

CS 152 Computer Architecture and Engineering
CS252 Graduate Computer Architecture

**Lecture 10 – Complex Pipelines,
Out-of-Order Issue, Register Renaming**

Krste Asanovic

Electrical Engineering and Computer Sciences
University of California at Berkeley

`http://www.eecs.berkeley.edu/~krste`
`http://inst.eecs.berkeley.edu/~cs152`

Last time in Lecture 9

- Modern page-based virtual memory systems provide:
 - Translation, Protection, Virtual memory.
- Translation and protection information stored in page tables, held in main memory
- Translation and protection information cached in “translation-lookaside buffer” (TLB) to provide single-cycle translation+protection check in common case
- Virtual memory interacts with cache design
 - Physical cache tags require address translation before tag lookup, or use untranslated offset bits to index cache.
 - Virtual tags do not require translation before cache hit/miss determination, but need to be flushed or extended with ASID to cope with context swaps. Also, must deal with virtual address aliases (usually by disallowing copies in cache).

Types of Data Hazards

Consider executing a sequence of

$$r_k \leftarrow r_i \text{ op } r_j$$

type of instructions

Data-dependence

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

Read-after-Write
(RAW) hazard

Anti-dependence

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

Write-after-Read
(WAR) hazard

Output-dependence

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

Write-after-Write
(WAW) hazard

Register vs. Memory Dependence


Data hazards due to register operands can be determined at the decode stage, but data hazards due to memory operands can be determined only after computing the effective address

Store: $M[r1 + disp1] \leftarrow r2$

Load: $r3 \leftarrow M[r4 + disp2]$

Does $(r1 + disp1) = (r4 + disp2)$?

Data Hazards: An Example



RAW Hazards

WAR Hazards

WAW Hazards

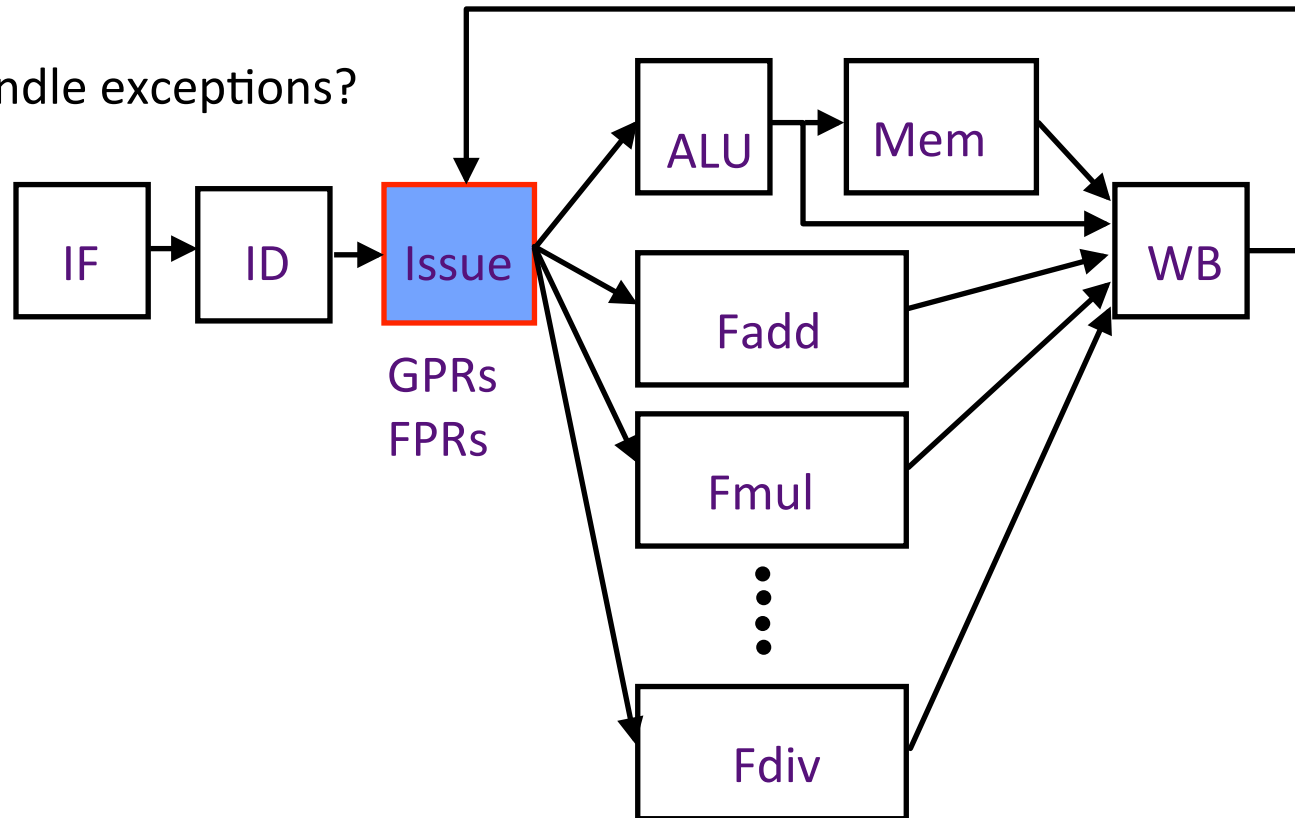
Complex Pipelining: Motivation

Pipelining becomes complex when we want high performance in the presence of:

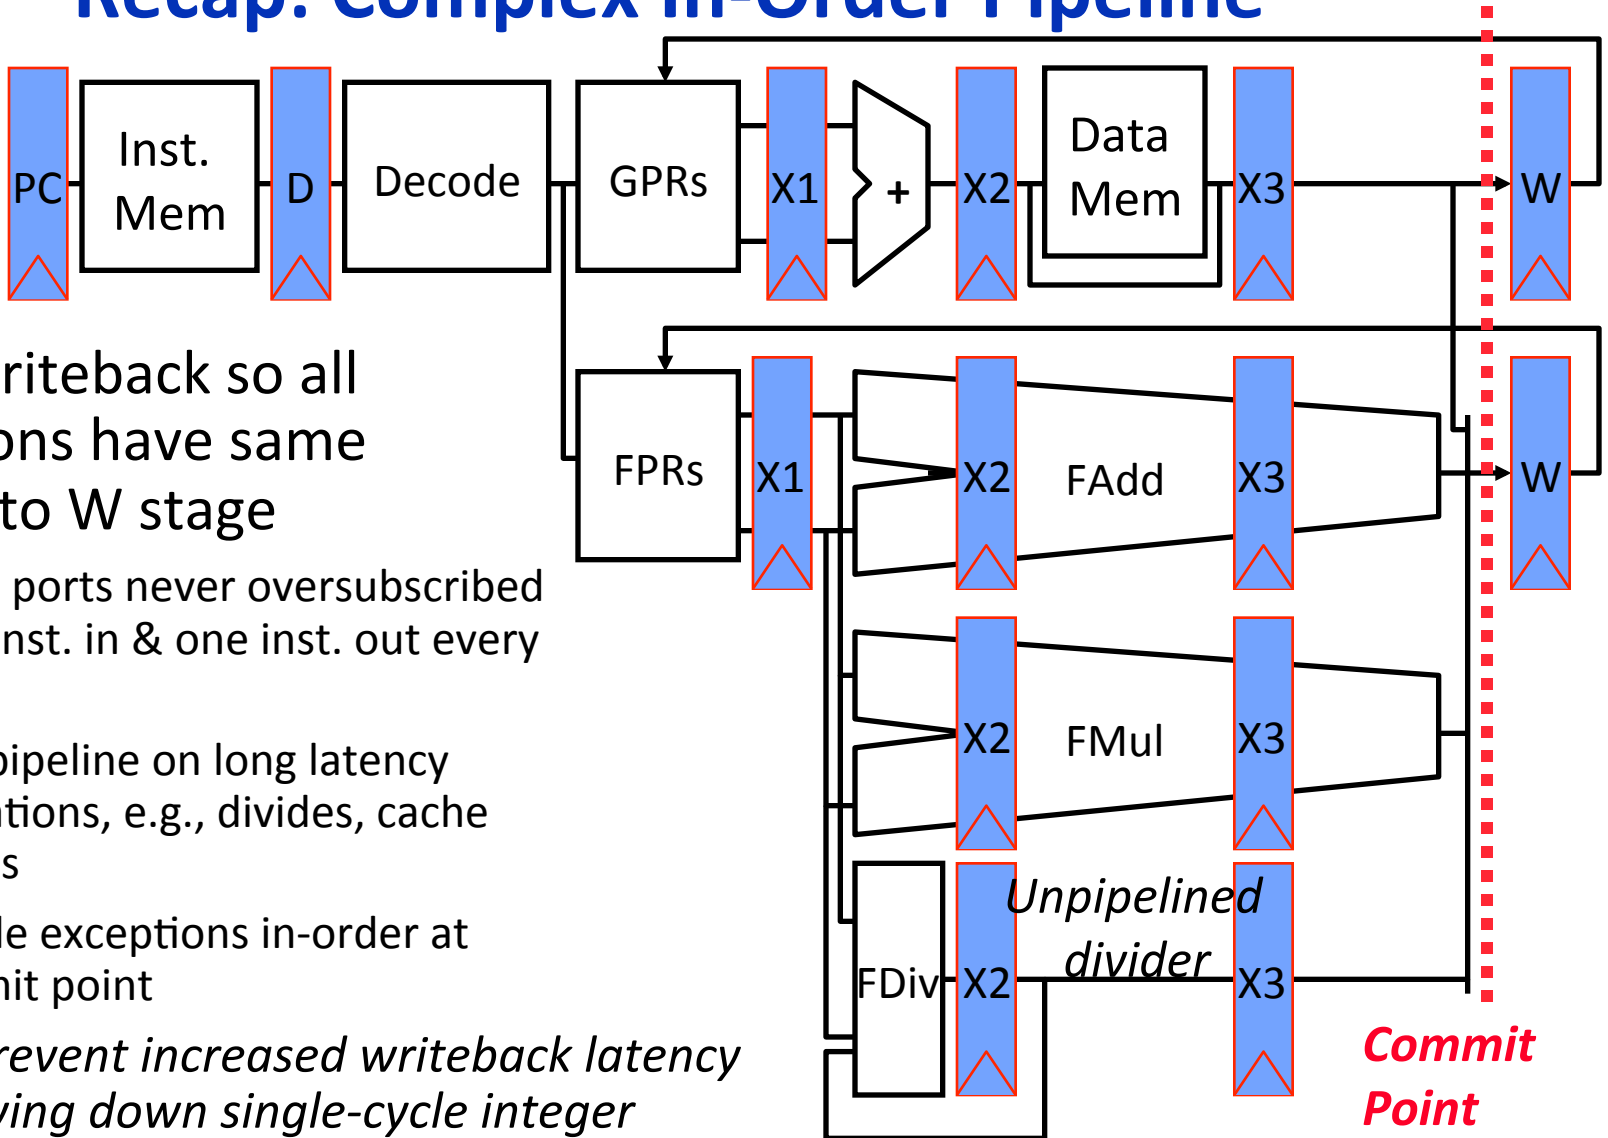
- Long latency or partially pipelined floating-point units
- Memory systems with variable access time
- Multiple arithmetic and memory units

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
- How to handle exceptions?



