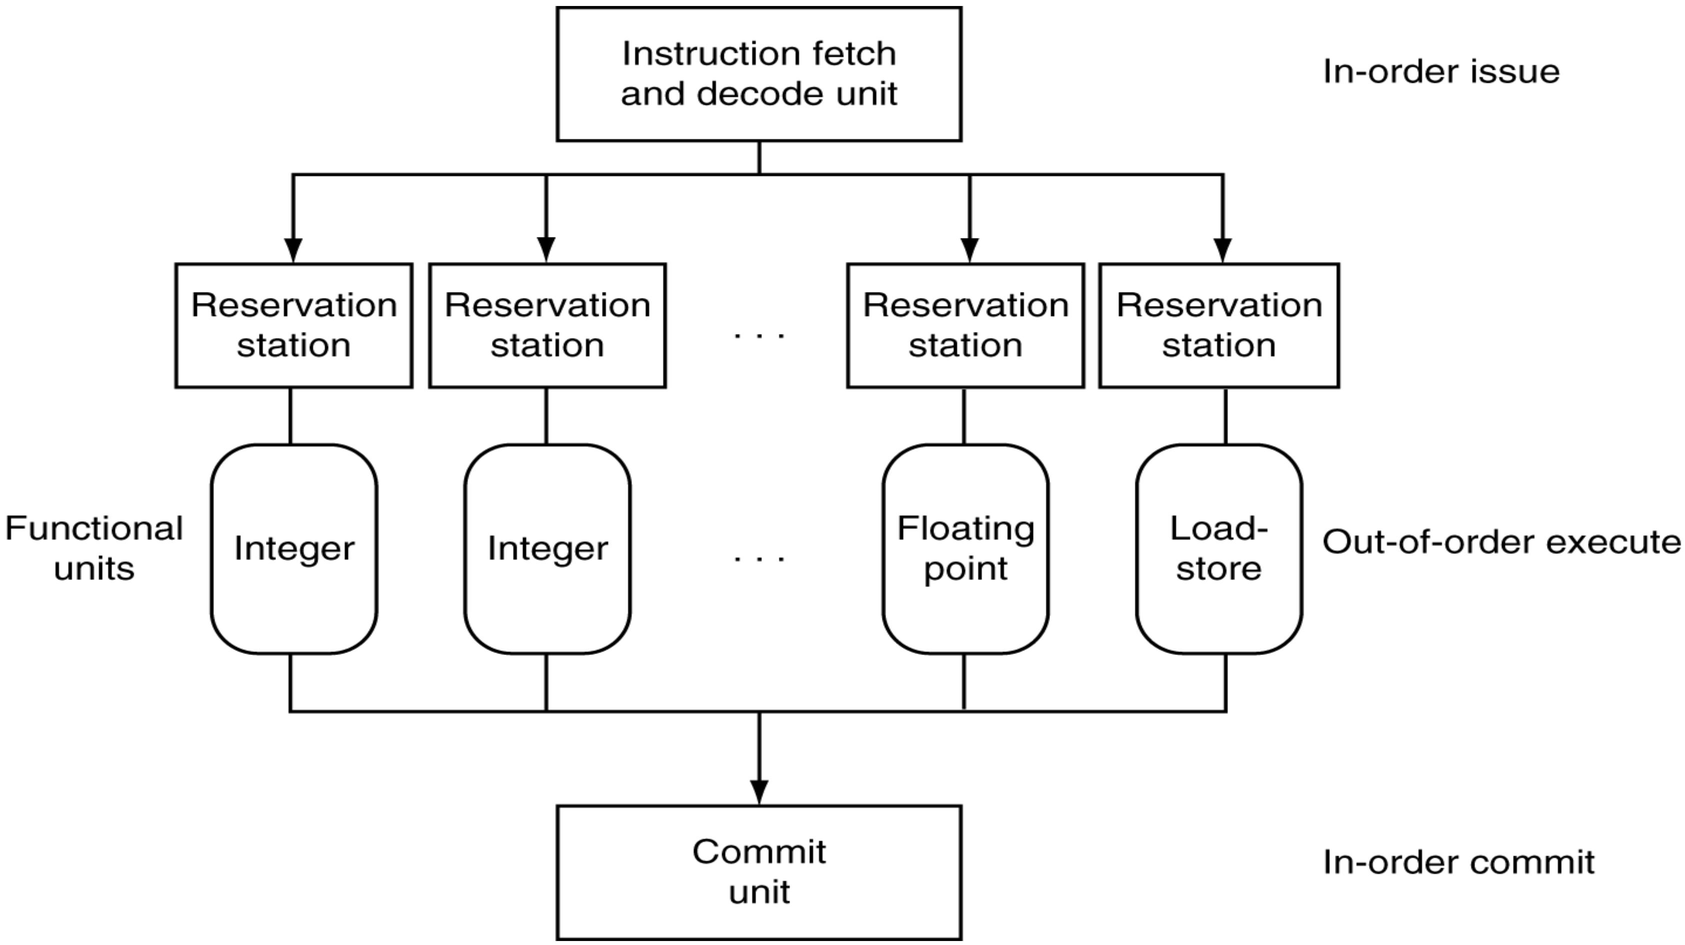

Superscalar = Multicore?

https://en.wikipedia.org/wiki/Superscalar_processor

- A superscalar processor is a CPU that implements a form of parallelism called instruction-level parallelism within a single processor. In contrast to a scalar processor that can execute at most one single instruction per clock cycle, a superscalar processor can execute more than one instruction during a clock cycle by simultaneously dispatching multiple instructions to different execution units on the processor. It therefore allows for more throughput (the number of instructions that can be executed in a unit of time) than would otherwise be possible at a given clock rate. **Each execution unit is not a separate processor (or a core if the processor is a multi-core processor), but an execution resource within a single CPU such as an arithmetic logic unit.**

- In Flynn's taxonomy, **a single-core superscalar processor is classified as an SISD** processor (Single Instruction stream, Single Data stream), though many superscalar processors support short vector operations and so could be classified as SIMD (Single Instruction stream, Multiple Data streams). A multi-core superscalar processor is classified as an MIMD processor (Multiple Instruction streams, Multiple Data streams).

