

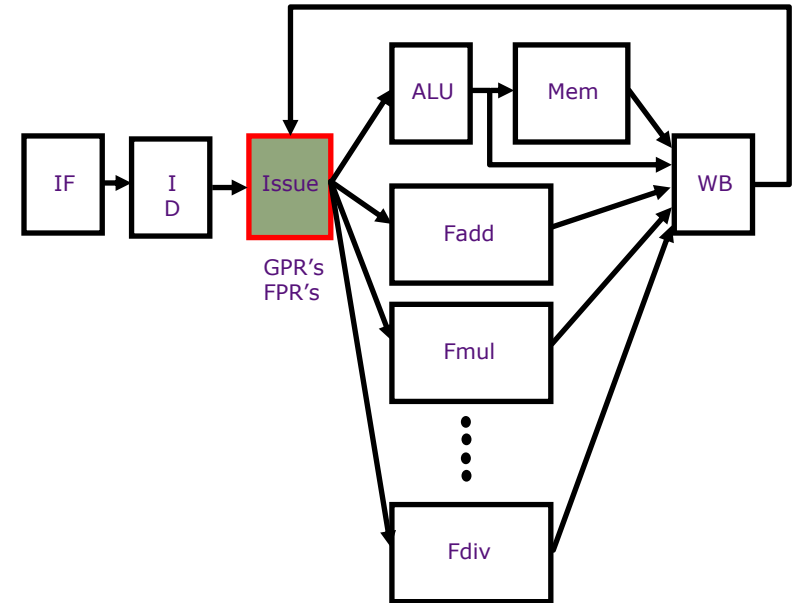
Complex Pipelining:
Out-of-Order Execution, Register
Renaming and Exceptions

Joel Emer

Computer Science and Artificial Intelligence Laboratory
M.I.T.

CDC 6600-style Scoreboard

- Instructions are issued in order;
An instruction is issued only if
- It cannot cause a RAW hazard
 ⇒ *if operands are read immediately then no need to remember sources of instructions in the execute phases*
 - It cannot cause a WAW hazard
 ⇒ *There can be at most instruction in the execute phase that can write in a particular register*




Scoreboard:
Two bit-vectors

Busy[FU#]: Indicates FU's availability
 These bits are hardwired to FU's.

WP[reg#]: Records if a write is pending for a register
 Set to true by the Issue stage and set to false by the WB stage

In-Order Issue Limitations: *an example*

					<i>latency</i>
1	LD	F2,	34(R2)		1
2	LD	F4,	45(R3)		<i>long</i>
3	MULTD	F6,	F4,	F2	3
4	SUBD	F8,	F2,	F2	1
5	DIVD	F4,	F2,	F8	4
6	ADDD	F10,	F6,	F4	1



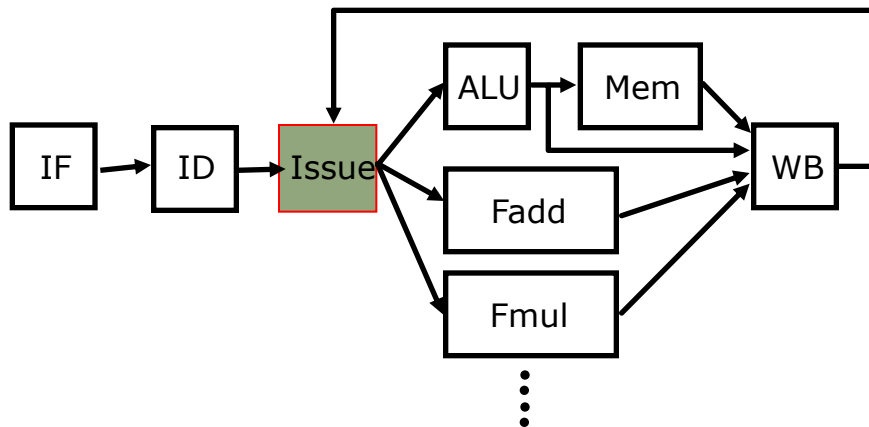
In-order:

1 (2,1) 2 3 4 4 3 5 5 6 6

In-order restriction prevents instruction 4 from being dispatched

Out-of-Order Issue

How can we address the delay caused by a RAW dependence associated with the next in-order instruction?




Find something else to do!

- 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.
- Can issue any instruction in buffer whose RAW hazards are satisfied (*for now at most one dispatch per cycle*). A writeback (WB) may enable more instructions.

In-Order Issue Limitations: *an example*

					<i>latency</i>
1	LD	F2,	34(R2)		1
2	LD	F4,	45(R3)		<i>long</i>
3	MULTD	F6,	F4,	F2	3
4	SUBD	F8,	F2,	F2	1
5	DIVD	F4,	F2,	F8	4
6	ADDD	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 feature of an ISA limits the number of instructions in the pipeline?

Number of Registers

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

Little's Law

Throughput (\bar{T}) = Number in Flight (\bar{N}) / Latency (\bar{L})



Example:

4 floating point registers

8 cycles per floating point operation

⇒ 1/2 issues per cycle!

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 ?

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

Instruction-level Parallelism via *Renaming*

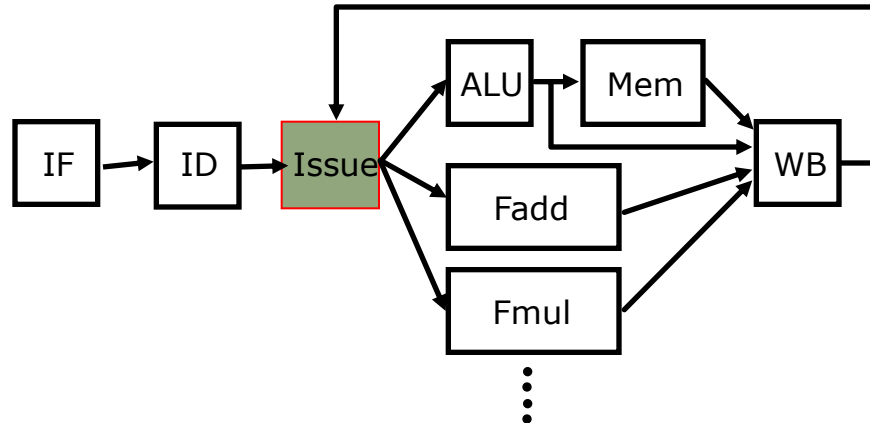
					<i>latency</i>
1	LD	F2,	34(R2)		1
2	LD	F4,	45(R3)		<i>long</i>
3	MULTD	F6,	F4,	F2	3
4	SUBD	F8,	F2,	F2	1
5	DIVD	F4' ,	F2,	F8	4
6	ADDD	F10,	F6,	F4'	1