Recap: Complex In-Order Pipeline

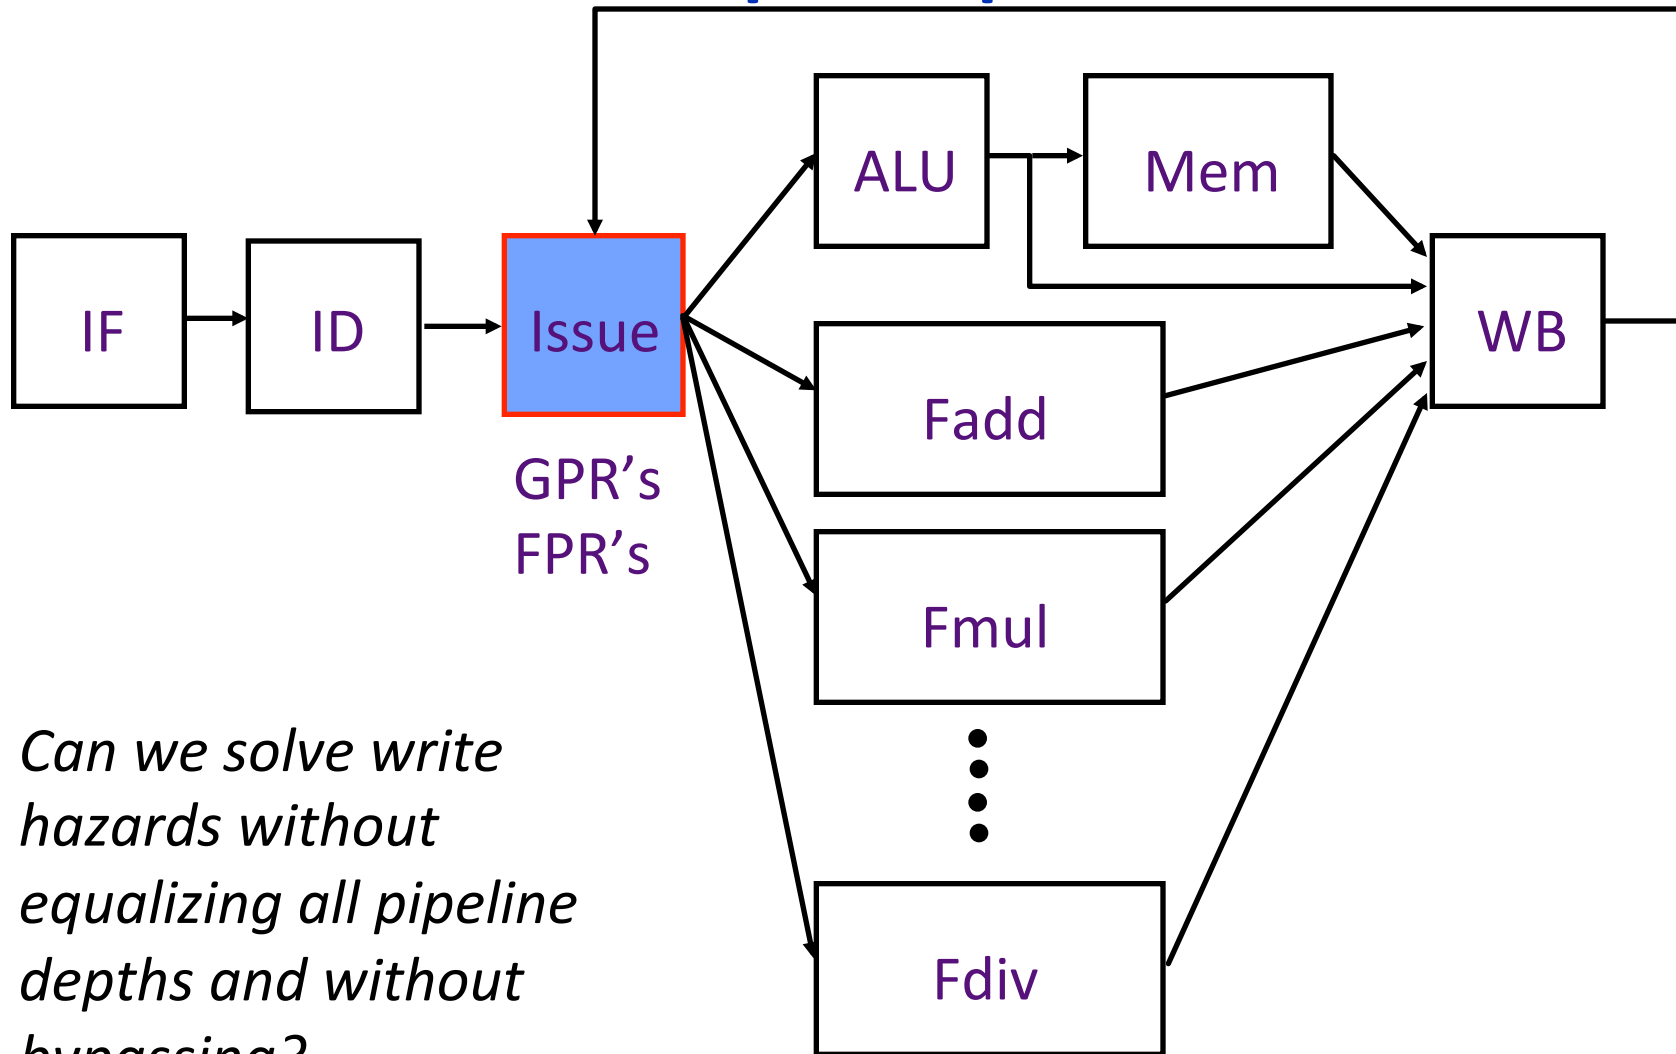


- Delay writeback so all operations have same latency to W stage
 - Write ports never oversubscribed (one inst. in & one inst. out every cycle)
 - Stall pipeline on long latency operations, e.g., divides, cache misses
 - Handle exceptions in-order at commit point

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


Bypassing

Complex Pipeline




Can we solve write hazards without equalizing all pipeline depths and without bypassing?

Instruction Scheduling



Valid orderings:

| | | | | | | |
|---------------------|-------|-------|-------|-------|-------|-------|
| <i>in-order</i> | I_1 | I_2 | I_3 | I_4 | I_5 | I_6 |
| <i>out-of-order</i> | I_2 | I_1 | I_3 | I_4 | I_5 | I_6 |
| <i>out-of-order</i> | I_1 | I_2 | I_3 | I_5 | I_4 | I_6 |



Out-of-order Completion

In-order Issue

| | | | | | | <i>Latency</i> |
|--------------------------|---------|----------------|---|-----------------------|-------------------|----------------|
| I_1 | FDIV.D | f6, | f6, | f4 | | 4 |
| I_2 | FLD | f2, | 45(x3) | | | 1 |
| I_3 | FMULT.D | f0, | f2, | f4 | | 3 |
| I_4 | FDIV.D | f8, | f6, | f2 | | 4 |
| I_5 | FSUB.D | f10, | f0, | f6 | | 1 |
| I_6 | FADD.D | f6, | f8, | f2 | | 1 |
| <i>in-order comp</i> | | 1 2 | <u>1</u> <u>2</u> 3 4 | <u>3</u> 5 <u>4</u> 6 | <u>5</u> <u>6</u> | |
| <i>out-of-order comp</i> | | 1 2 <u>2</u> 3 | <u>1</u> 4 <u>3</u> 5 <u>5</u> 4 6 <u>6</u> | | | |

When is it Safe to Issue an Instruction?

Suppose a data structure keeps track of all the instructions in all the functional units

The following checks need to be made before the Issue stage can dispatch an instruction

- Is the required function unit available?
- Is the input data available? \Rightarrow RAW?
- Is it safe to write the destination? \Rightarrow WAR?
WAW?
- Is there a structural conflict at the WB stage?

A Data Structure for Correct Issues

Keeps track of the status of Functional Units

| <i>Name</i> | <i>Busy</i> | <i>Op</i> | <i>Dest</i> | <i>Src1</i> | <i>Src2</i> |
|-------------|-------------|-----------|-------------|-------------|-------------|
| Int | | | | | |
| Mem | | | | | |
| Add1 | | | | | |
| Add2 | | | | | |
| Add3 | | | | | |
| Mult1 | | | | | |
| Mult2 | | | | | |
| Div | | | | | |

The instruction i at the Issue stage consults this table

FU available?

check the busy column

RAW?

search the dest column for i 's sources

WAR?

search the source columns for i 's destination

WAW?

search the dest column for i 's destination

*An entry is added to the table if no hazard is detected;
An entry is removed from the table after Write-Back*

Simplifying the Data Structure

Assuming In-order Issue

Suppose the instruction is not dispatched by the Issue stage if a RAW hazard exists or the required FU is busy, and that operands are latched by functional unit on issue:

Can the dispatched instruction cause a
WAR hazard ?

NO: Operands read at issue

WAW hazard ?

YES: Out-of-order completion

Simplifying the Data Structure ...