"Iron Law" of Processor Performance

CPI = Cycles Per Instruction

Can time | **Can count** | **Can look up**

$$\frac{\text{Time}}{\text{Program}} = \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Cycles}}{\text{Instruction}} \times \frac{\text{Time}}{\text{Cycle}}$$

$$\textbf{CPI} = \frac{\text{Cycles}}{\text{Instruction}} = \frac{\text{Time}}{\text{Program}} \div \left(\frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Time}}{\text{Cycle}} \right)$$

Benchmark: CPI of Intel Core i7

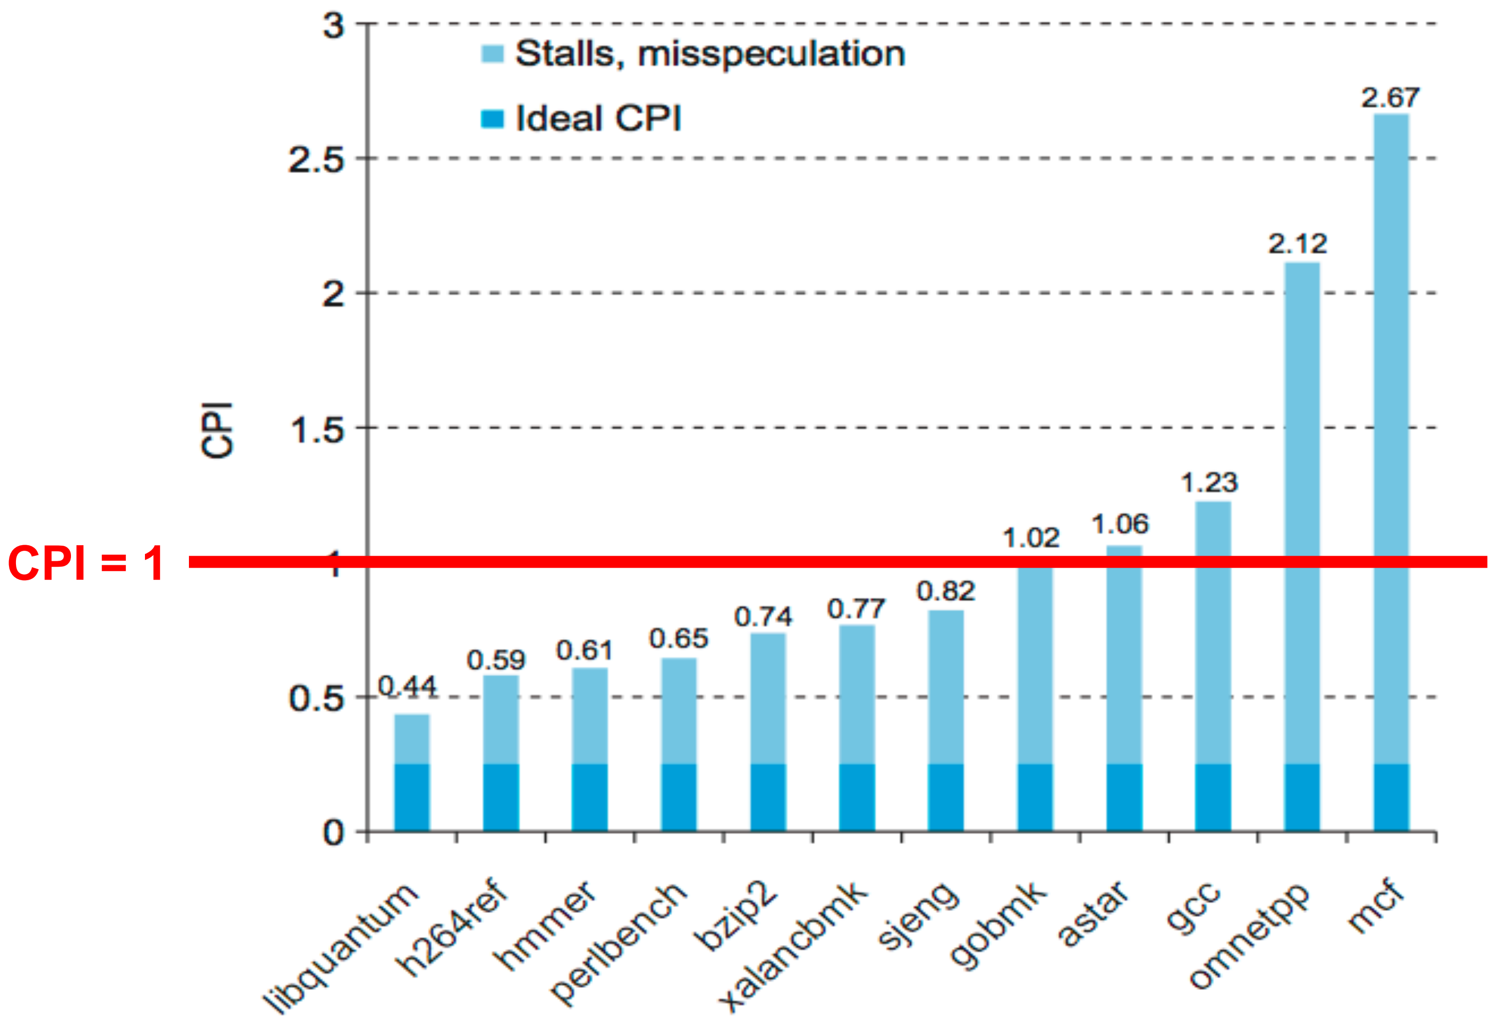

CPI of Intel Core i7 920 running SPEC2006 integer benchmarks.

Calculating CPI Another Way

- First calculate CPI for each individual instruction (**`add`**, **`sub`**, **`and`**, etc.)
- Next calculate frequency of each individual instruction
- Finally multiply these two for each instruction and add them up to get final CPI (the weighted sum)

Example (RISC processor)

Op	$Freq_i$	CPI_i	Prod	(% Time)
ALU	50%	1	.5	(23%)
Load	20%	5	1.0	(45%)
Store	10%	3	.3	(14%)
Branch	20%	2	.4	(18%)
	Instruction Mix		2.2	(Where time spent)

Agenda

- Control Hazards
- Processor Performance
- Complex Pipelines
  – Static Multiple Issues (VLIW)
  – Dynamic Multiple Issues (Superscalar)

Complex Pipeline

- More than one Functional Unit
- Floating point execution!
  - Fadd & Fmul: fixed number of cycles; > 1
  - Fdiv: unknown number of cycles!
- Memory access: on Cache miss unknown number of cycles
- Issue: Assign instruction
  to functional unit

