




Complex Pipelining: Out-of-Order Execution & Register Renaming

Arvind

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

*Based on the material prepared by
Krste Asanovic and Arvind*

In-Order Issue Pipeline



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 to true by the Issue stage and set to false by the WB stage

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

FU available?

RAW?

WAR?

WAW?

Busy[FU#]


WP[src1] or WP[src2]

cannot arise

WP[dest]

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



- 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.
- 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.

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

Number of Registers

Which features of a program limit the number of instructions in the pipeline?

Control transfers

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

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 based on on-the-fly *register renaming*

Instruction-Level Parallelism with Renaming

					latency
1	LD	F2,	34(R2)		1
2	LD	F4,	45(R3)		long
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

Any antidependence can be eliminated by renaming.

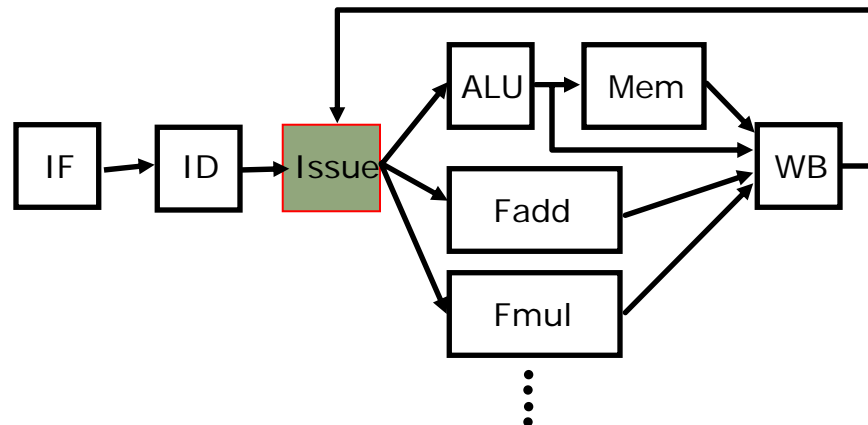
(renaming \Rightarrow additional storage)

Can it be done in hardware?

yes!



Register Renaming



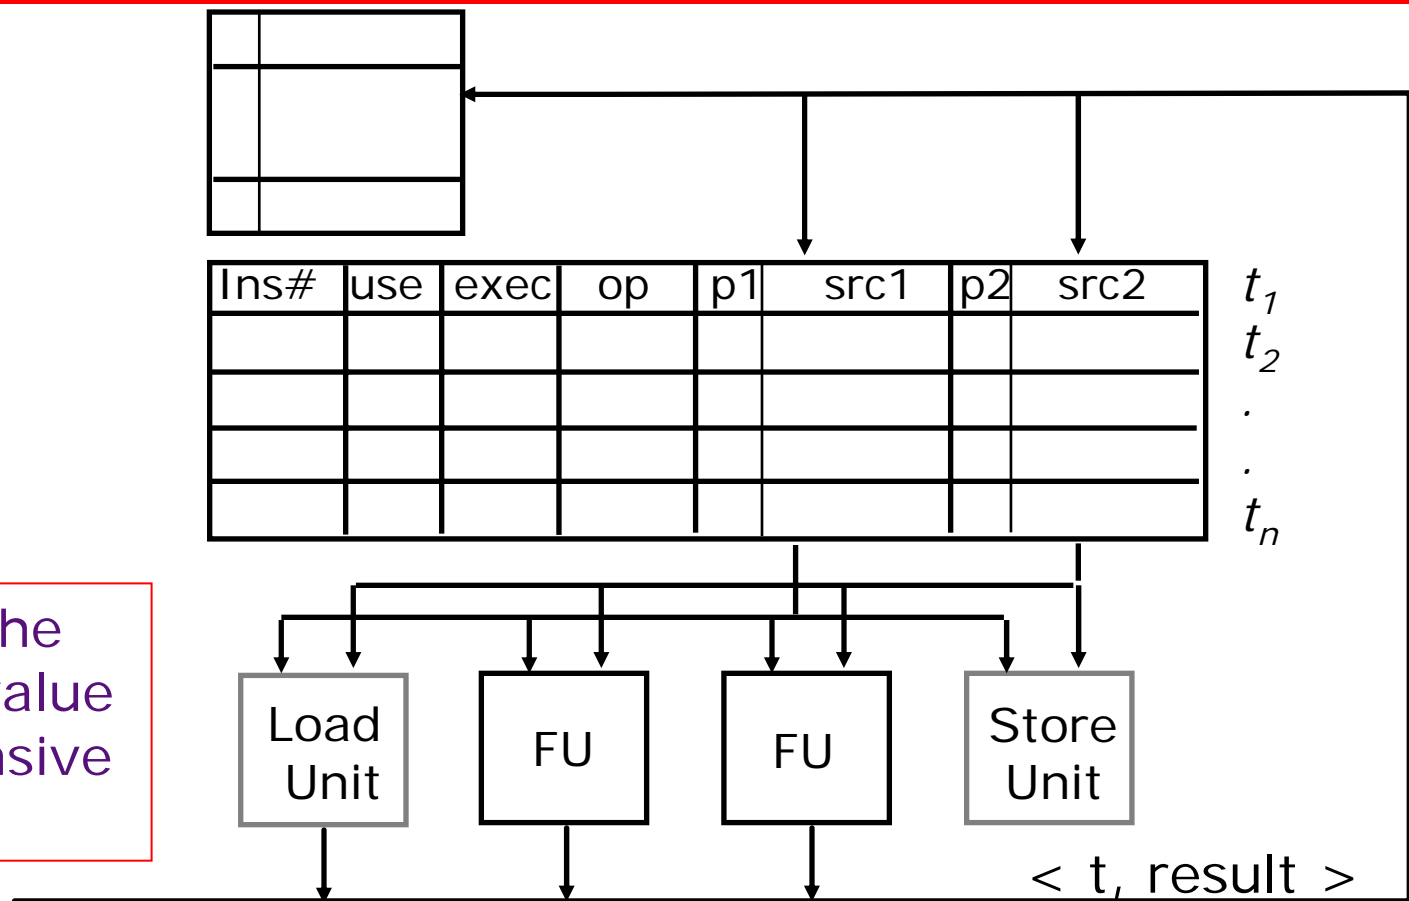
- Decode does register renaming and adds instructions to the issue stage reorder buffer (ROB)
 - ⇒ renaming makes WAR or WAW hazards impossible
- Any instruction in ROB whose RAW hazards have been satisfied can be dispatched.
 - ⇒ Out-of-order or dataflow execution