- No WAR hazard
 - ⇒ no need to keep src1 and src2
- The Issue stage does not dispatch an instruction in case of a WAW hazard
 - ⇒ a register name can occur at most once in the dest column
- WP[reg#] : a bit-vector to record the registers for which writes are pending
 - These bits are set by the Issue stage and cleared by the WB stage
 - ⇒ Each pipeline stage in the FU's must carry the register destination field and a flag to indicate if it is valid

Scoreboard for In-order Issues

Busy[FU#] : a bit-vector to indicate FU's availability.
(FU = Int, Add, Mult, Div)

These bits are hardwired to FU's.

WP[reg#] : a bit-vector to record the registers for which writes are pending.

These bits are set by Issue stage and cleared by WB stage

Issue checks the instruction (opcode dest src1 src2) against the scoreboard (Busy & WP) to dispatch

| | |
|---------------|----------------------|
| FU available? | Busy[FU#] |
| RAW? | WP[src1] or WP[src2] |
| WAR? | <i>cannot arise</i> |
| WAW? | WP[dest] |


Scoreboard Dynamics

| | Functional Unit Status | | | | | Registers Reserved for Writes | |
|-----|------------------------|--------|---------|--------|----|----------------------------------|---------------|
| | Int(1) | Add(1) | Mult(3) | Div(4) | WB | | |
| t0 | I_1 | | | f6 | | | f6 |
| t1 | I_2 f2 | | | f6 | | | f6, f2 |
| t2 | | | | | f6 | f2 | f6, f2 I_2 |
| t3 | I_3 | | f0 | | | f6 | f6, f0 |
| t4 | | | | f0 | | f6 | f6, f0 I_1 |
| t5 | I_4 | | | f0 f8 | | | f0, f8 |
| t6 | | | | | f8 | f0 | f0, f8 I_3 |
| t7 | I_5 | f10 | | | f8 | | f8, f10 |
| t8 | | | | | | f8 f10 | f8, f10 I_5 |
| t9 | | | | | | f8 | f8 I_4 |
| t10 | I_6 | f6 | | | | | f6 |
| t11 | | | | | | f6 | f6 I_6 |

| | | | | |
|-------|---------|------|--------|----|
| I_1 | FDIV.D | f6, | f6, | f4 |
| I_2 | FLD | f2, | 45(x3) | |
| I_3 | FMULT.D | f0, | f2, | f4 |
| I_4 | FDIV.D | f8, | f6, | f2 |
| I_5 | FSUB.D | f10, | f0, | f6 |
| I_6 | FADD.D | f6, | f8, | f2 |

In-Order Issue Limitations: *an example*

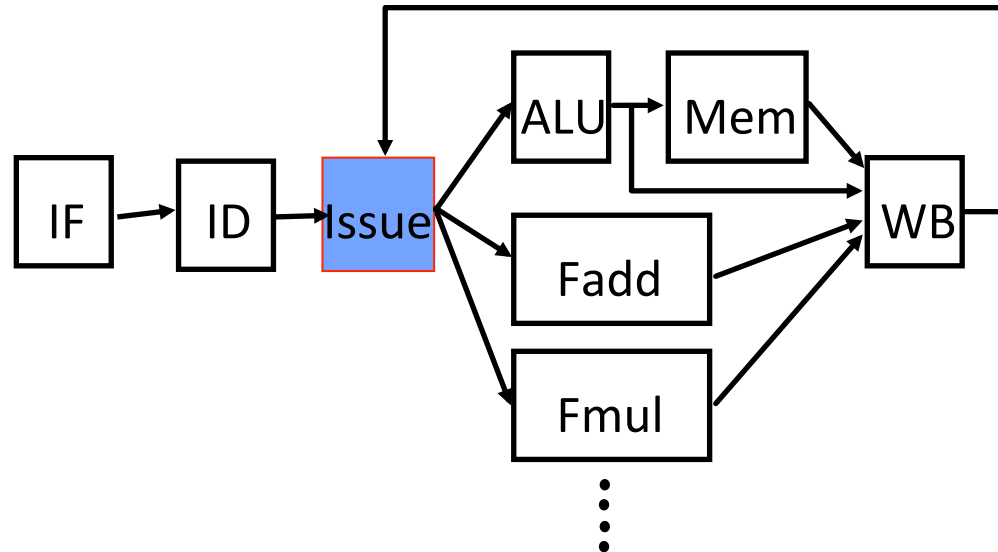
| | | | | | |
|---|---------|------|--------|----|----------------|
| | | | | | <i>latency</i> |
| 1 | FLD | f2, | 34(x2) | | 1 |
| 2 | FLD | f4, | 45(x3) | | <i>long</i> |
| 3 | FMULT.D | f6, | f4, | f2 | 3 |
| 4 | FSUB.D | f8, | f2, | f2 | 1 |
| 5 | FDIV.D | f4, | f2, | f8 | 4 |
| 6 | FADD.D | f10, | f6, | f4 | 1 |



In-order: 1 (2,1) 2 3 4 4 3 5 5 6 6

In-order issue restriction prevents instruction 4 from being dispatched


Out-of-Order Issue



- Issue stage buffer holds multiple instructions waiting to issue.
- Decode adds next instruction to buffer if there is space and the instruction does not cause a WAR or WAW hazard.
 - Note: WAR possible again because issue is out-of-order (WAR not possible with in-order issue and latching of input operands at functional unit)
- Any instruction in buffer whose RAW hazards are satisfied can be issued (for now, at most one dispatch per cycle). On a write back (WB), new instructions may get enabled.

Issue Limitations: In-Order and Out-of-Order

| | | | | | <i>latency</i> |
|---|---------|------|--------|----|----------------|
| 1 | FLD | f2, | 34(x2) | | 1 |
| 2 | FLD | f4, | 45(x3) | | <i>long</i> |
| 3 | FMULT.D | f6, | f4, | f2 | 3 |
| 4 | FSUB.D | f8, | f2, | f2 | 1 |
| 5 | FDIV.D | f4, | f2, | f8 | 4 |
| 6 | FADD.D | f10, | f6, | f4 | 1 |



In-order: 1 (2,1) 2 3 4 4 3 5 . . . 5 6 6
 Out-of-order: 1 (2,1) 4 4 2 3 . . 3 5 . . . 5 6 6

Out-of-order execution did not allow any significant improvement!

How many instructions can be in the pipeline?

Which features of an ISA limit the number of instructions in the pipeline?

Number of Registers

Out-of-order dispatch by itself does not provide any significant performance improvement!

CS152 Administrivia

- Midterm in class Monday Feb 26
 - Covers lectures 1 – 9, plus assigned problem sets, labs, readings

CS252 Administrivia

- Project proposal due Monday March 5th
 - Mail PDF of proposal to instructors
 - Give a <5-minute presentation in class in discussion section time on March 5th
- No reading on Monday February 26th due to midterm

Overcoming the Lack of Register Names

Floating Point pipelines often cannot be kept filled with small number of registers.


IBM 360 had only 4 floating-point registers

Can a microarchitecture use more registers than specified by the ISA without loss of ISA compatibility ?