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

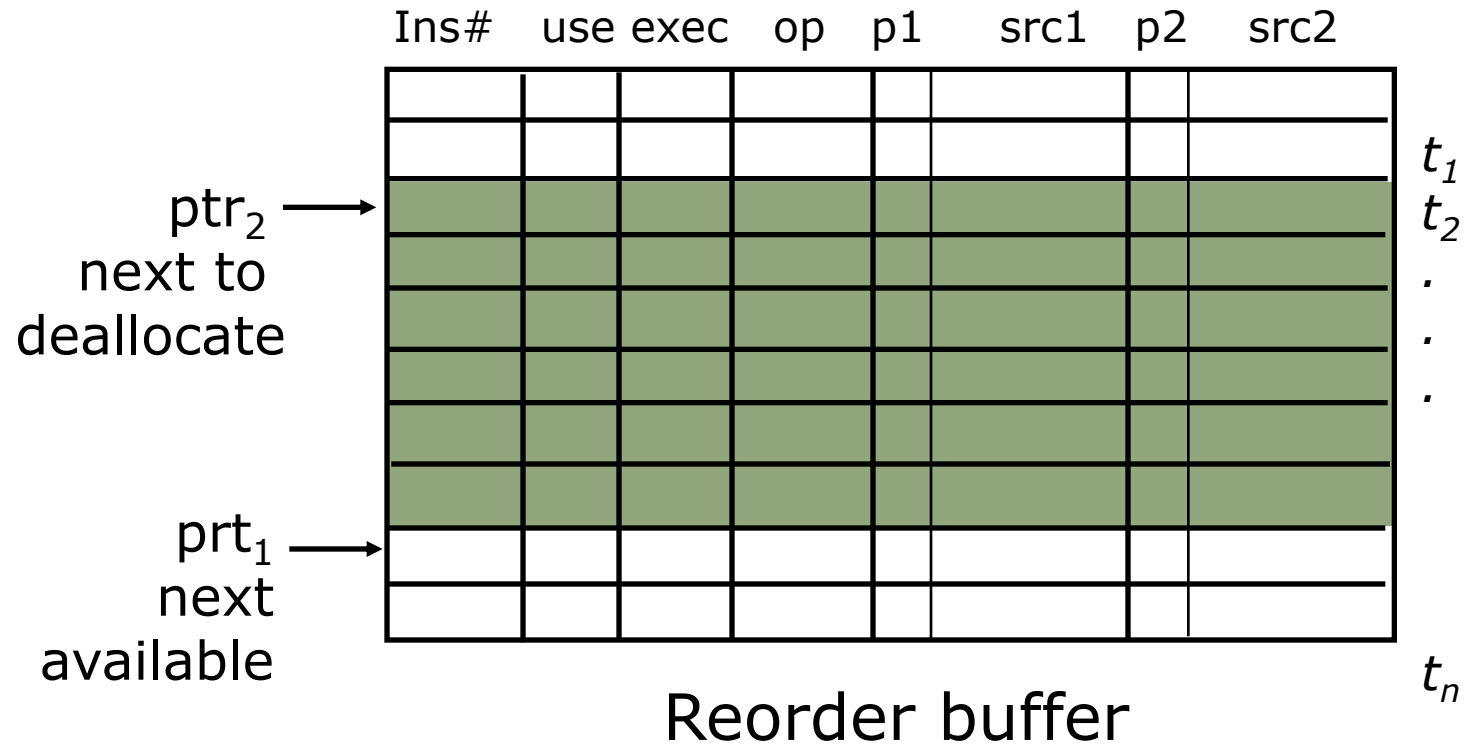
*Renaming eliminates WAR and WAW hazards
 (renaming \Rightarrow additional storage)*

Handling register dependencies



- Decode does register renaming, providing a new spot for each register write
 - ⇒ Renaming eliminates structural hazards (WAR and WAW) by allowing use of more storage space.
- Renamed instructions added to an issue stage structure, called the reorder buffer (ROB). Any instruction in ROB whose RAW hazards have been satisfied can be dispatched.
 - ⇒ Out-of-order or dataflow execution handles RAW hazards

Dataflow execution



Instruction slot is candidate for execution when:

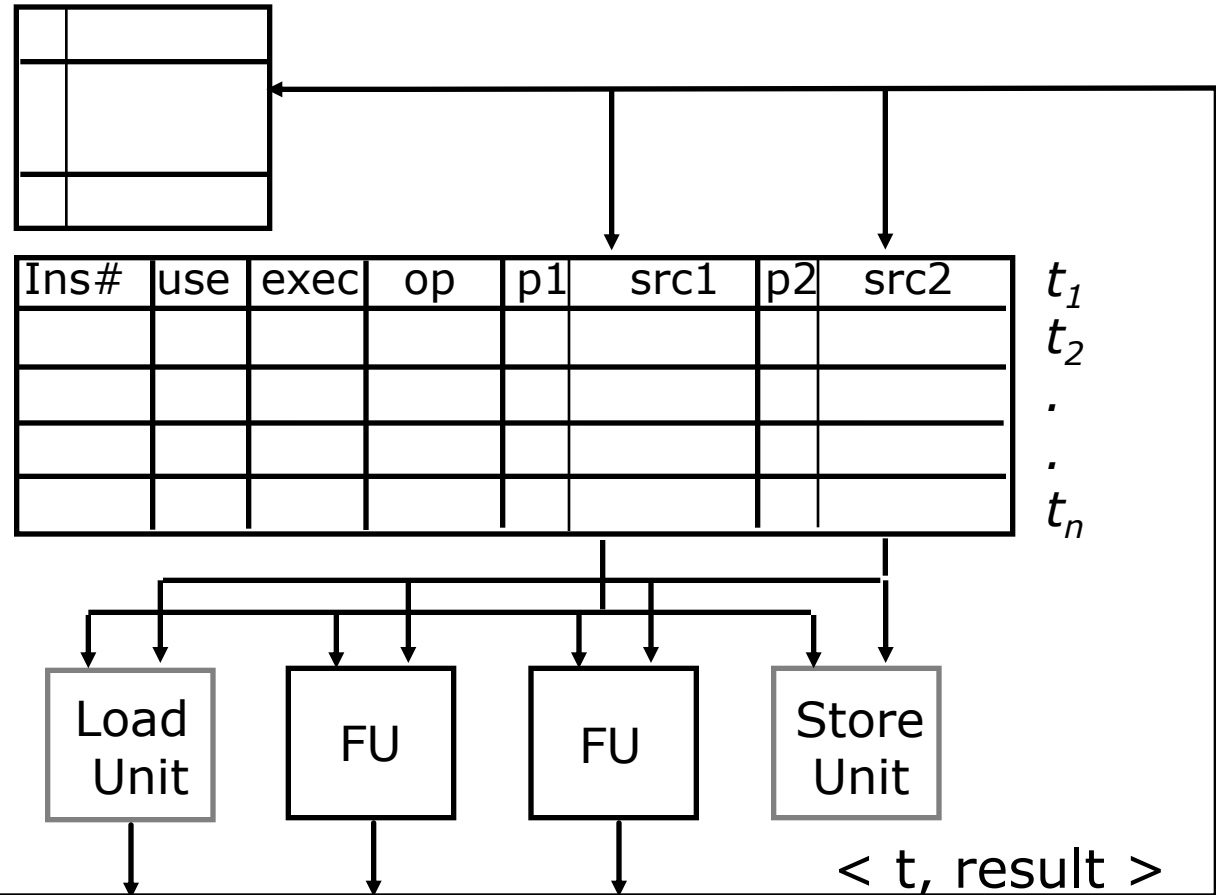
- It holds a valid instruction ("use" bit is set)
- It has not already started execution ("exec" bit is clear)
- Both operands are available (p1 and p2 are set)