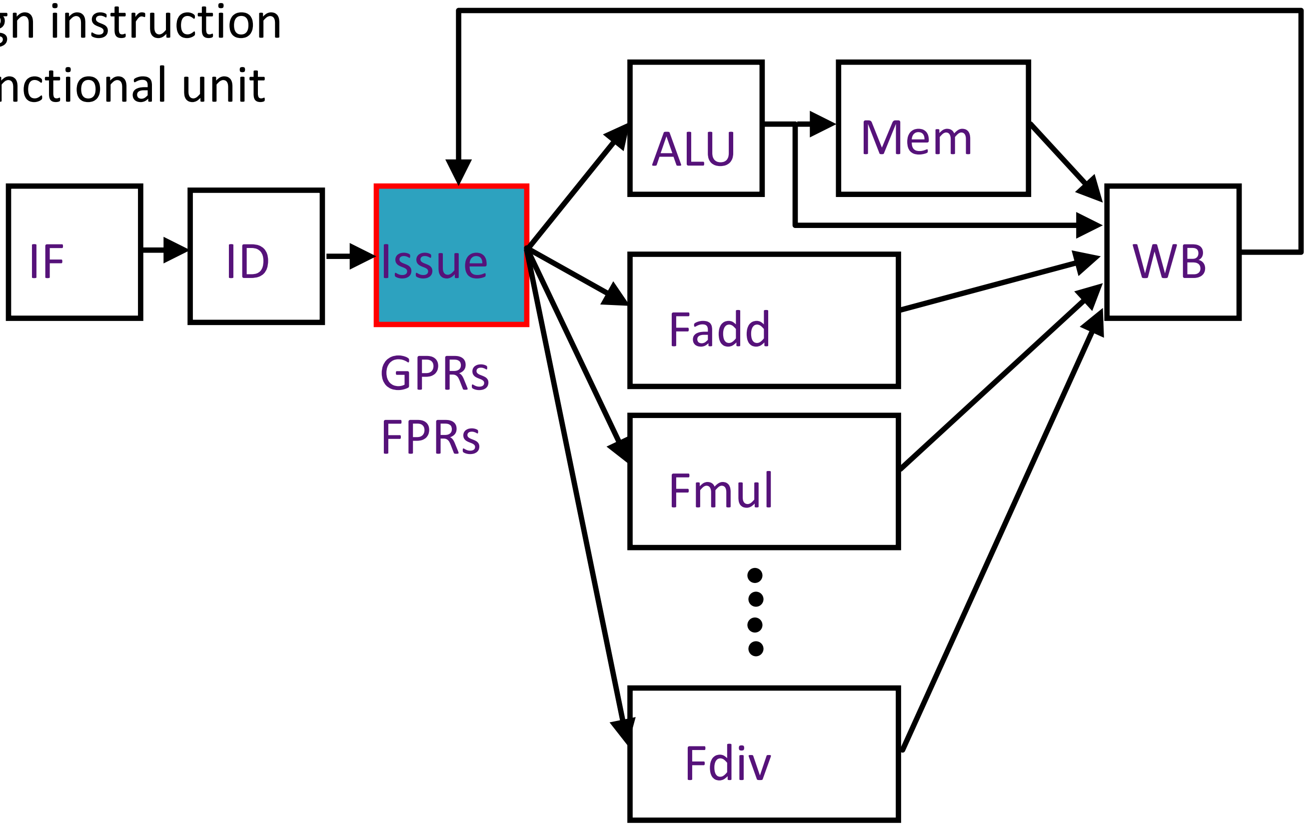

Issues in Complex Pipeline Control

- Structural conflicts at the execution stage if some FPU or memory unit is not pipelined and takes more than one cycle
- Structural conflicts at the write-back stage due to variable latencies of different functional units
- Out-of-order write hazards due to variable latencies of different functional units