Robert Tomasulo of IBM suggested an ingenious solution in 1967 using on-the-fly *register renaming*

Issue Limitations: In-Order and Out-of-Order

| | | | | | <i>latency</i> |
|---|---------|--------------|--------|------------|----------------|
| 1 | FLD | f2, | 34(x2) | | 1 |
| 2 | FLD | f4, | 45(x3) | | <i>long</i> |
| 3 | FMULT.D | f6, | f4, | f2 | 3 |
| 4 | FSUB.D | f8, | f2, | f2 | 1 |
| 5 | FDIV.D | f4' , | f2, | f8 | 4 |
| 6 | FADD.D | f10, | f6, | f4' | 1 |



In-order: 1 (2,1) 2 3 4 4 3 5 . . . 5 6 6

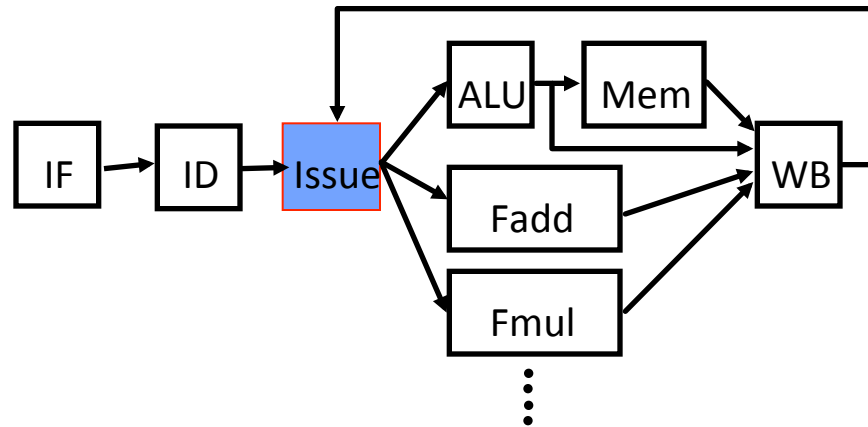
Out-of-order: 1 (2,1) 4 4 5 . . . 2 (3,5) 3 6 6

Any antidependence can be eliminated by renaming.

(renaming \Rightarrow additional storage)

*Can it be done in hardware? **yes!***

Register Renaming



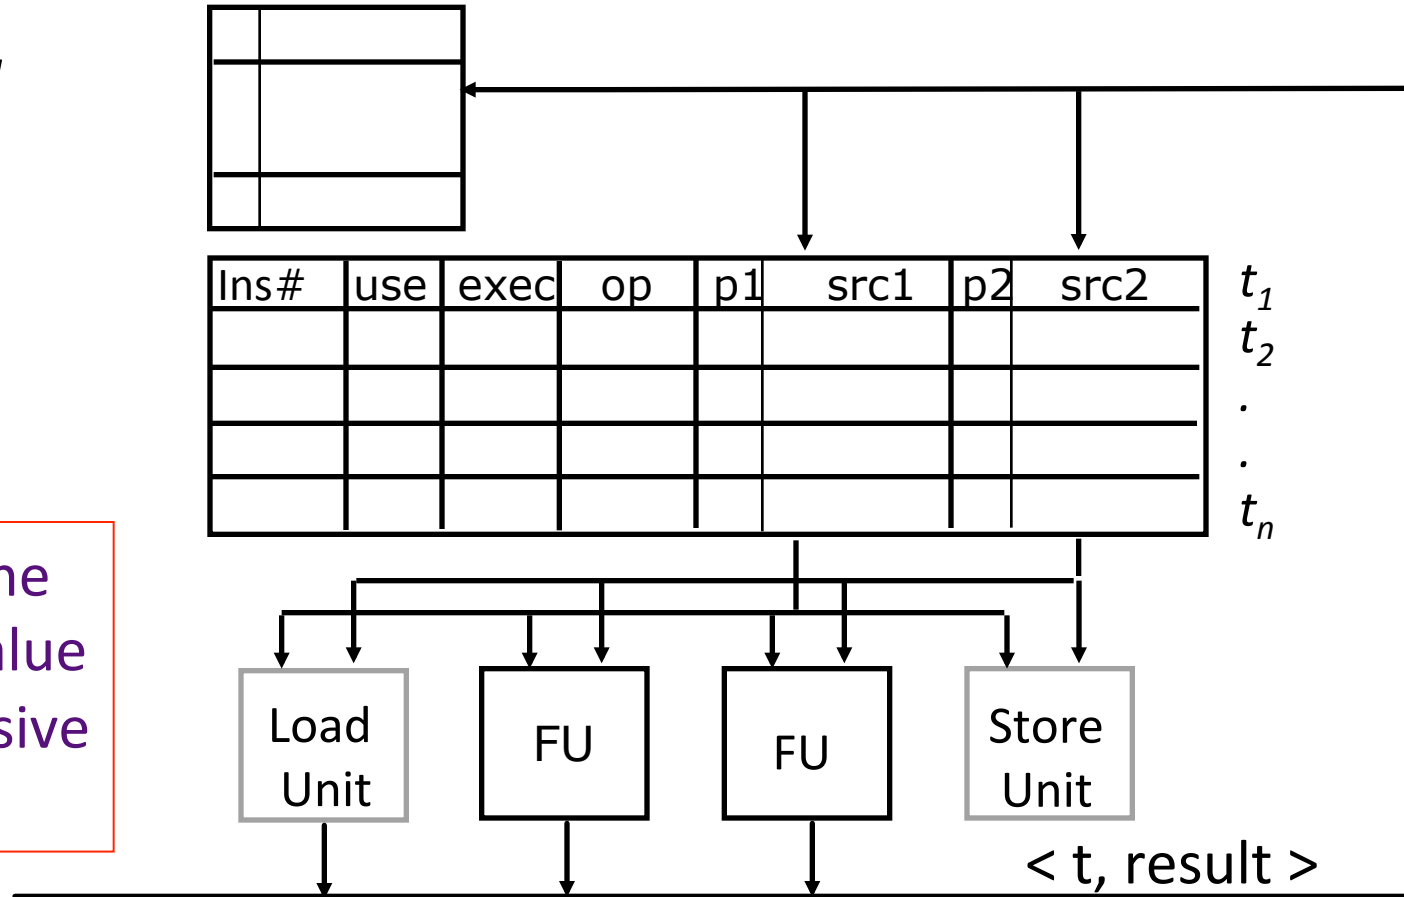
- Decode does register renaming and adds instructions to the issue-stage instruction reorder buffer (ROB)
 - ⇒ renaming makes WAR or WAW hazards impossible
- Any instruction in ROB whose RAW hazards have been satisfied can be dispatched.
 - ⇒ Out-of-order or dataflow execution

Renaming Structures

Renaming table & regfile


Reorder buffer

Replacing the tag by its value is an expensive operation



- Instruction template (i.e., tag t) is allocated by the Decode stage, which also associates tag with register in regfile
- When an instruction completes, its tag is deallocated