Data-Driven Execution

*Renaming
table &
reg file*

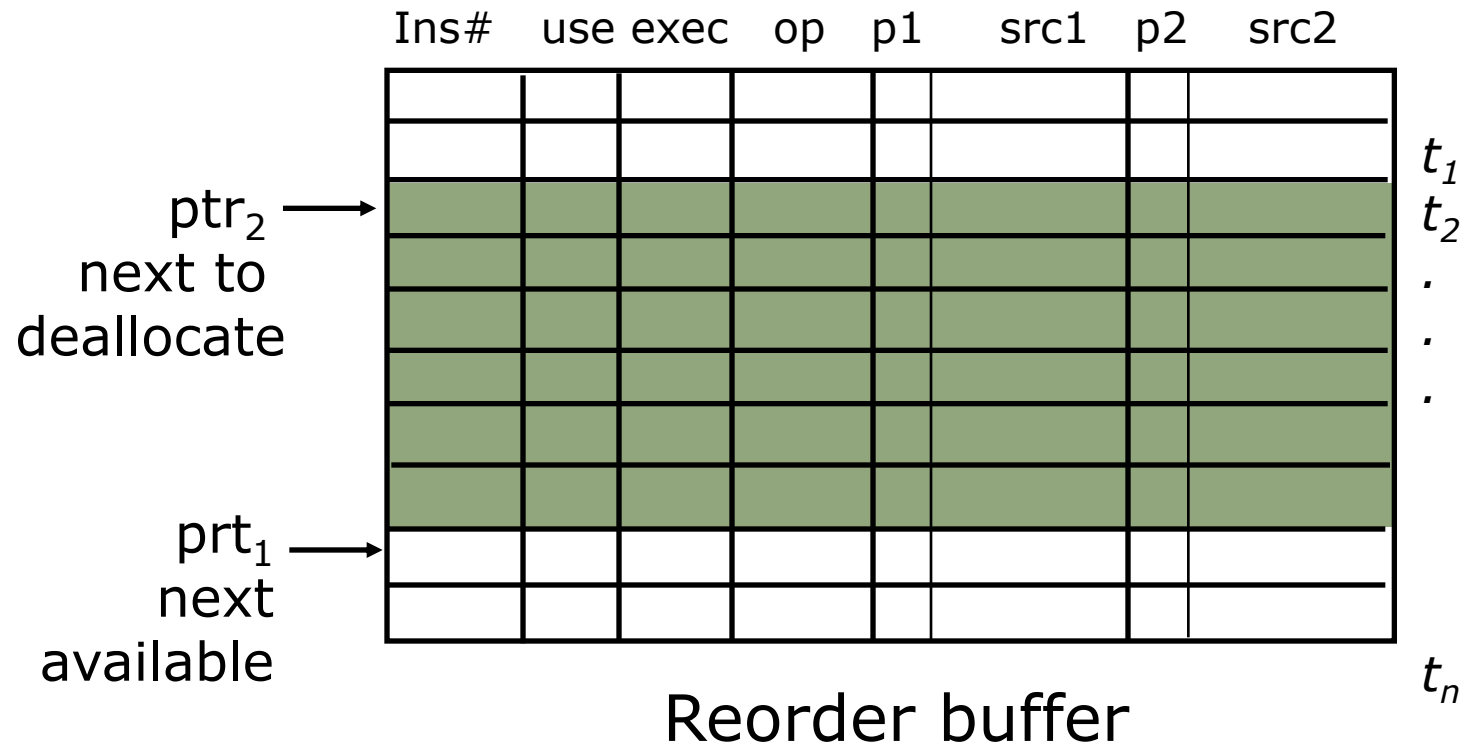
*Reorder
buffer*

Replacing the
tag by its value
is an expensive
operation



- Instruction template (i.e., tag t) is allocated by the Decode stage, which also stores the tag in the reg file
- When an instruction completes, its tag is deallocated