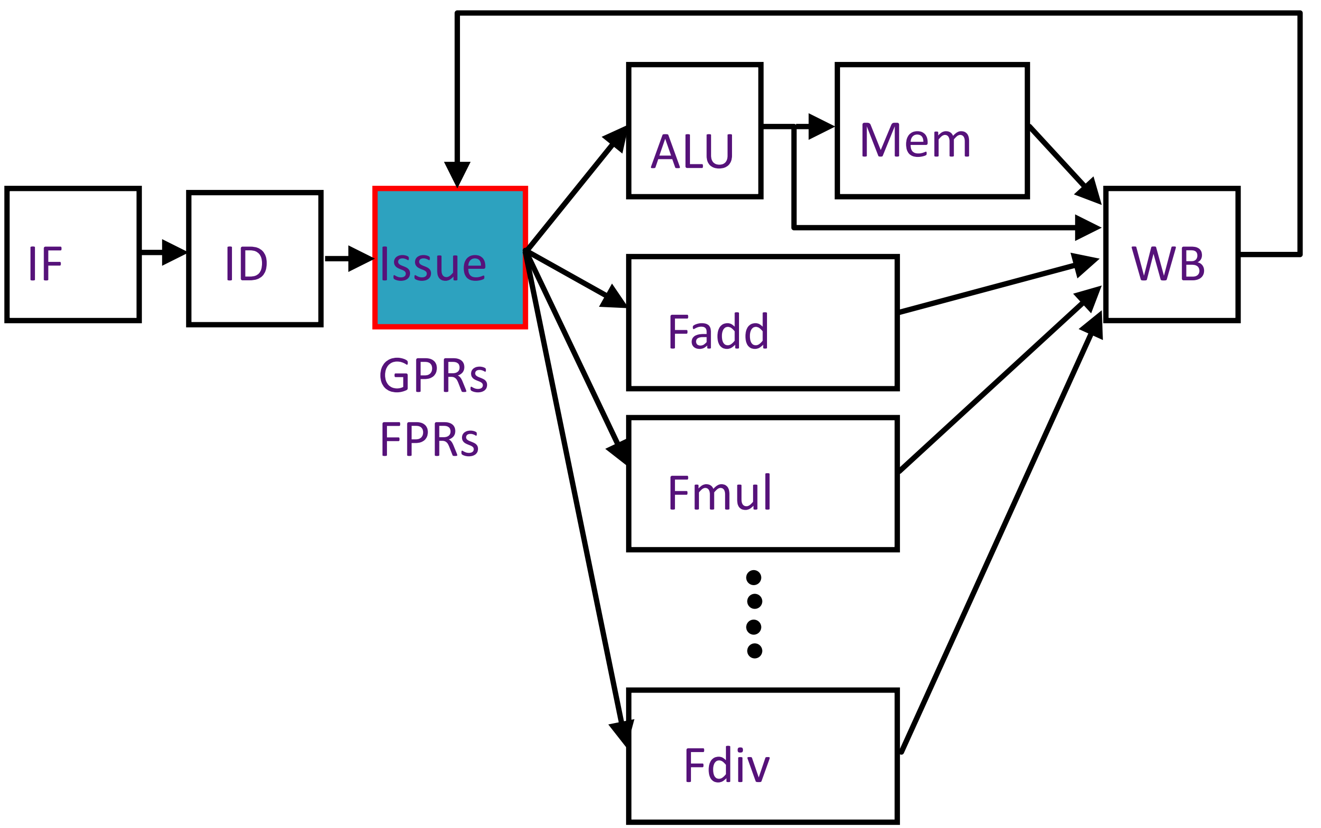

Modern Complex In-Order Pipeline

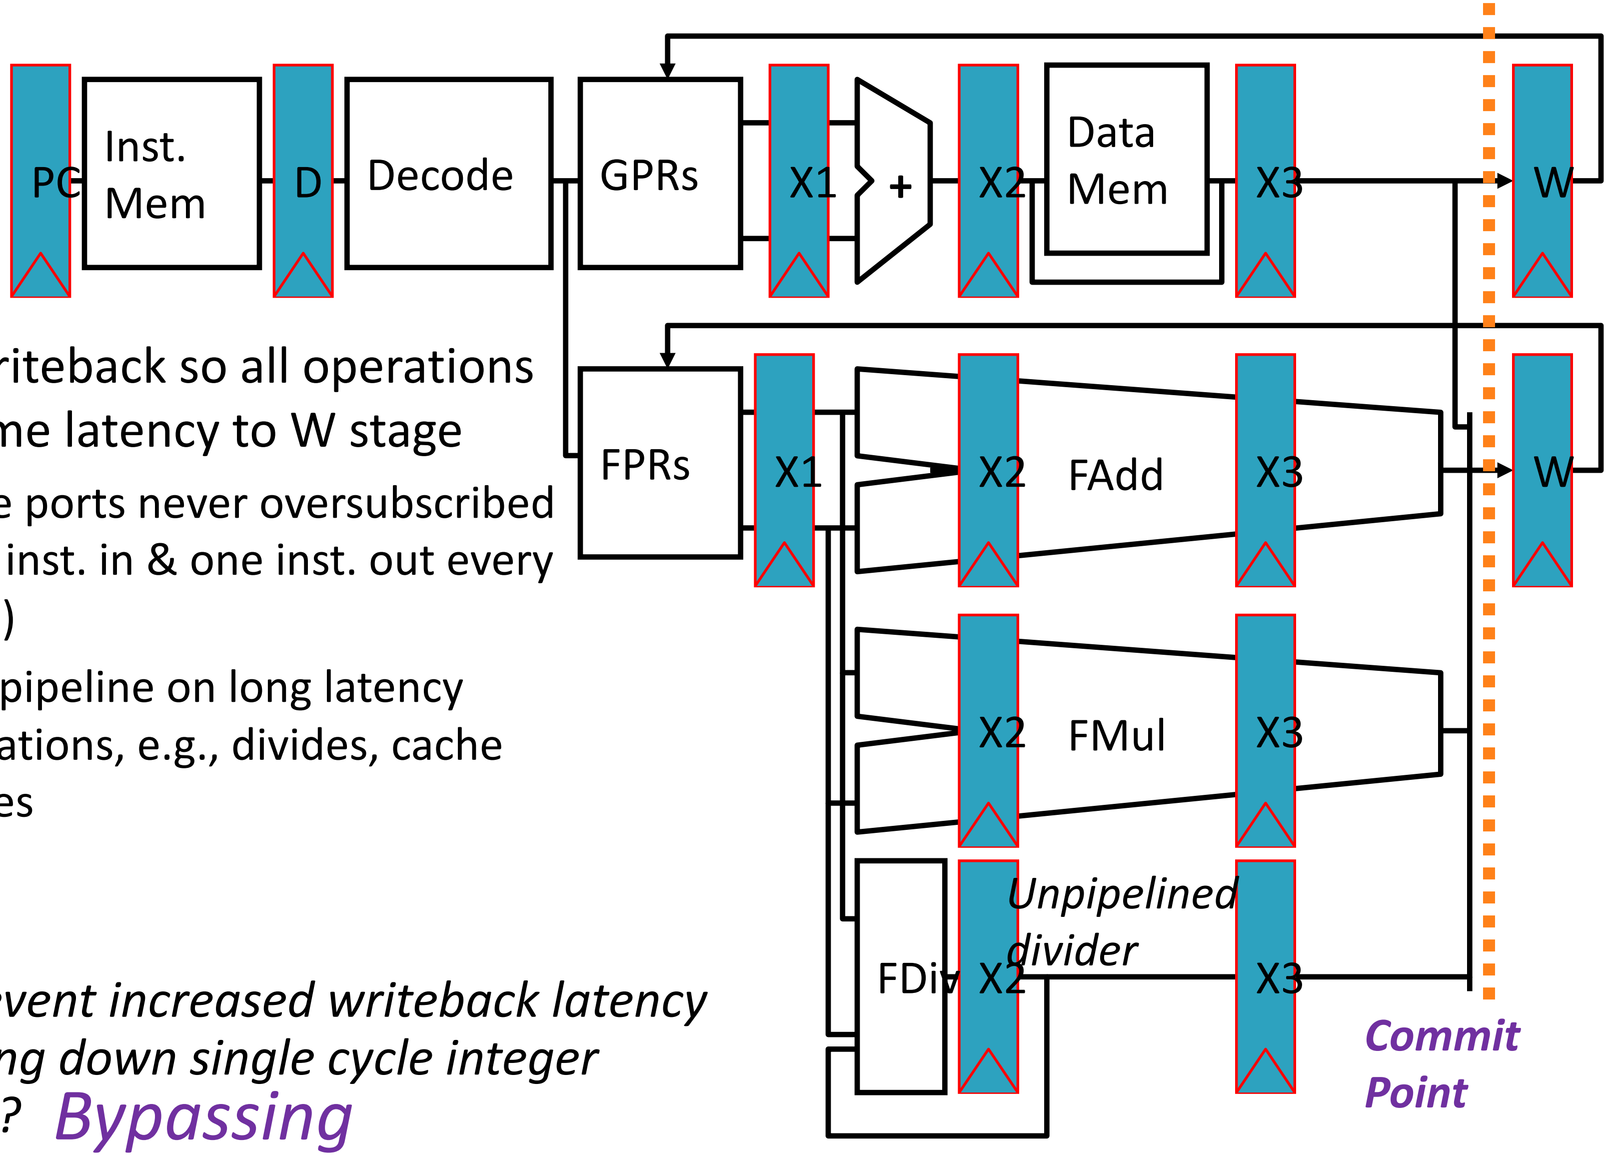

- Delay writeback so all operations have same latency to W stage
  - Write ports never oversubscribed (one inst. in & one inst. out every cycle)
  - Stall pipeline on long latency operations, e.g., divides, cache misses

How to prevent increased writeback latency from slowing down single cycle integer operations? Bypassing

Agenda

- Control Hazards
- Processor Performance
- Complex Pipelines
  - Static Multiple Issues (VLIW)
  - Dynamic Multiple Issues (Superscalar)