Reorder Buffer Management



ROB managed circularly

- “exec” bit is set when instruction begins execution
- When an instruction completes its “use” bit is marked free
- ptr_2 is incremented only if the “use” bit is marked free

Instruction slot is candidate for execution when:

- It holds a valid instruction (“use” bit is set)
- It has not already started execution (“exec” bit is clear)
- Both operands are available (p1 and p2 are set)

Renaming & Out-of-order Issue

An example

Renaming table

| | p | data |
|----|---|------|
| f1 | | |
| f2 | | v1 |
| f3 | | |
| f4 | | t5 |
| f5 | | |
| f6 | | t3 |
| f7 | | |
| f8 | | v4 |

v1

data / t_i

Reorder buffer

| Ins# | use | exec | op | p1 | src1 | p2 | src2 |
|------|-----|------|-----|----|------|----|------|
| 1 | 0 | 0 | LD | | | | |
| 2 | 0 | 0 | LD | | | | |
| 3 | 1 | 0 | MUL | 0 | v2 | 1 | v1 |
| 4 | 0 | 0 | SUB | 1 | v1 | 1 | v1 |
| 5 | 1 | 0 | DIV | 1 | v1 | 0 | t4 |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |
| | | | | | | | |

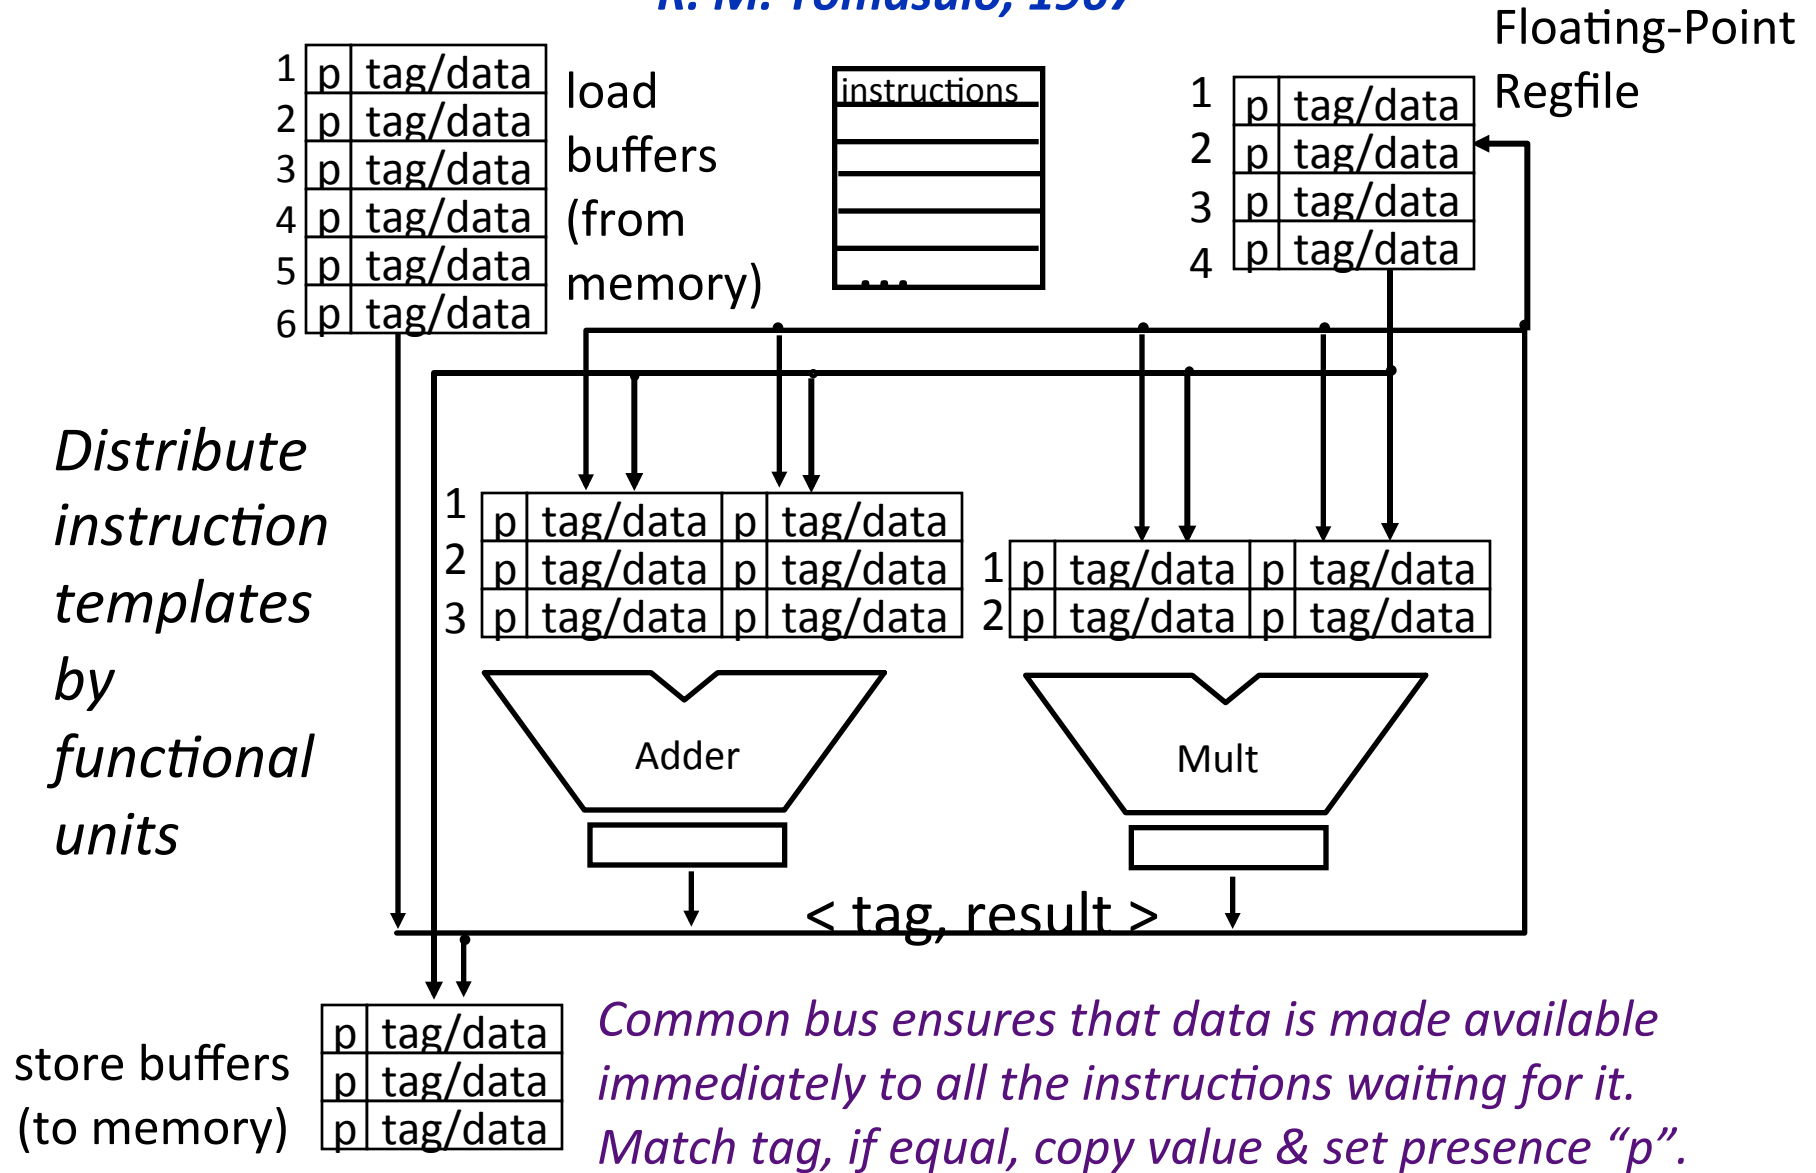
t₁
t₂
t₃
t₄
t₅
.
.