Simplifying Allocation/Deallocation

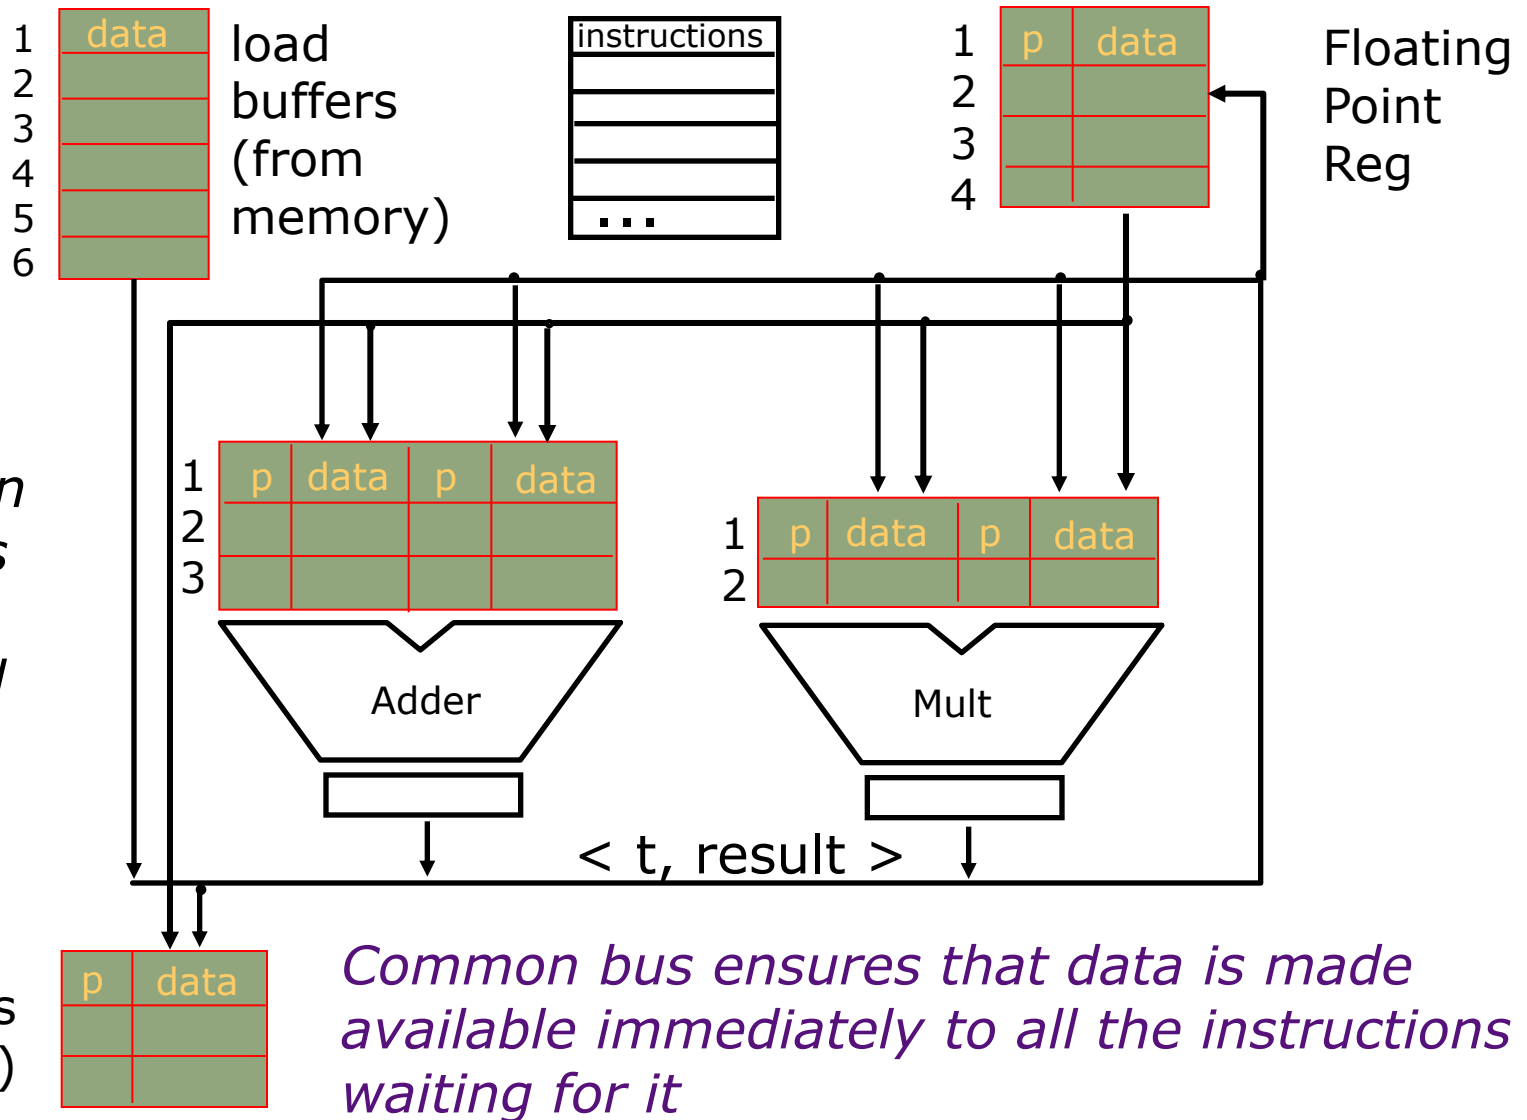


Instruction buffer is managed circularly

- "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

IBM 360/91 Floating Point Unit

R. M. Tomasulo, 1967



Effectiveness?

Renaming and Out-of-order execution was first implemented in 1969 in IBM 360/91 but was effective only on a very small class of problems and thus did not show up in the subsequent models until mid-nineties.

Why?

1. Did not address the memory latency problem which turned out to be a much bigger issue than FU latency
2. Made exceptions imprecise

One more problem needed to be solved

Control transfers

More on this in the next lecture

Precise Exceptions

Exceptions are relatively unlikely events that need special processing, but where adding explicit control flow instructions is not desired, e.g., divide by 0, page fault

Exceptions can be viewed as an implicit conditional subroutine call that is inserted between two instructions.

Therefore, it must appear as if the exception is taken between two instructions (say I_i and I_{i+1})


- the effect of all instructions up to and including I_i is complete
- no effect of any instruction after I_i has taken place

The handler either aborts the program or restarts it at I_{i+1} .

Effect on Exceptions

Out-of-order Completion

I_1	DIVD	f6,	f6,	f4
I_2	LD	f2,	45(r3)	
I_3	MULTD	f0,	f2,	f4
I_4	DIVD	f8,	f6,	f2
I_5	SUBD	f10,	f0,	f6
I_6	ADDD	f6,	f8,	f2