Static Multiple Issue

- aka.: Very Long Instruction Word (VLIW)
- Compiler bundles instructions together
- Compiler takes care of hazards
- CPU executes at the same time

Instruction type	Pipe stages							
ALU or branch instruction	IF	ID	EX	MEM	WB			
Load or store instruction	IF	ID	EX	MEM	WB			
ALU or branch instruction		IF	ID	EX	MEM	WB		
Load or store instruction		IF	ID	EX	MEM	WB		
ALU or branch instruction			IF	ID	EX	MEM	WB	
Load or store instruction			IF	ID	EX	MEM	WB	
ALU or branch instruction				IF	ID	EX	MEM	WB
Load or store instruction				IF	ID	EX	MEM	WB

Static Two-Issue RISC-V Datapath

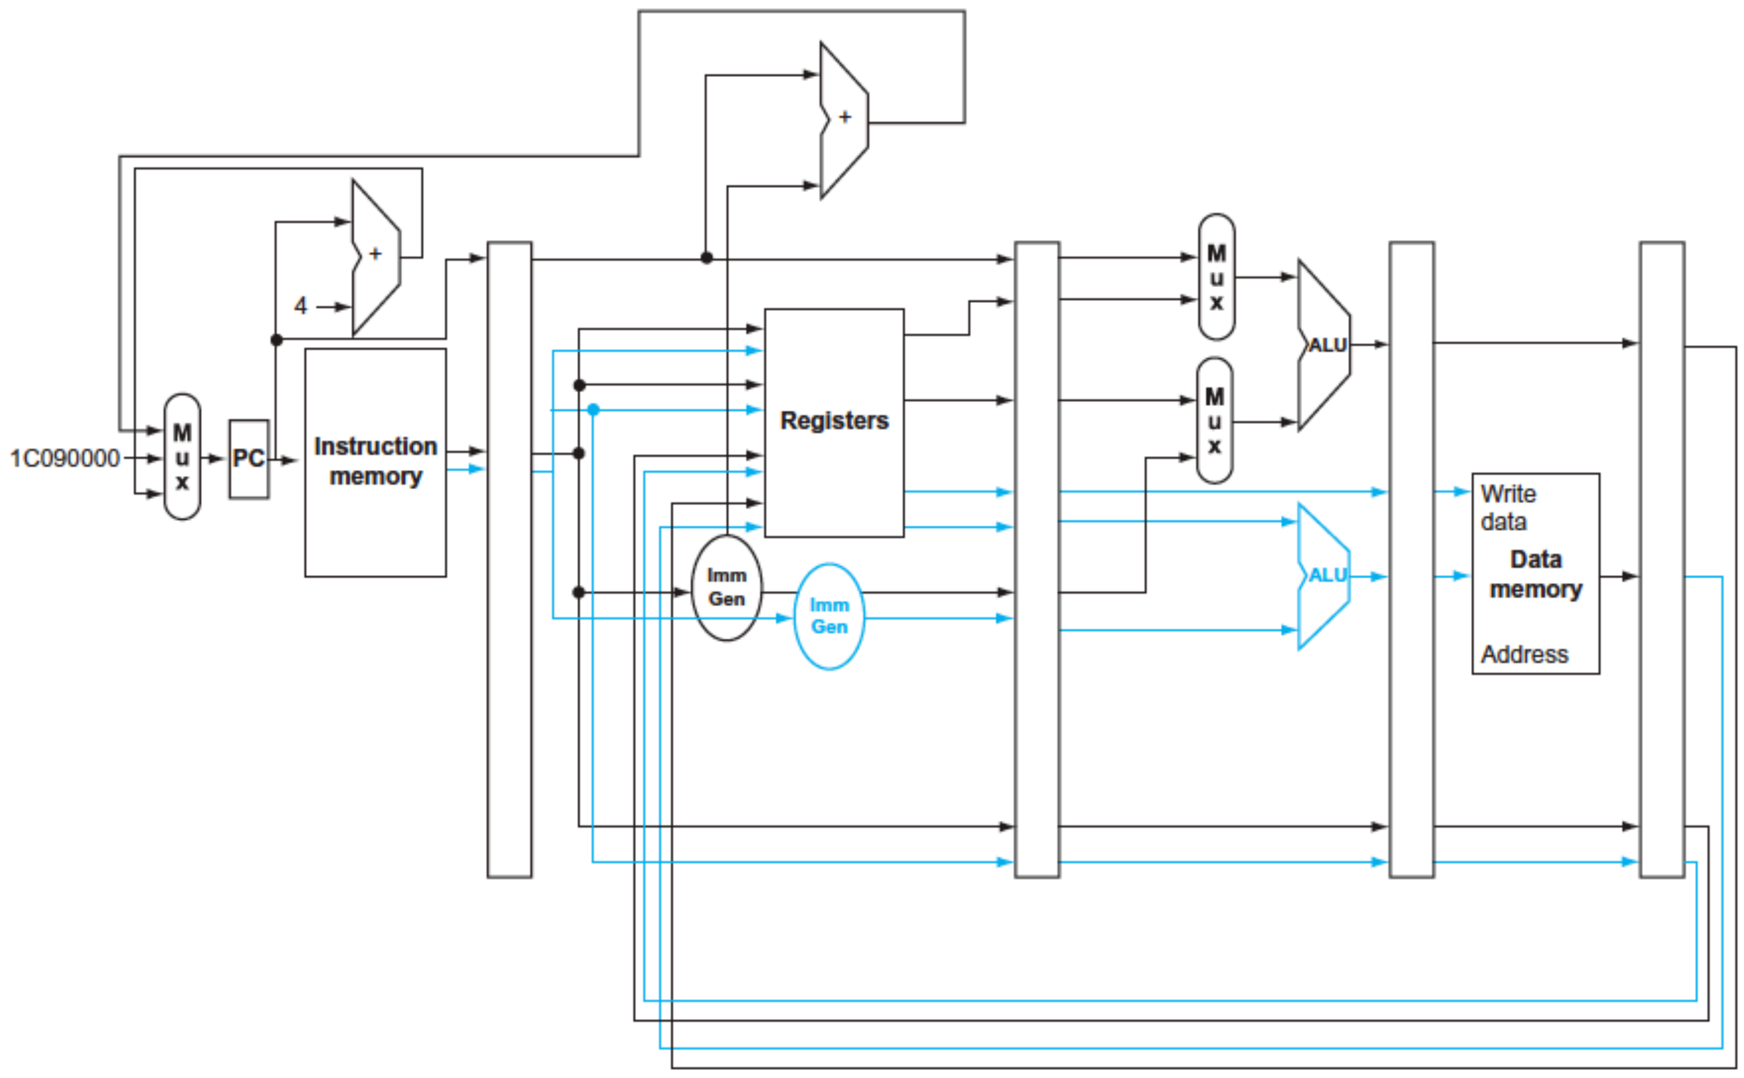

In-Order Superscalar Pipeline

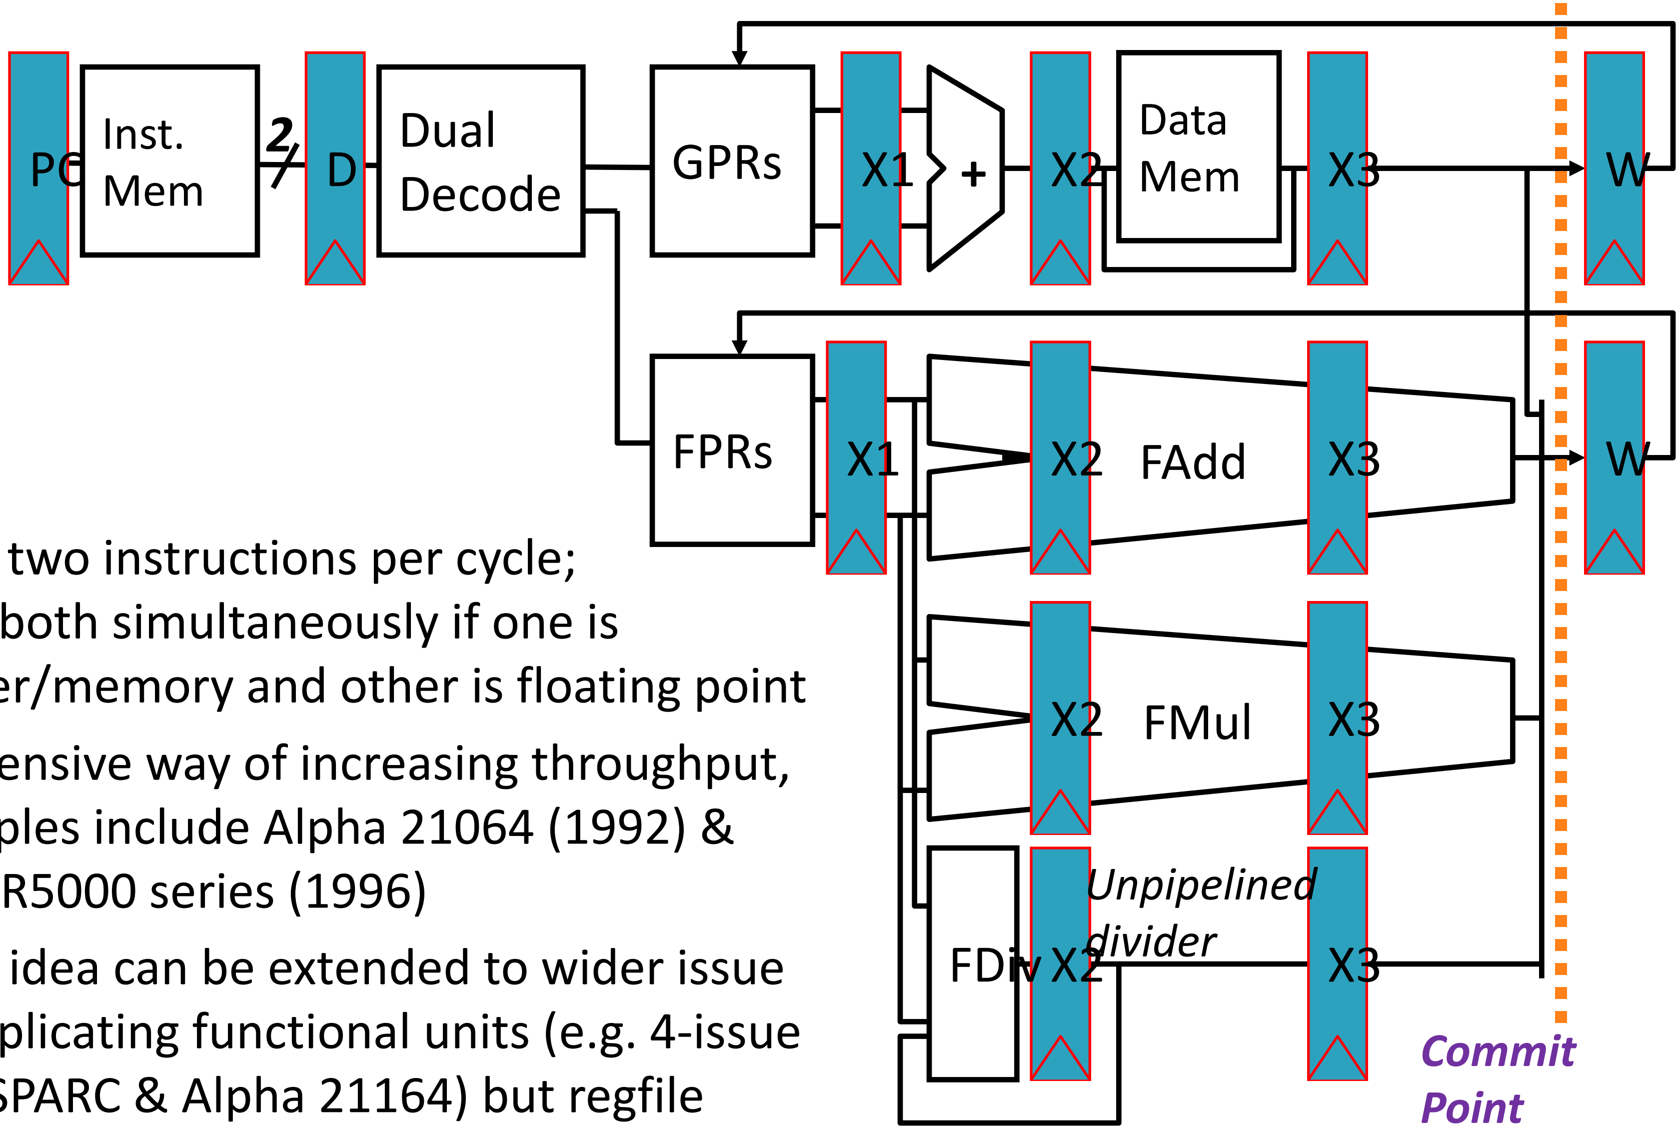

- Fetch two instructions per cycle; issue both simultaneously if one is integer/memory and other is floating point
- Inexpensive way of increasing throughput, examples include Alpha 21064 (1992) & MIPS R5000 series (1996)
- Same idea can be extended to wider issue by duplicating functional units (e.g. 4-issue UltraSPARC & Alpha 21164) but regfile ports and bypassing costs grow quickly