| | | | |
|---|---------|------|--------|
| 1 | FLD | f2, | 34(x2) |
| 2 | FLD | f4, | 45(x3) |
| 3 | FMULT.D | f6, | f4, f2 |
| 4 | FSUB.D | f8, | f2, f2 |
| 5 | FDIV.D | f4, | f2, f8 |
| 6 | FADD.D | f10, | f6, f4 |

- When are tags in sources replaced by data?
Whenever an FU produces data
- When can a name be reused?
Whenever an instruction completes

IBM 360/91 Floating-Point Unit

R. M. Tomasulo, 1967



IBM ACS

- Second supercomputer project (Y) started at IBM in response to CDC6600
- Multiple Dynamic instruction Scheduling (DIS) invented by Lynn Conway for ACS
 - Used unary encoding of register specifiers and wired-OR logic to detect any hazards (similar design used in Alpha 21264 in 1995!)
- Seven-issue, out-of-order processor
 - Two decoupled streams, each with DIS
- Cancelled in favor of IBM360-compatible machines

Out-of-Order Fades into Background

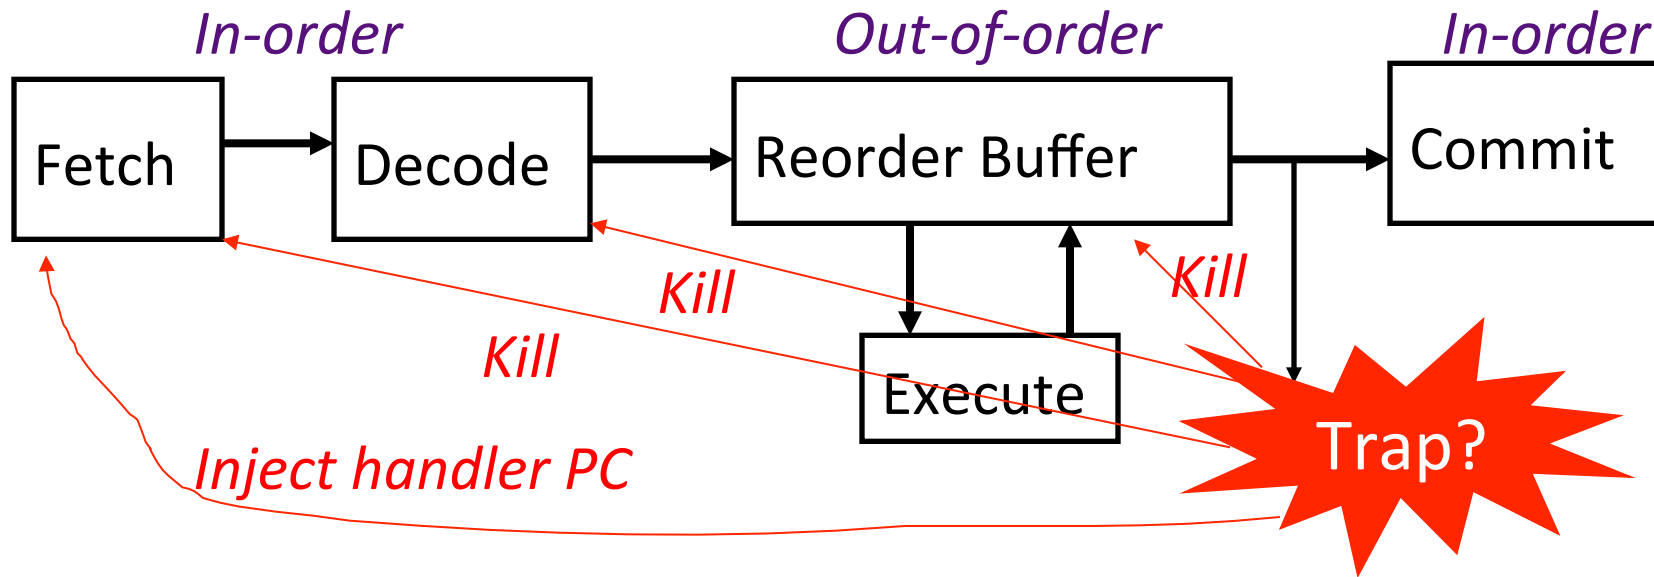
Out-of-order processing implemented commercially in 1960s, but disappeared again until 1990s as two major problems had to be solved:

- Precise traps
 - Imprecise *traps* complicate debugging and OS code
 - Note, precise *interrupts* are relatively easy to provide
- Branch prediction
 - Amount of exploitable instruction-level parallelism (ILP) limited by control hazards

Also, simpler machine designs in new technology beat complicated machines in old technology

- Big advantage to fit processor & caches on one chip
- Microprocessors had era of 1%/week performance scaling