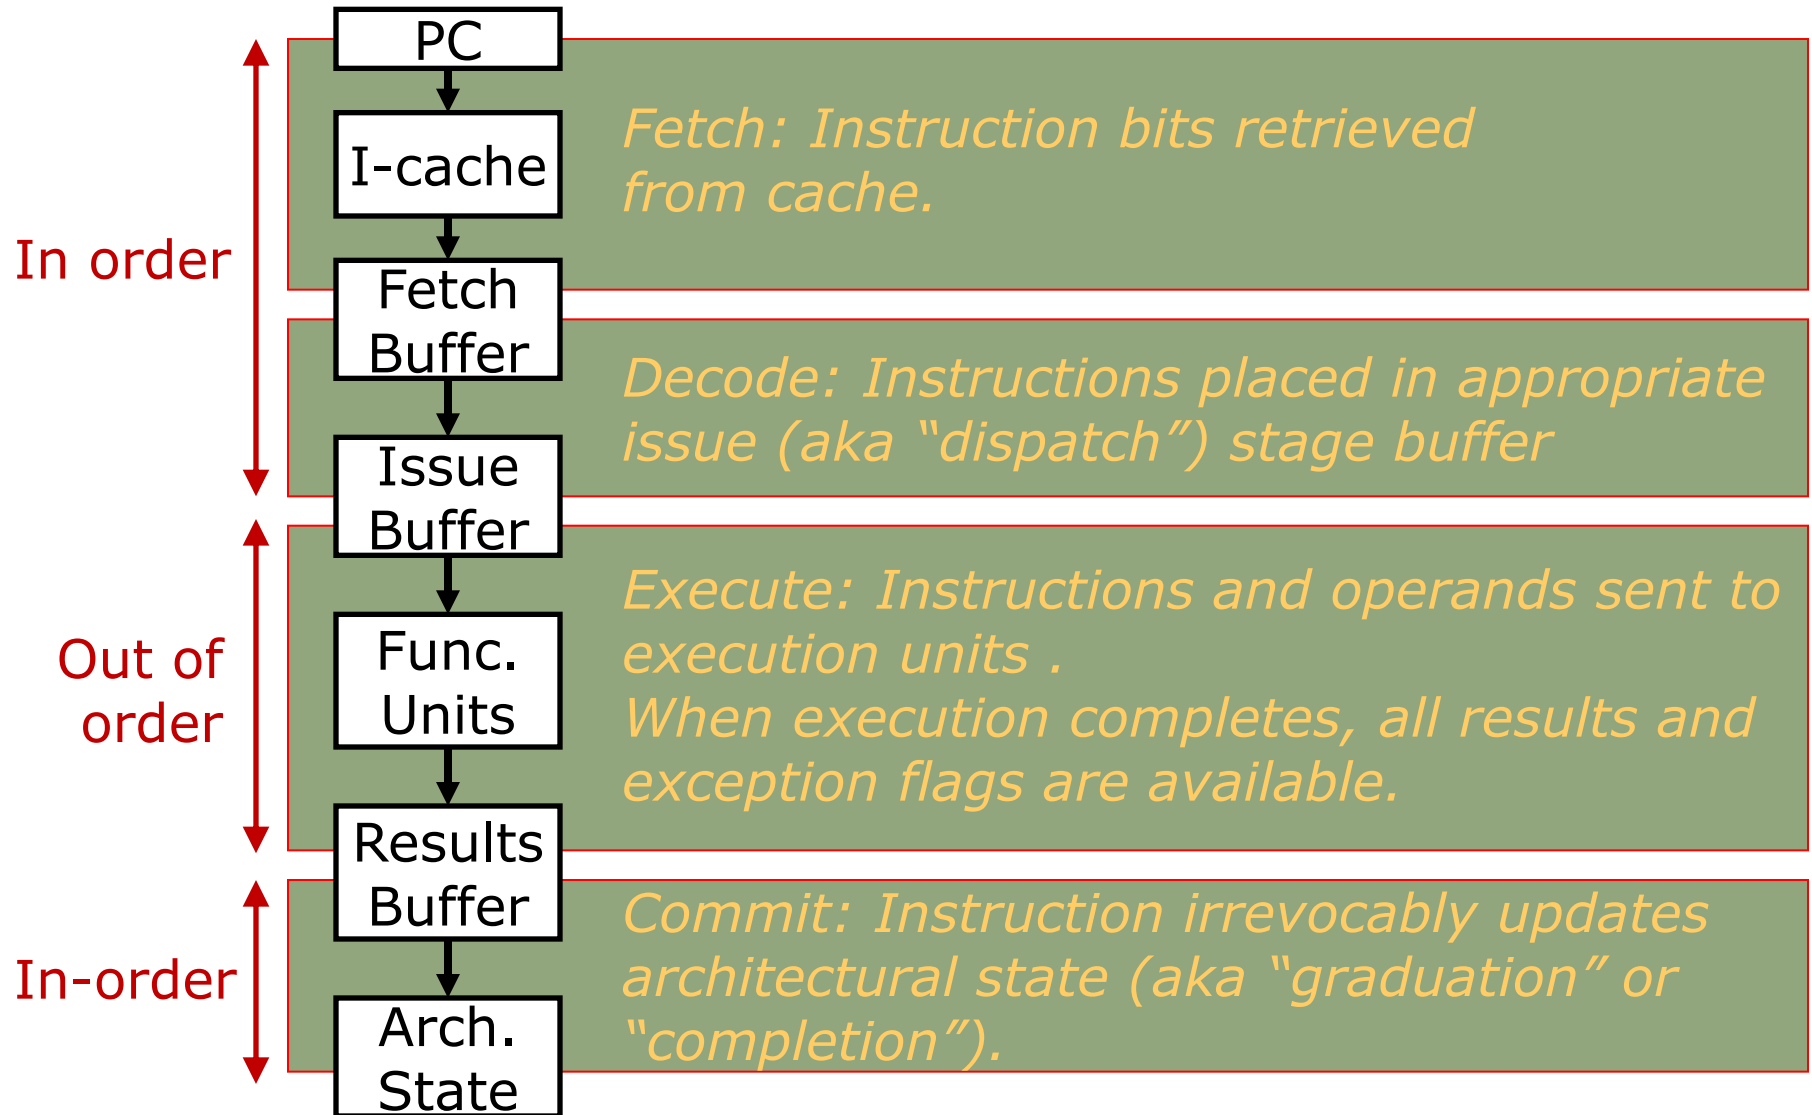


Precise exceptions are difficult to implement at high speed
- want to start execution of later instructions before
exception checks finished on earlier instructions