Agenda

- Control Hazards
- Processor Performance
- Complex Pipelines
  - Static Multiple Issues (VLIW)
  - Dynamic Multiple Issues (Superscalar)

Superscalar: Dynamic Multiple Issue

- Hardware guarantees correct execution =>
  - Compiler does not need to (but can) optimize
- Dynamic pipeline scheduling:
  - Re-order instructions based on:
    - What functional units are free
    - Avoiding of data hazards
  - Reservation Station
    - Buffer of instructions waiting to be executed
    - With operands (Registers) needed
    - Once all operands are available: execute!
  - Commit Unit (Reorder buffer): supply the operands to reservation station; write to register
  - OR: Unified Physical Register File :
    Registers are renamed for use in reservation station and commit unit

Phases of Instruction Execution

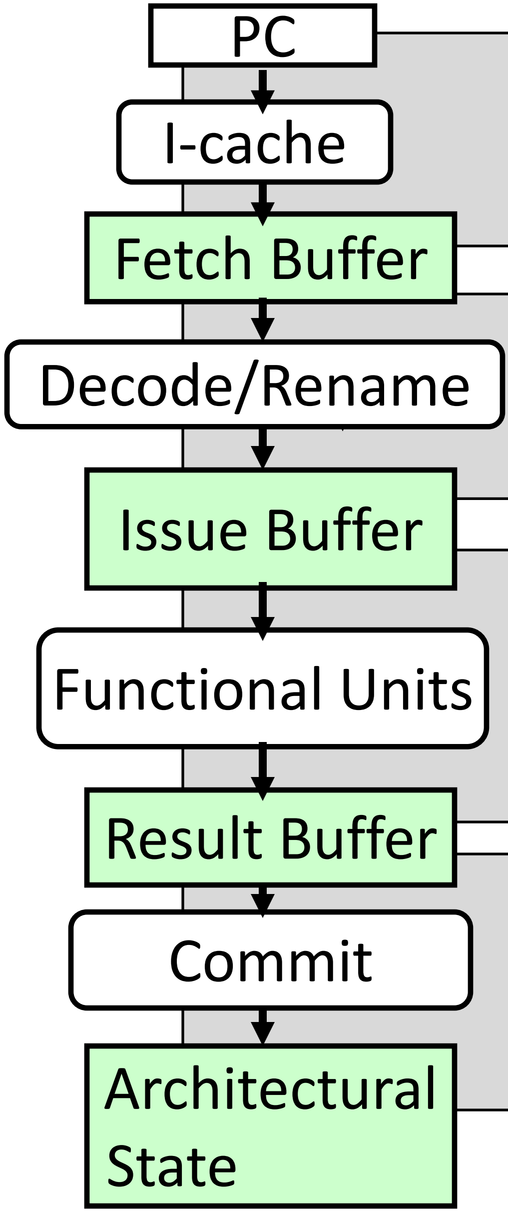

Fetch: *Instruction bits retrieved from instruction cache.*

Decode: *Instructions dispatched to appropriate issue buffer*

Execute: *Instructions and operands issued to functional units. When execution completes, all results and exception flags are available.*

Commit: *Instruction irrevocably updates architectural state (aka "graduation"), or takes precise trap/interrupt.*

Separating Completion from Commit

- Re-order buffer holds register results from completion until commit
  - Entries allocated in program order during decode
  - Buffers completed values and exception state until in-order commit point
  - Completed values can be used by dependents before committed (bypassing)
  - Each entry holds program counter, instruction type, destination register specifier and value if any, and exception status (info often compressed to save hardware)

In-Order versus Out-of-Order Phases

- Instruction fetch/decode/rename always in-order
  - Need to parse ISA sequentially to get correct semantics
  - *Proposals for speculative OoO instruction fetch, e.g., Multiscalar. Predict control flow and data dependencies across sequential program segments fetched/decoded/executed in parallel, fixup if prediction wrong*
- Dispatch (place instruction into machine buffers to wait for issue) also always in-order
  - Some use “Dispatch” to mean “Issue”

In-Order Versus Out-of-Order Issue

- In-order (InO) issue:
  - Issue **stalls** on read after write (RAW), dependencies or structural hazards, or possibly write after read (WAR), write after write (WAW) hazards
  - Instruction cannot issue to execution units unless all preceding instructions have issued to execution units
- Out-of-order (OoO) issue:
  - Instructions dispatched in program order to *reservation stations (or other forms of instruction buffer)* to wait for operands to arrive, or other hazards to clear
  - While earlier instructions wait in issue buffers, following instructions can be dispatched and issued out-of-order

In-Order versus Out-of-Order Completion

- All but simplest machines have out-of-order completion, due to different latencies of functional units and desire to bypass values as soon as available
- Classic RISC V-stage integer pipeline just barely has in-order completion
  - Load takes two cycles, but following one-cycle integer op completes at same time, not earlier
  - Adding pipelined FPU immediately brings OoO completion

Superscalar Intel Processors

- Pentium 4: Marketing demanded higher clock rate => deeper pipelines & high power consumption
- Afterwards: Multi-core processors

Microprocessor	Year	Clock Rate	Pipeline Stages	Issue Width	Out-of-Order/ Speculation	Cores/ Chip	Power	
Intel 486	1989	25 MHz	5	1	No	1	5	W
Intel Pentium	1993	66 MHz	5	2	No	1	10	W
Intel Pentium Pro	1997	200 MHz	10	3	Yes	1	29	W
Intel Pentium 4 Willamette	2001	2000 MHz	22	3	Yes	1	75	W
Intel Pentium 4 Prescott	2004	3600 MHz	31	3	Yes	1	103	W
Intel Core	2006	2930 MHz	14	4	Yes	2	75	W
Intel Core i5 Nehalem	2010	3300 MHz	14	4	Yes	2–4	87	W
Intel Core i5 Ivy Bridge	2012	3400 MHz	14	4	Yes	8	77	W

Arm Cortex A53 & Intel Core i7 920

Processor	ARM A53	Intel Core i7 920
Market	Personal Mobile Device	Server, Cloud
Thermal design power	100 milliWatts (1 core @ 1 GHz)	130 Watts
Clock rate	1.5 GHz	2.66 GHz
Cores/Chip	4 (configurable)	4
Floating point?	Yes	Yes
Multiple Issue?	Dynamic	Dynamic
Peak instructions/clock cycle	2	4
Pipeline Stages	8	14
Pipeline schedule	Static In-order	Dynamic Out-of-order with Speculation
Branch prediction	Hybrid	2-level
1st level caches/core	16-64 KiB I, 16-64 KiB D	32 KiB I, 32 KiB D
2nd level cache/core	128–2048 KiB (shared)	256 KiB (per core)
3rd level cache (shared)	(platform dependent)	2–8 MiB