In-Order Commit for Precise Traps

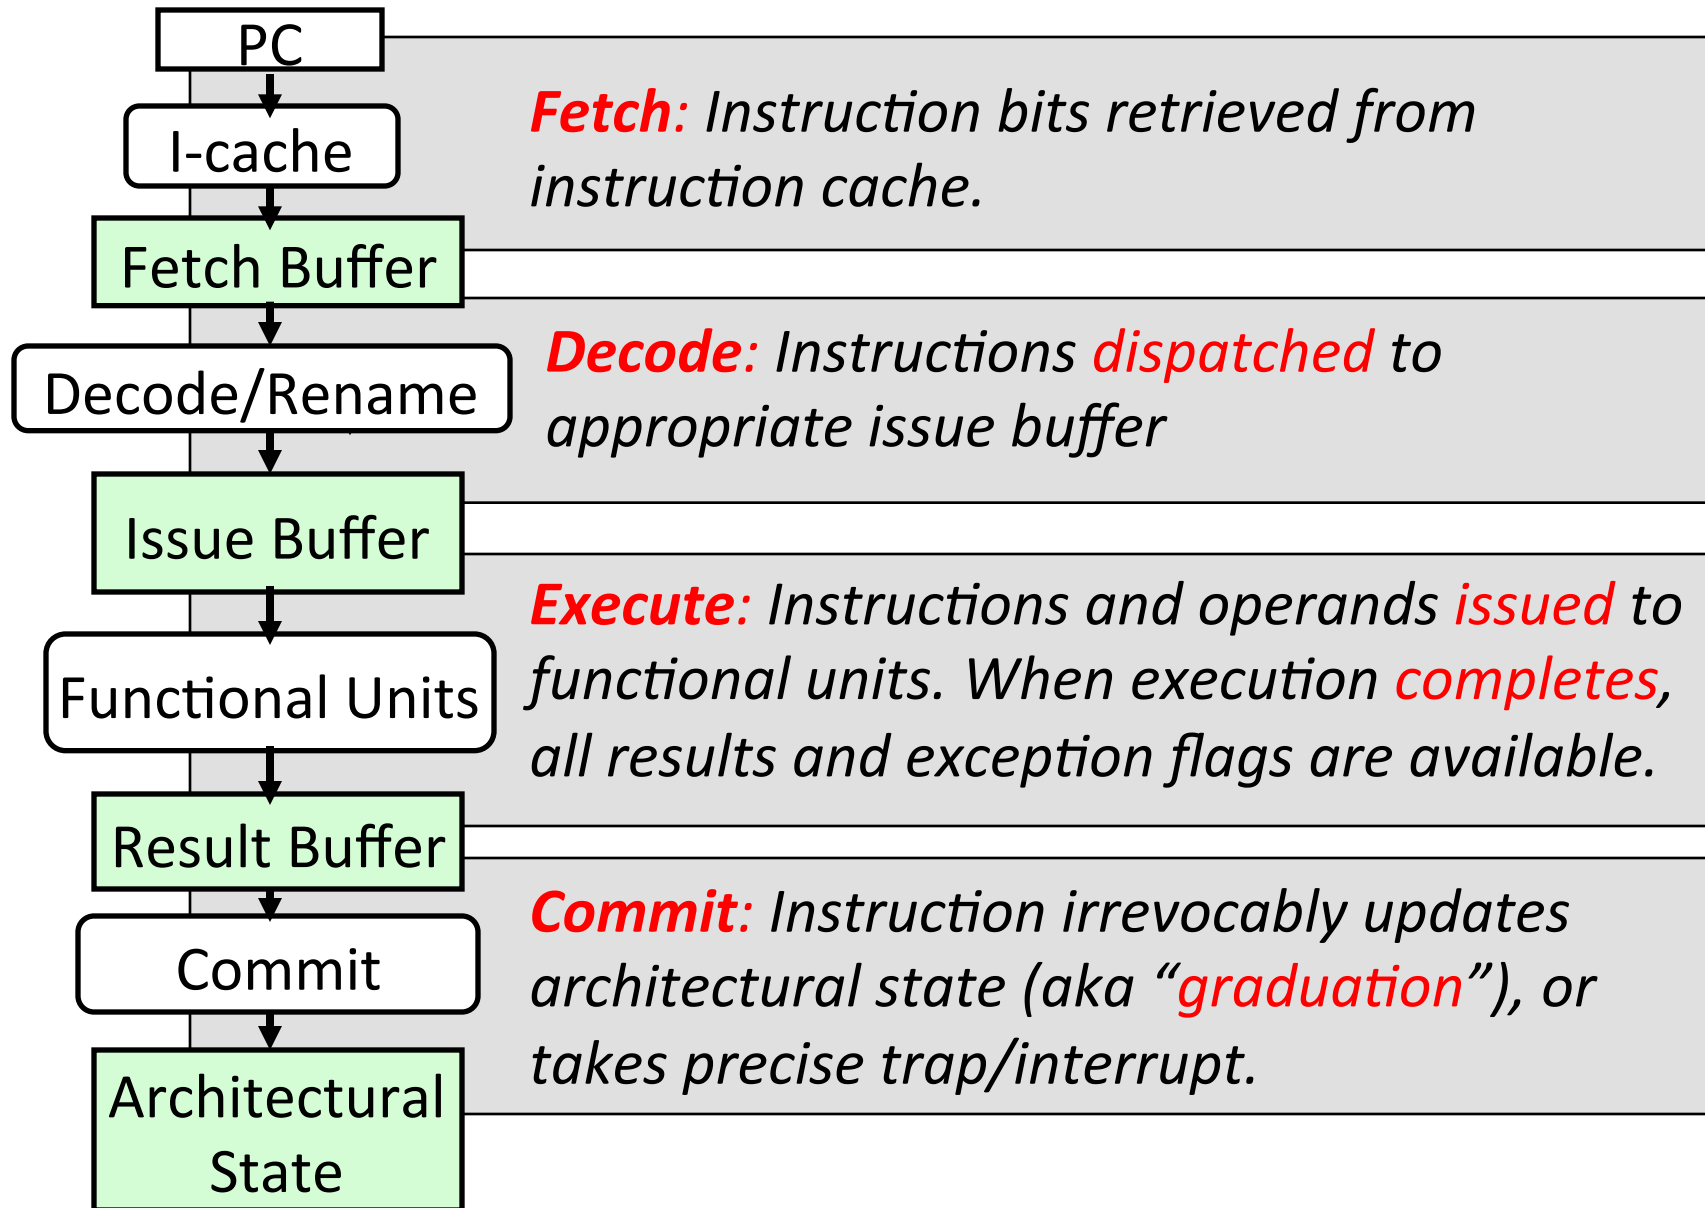


- In-order instruction fetch and decode, and dispatch to reservation stations inside reorder buffer
- Instructions issue from reservation stations out-of-order
- Out-of-order completion, values stored in temporary buffers
- Commit is in-order, checks for traps, and if none updates architectural state

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)
- Memory reordering needs special data structures
 - Speculative store address and data buffers
 - Speculative load address and data buffers

Phases of Instruction Execution



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”, but not in these lectures

In-Order Versus Out-of-Order Issue

- In-order issue:
 - Issue stalls on RAW dependencies or structural hazards, or possibly WAR/WAW hazards
 - Instruction cannot issue to execution units unless all preceding instructions have issued to execution units
- Out-of-order 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 the 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 5-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

Acknowledgements

- This course is partly inspired by previous MIT 6.823 and Berkeley CS252 computer architecture courses created by my collaborators and colleagues:
 - Arvind (MIT)
 - Joel Emer (Intel/MIT)
 - James Hoe (CMU)
 - John Kubiatowicz (UCB)
 - David Patterson (UCB)