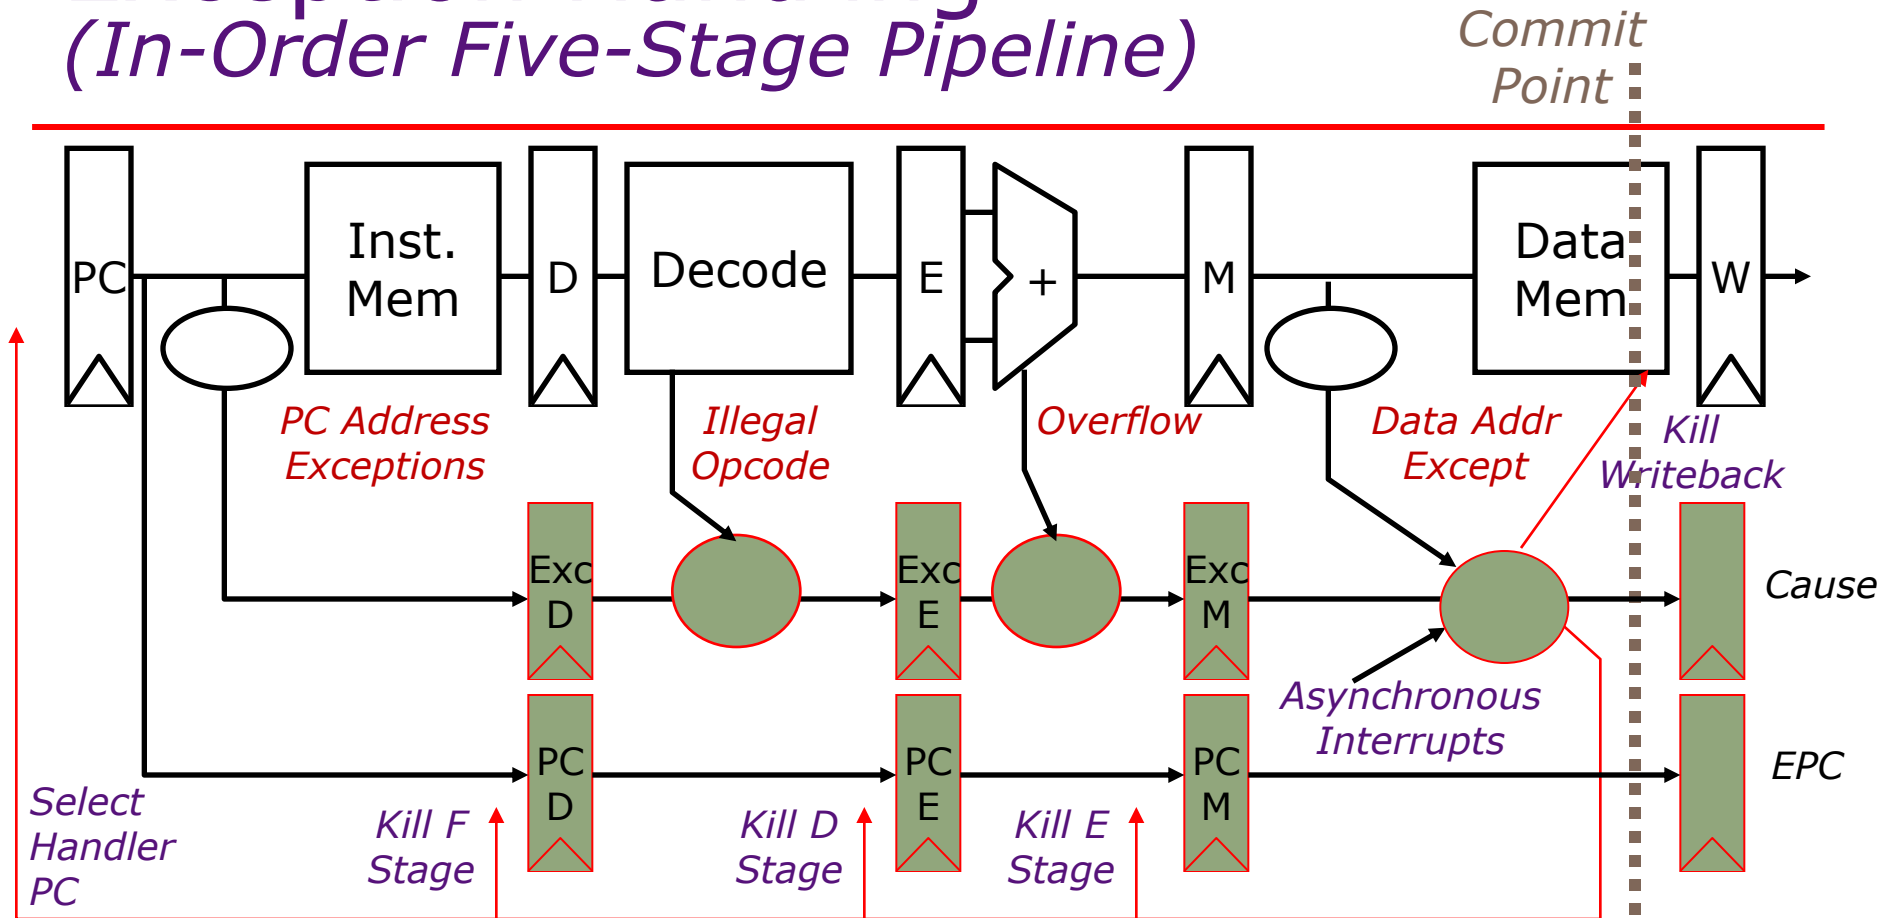
Exceptions

- Exceptions create a dependence on the value of the next PC
 - Options for handling this dependence:
 - Stall No
 - Bypass No
 - Find something else to do No
 - Change the architecture Sometimes: Alpha, Multiflow
 - Speculate! Most common approach!
 - How can we handle rollback on mis-speculation
- Delay state update until commit on speculated instructions
- Note: earlier exceptions must override later ones

Phases of Instruction Execution



Exception Handling (In-Order Five-Stage Pipeline)



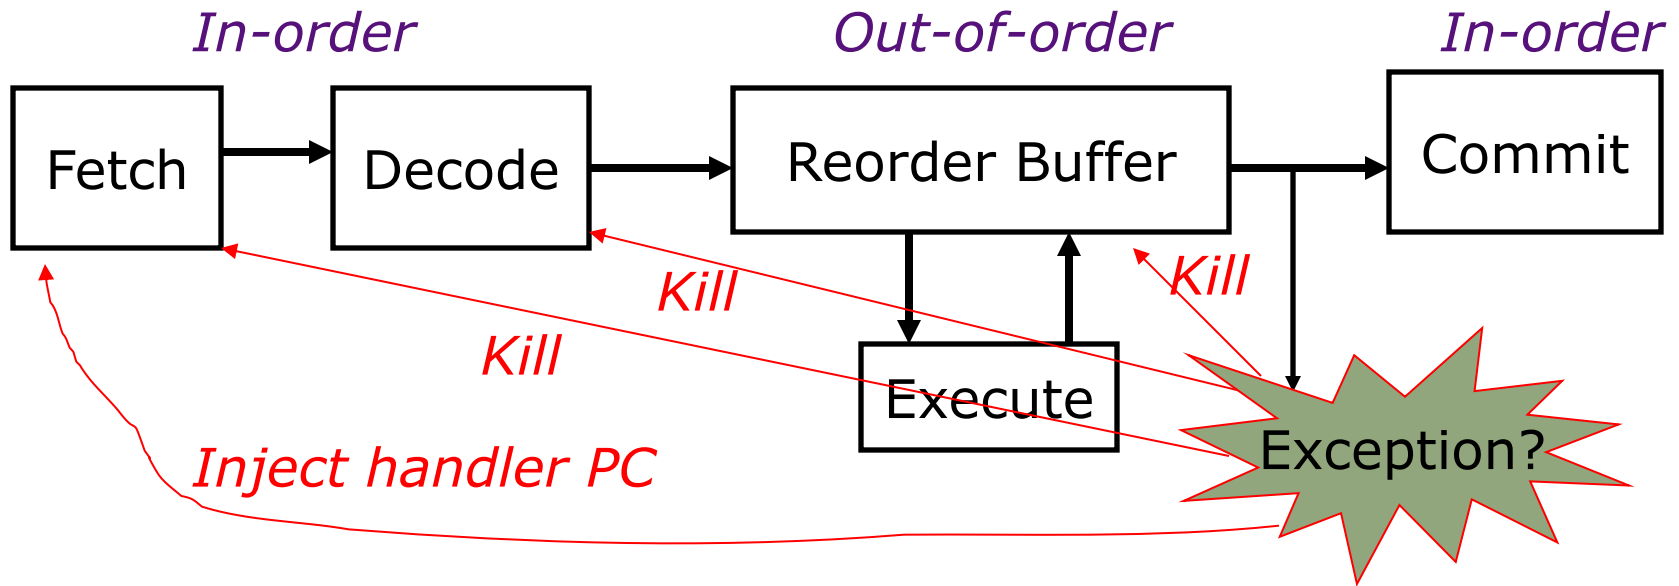
Hold exception flags in pipeline until commit point (M stage)

• If exception at commit:

- update Cause/EPC registers
- kill all stages
- fetch at handler PC

Inject external interrupts at commit point

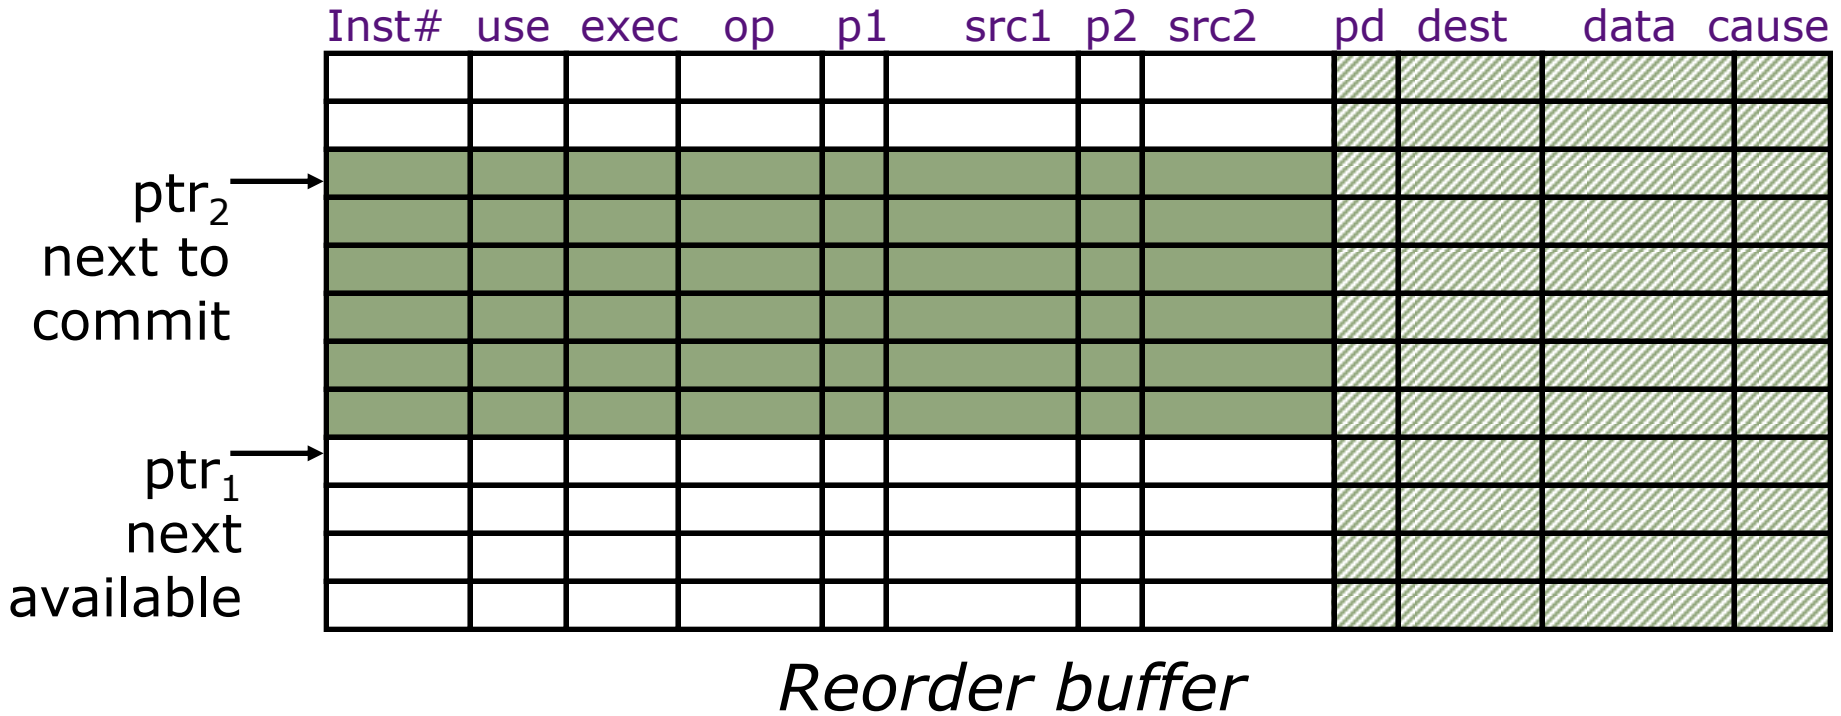
In-Order Commit for Precise Exceptions



- Instructions fetched and decoded into instruction reorder buffer in-order
- Execution is out-of-order (\Rightarrow out-of-order completion)
- *Commit* (write-back to architectural state, i.e., regfile & memory, is in-order)

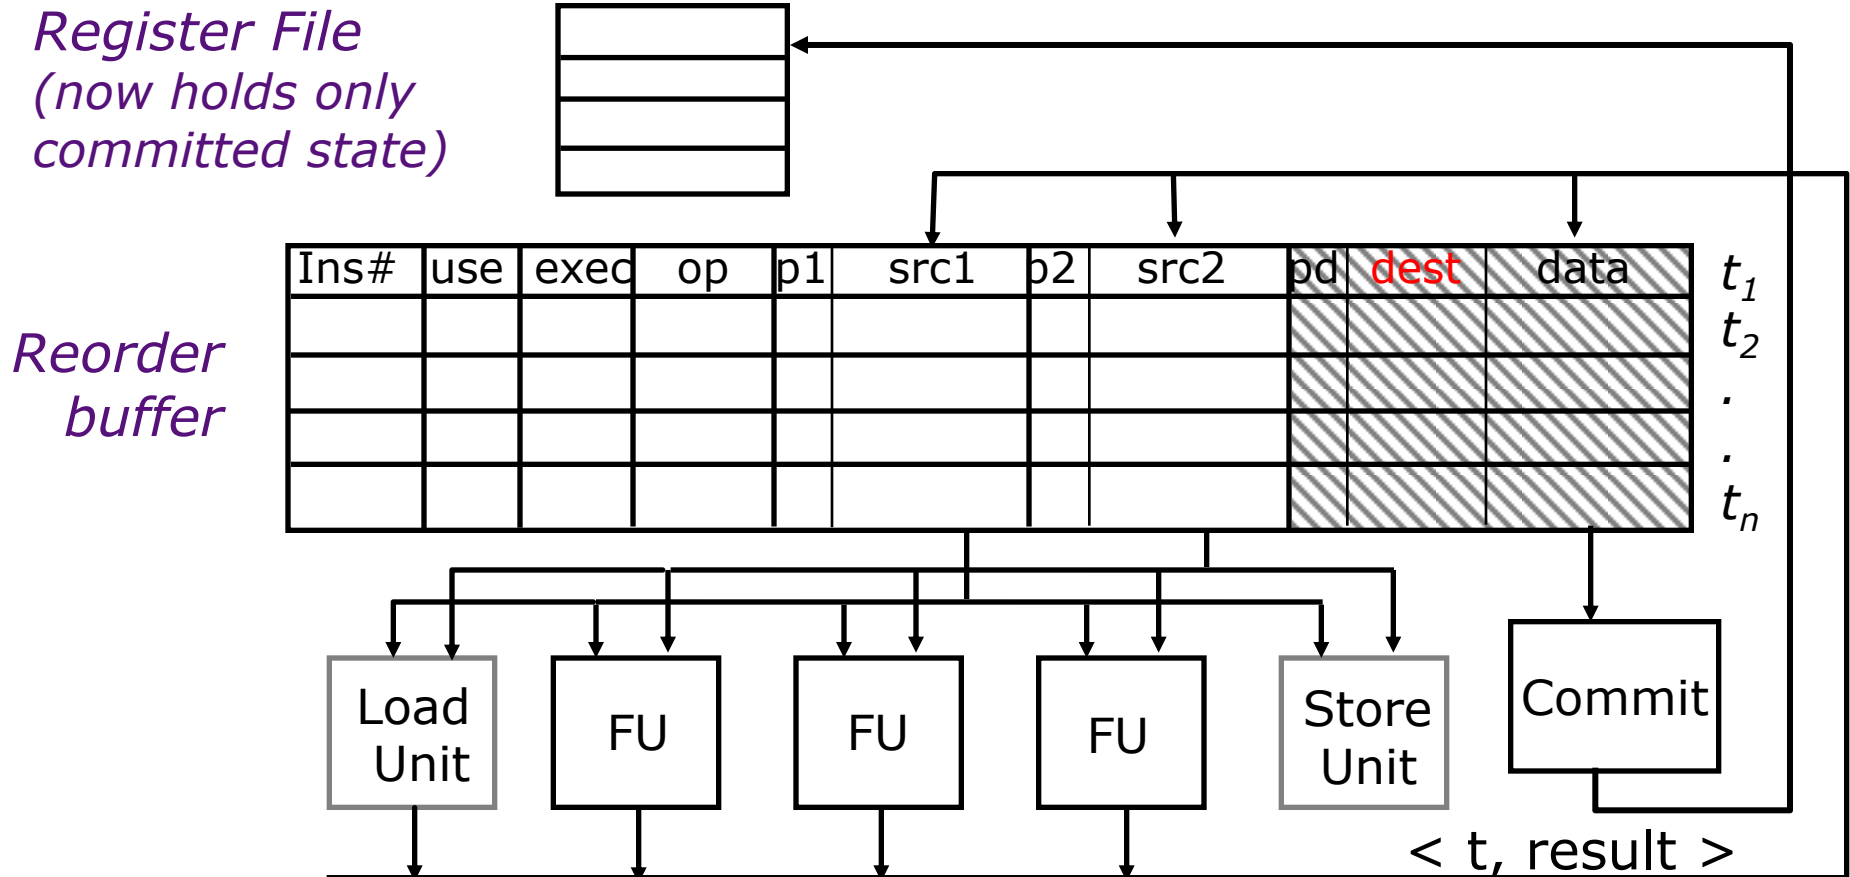
Temporary storage needed to hold results before commit (shadow registers and store buffers)