ARM Cortex A53 Pipeline

- Prediction 1 clock cycle! Predict: branches, future function returns; 8 clock cycles on mis-prediction (flush pipeline)

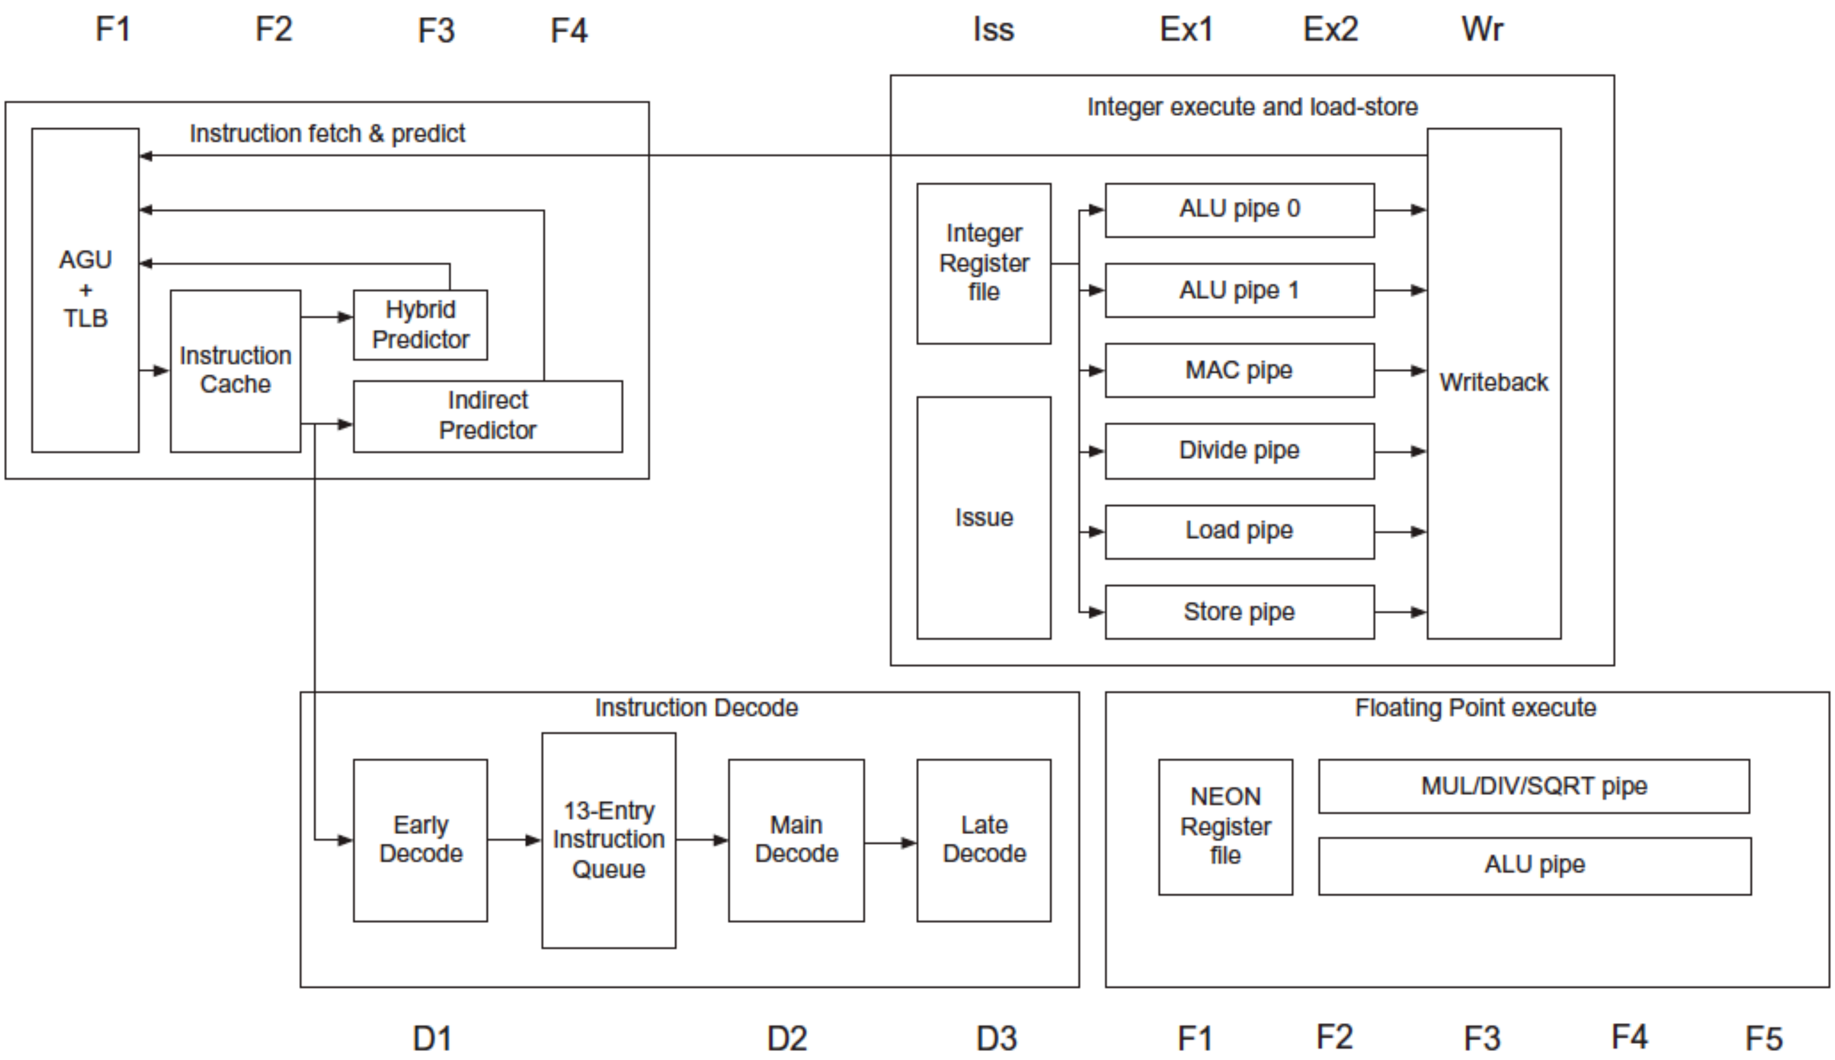

Speculative & Out-of-Order Execution