Extensions for Precise Exceptions



- add $\langle \text{pd}, \text{dest}, \text{data}, \text{cause} \rangle$ fields in the instruction template
- commit instructions to reg file and memory in program order \Rightarrow buffers can be maintained circularly
- on exception, clear reorder buffer by resetting $\text{ptr}_1 = \text{ptr}_2$
(stores must wait for commit before updating memory)

Rollback and Renaming

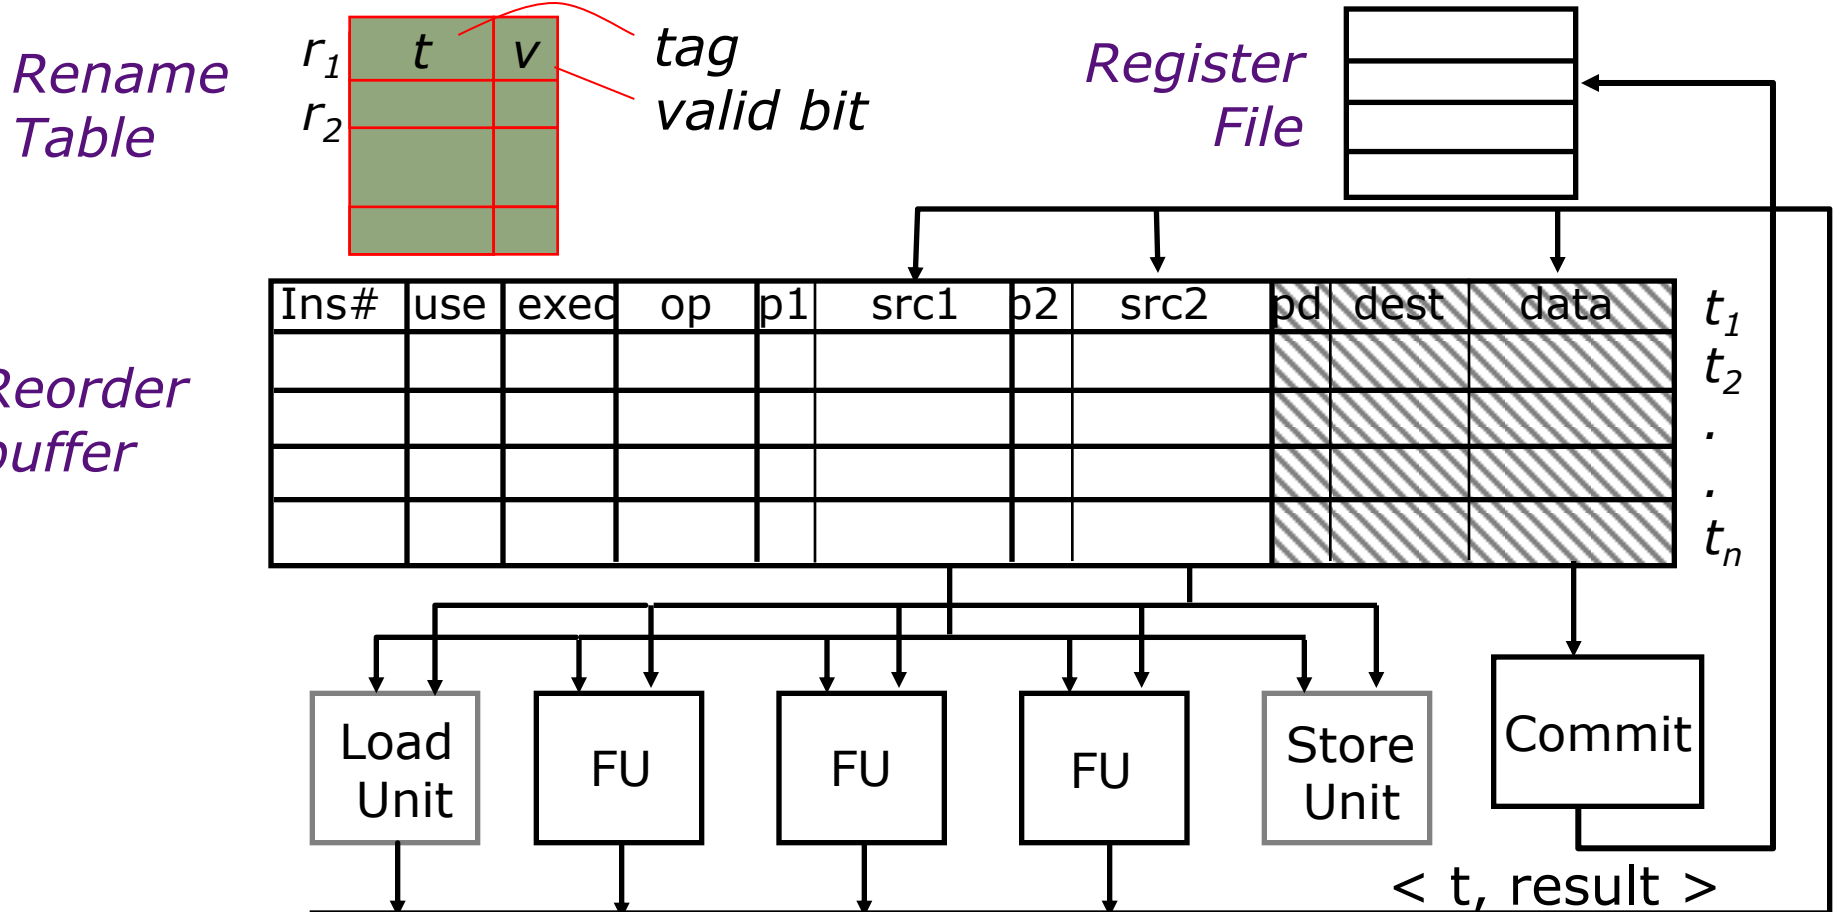


Register file does not contain renaming tags any more.

How does the decode stage find the tag of a source register?

Search the "dest" field in the reorder buffer

Renaming Table



Renaming table is a cache to speed up register name lookup.
It needs to be cleared after each exception taken.

When else are valid bits cleared?

Control transfers

Physical Register Files

- Reorder buffers are space inefficient – a data value may be stored in multiple places in the reorder buffer
- idea – keep all data values in a physical register file
 - Tag represents the name of the data value and name of the physical register that holds it
 - Reorder buffer contains only tags

Thus, 64 data values may be replaced by 8-bit tags for a 256 element physical register file

More on this in later lectures ...

Branch Penalty

Next fetch started

How many instructions need to be killed on a misprediction?

Modern processors may have > 10 pipeline stages between nextPC calculation and branch resolution !

Branch executed

next lecture:
Branch prediction &
Speculative execution