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

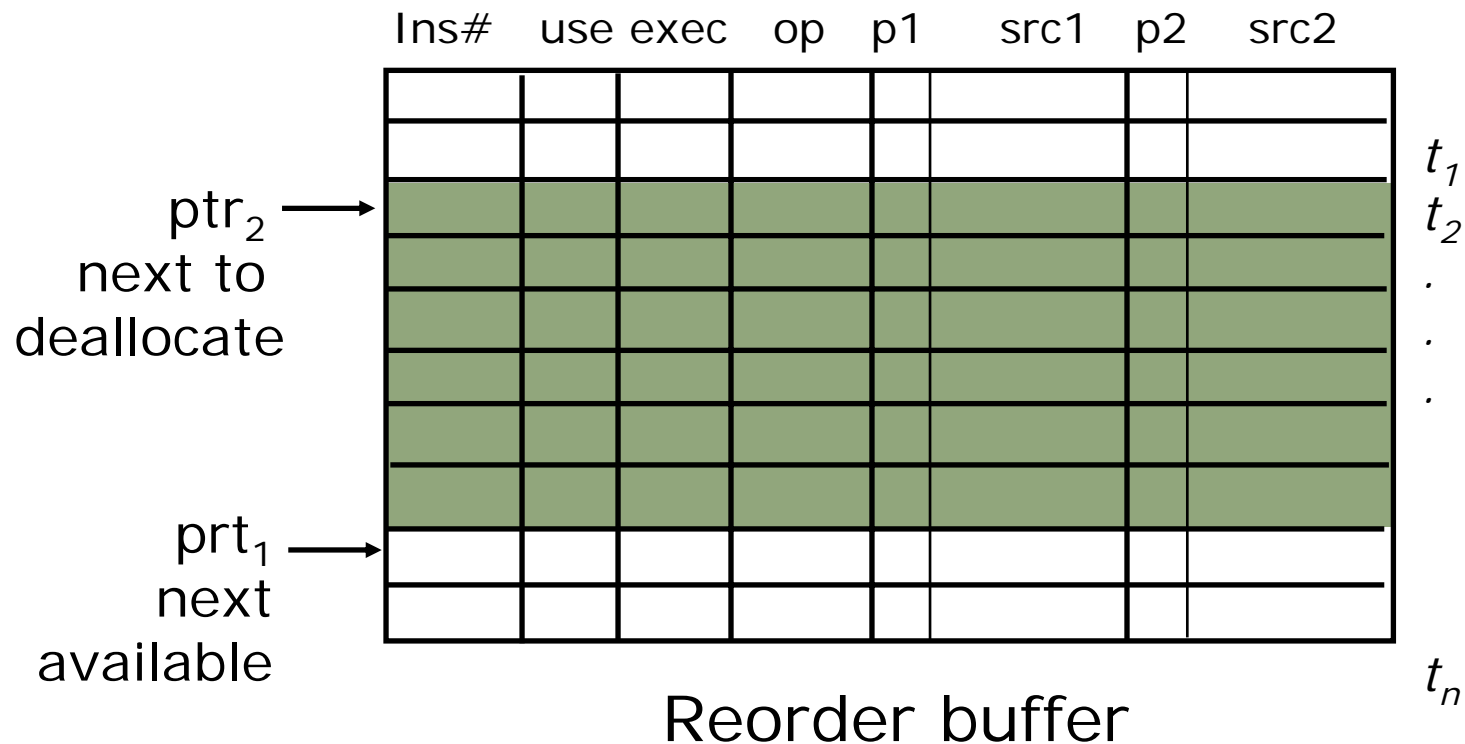
Dataflow execution



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)

Simplifying Allocation/Deallocation

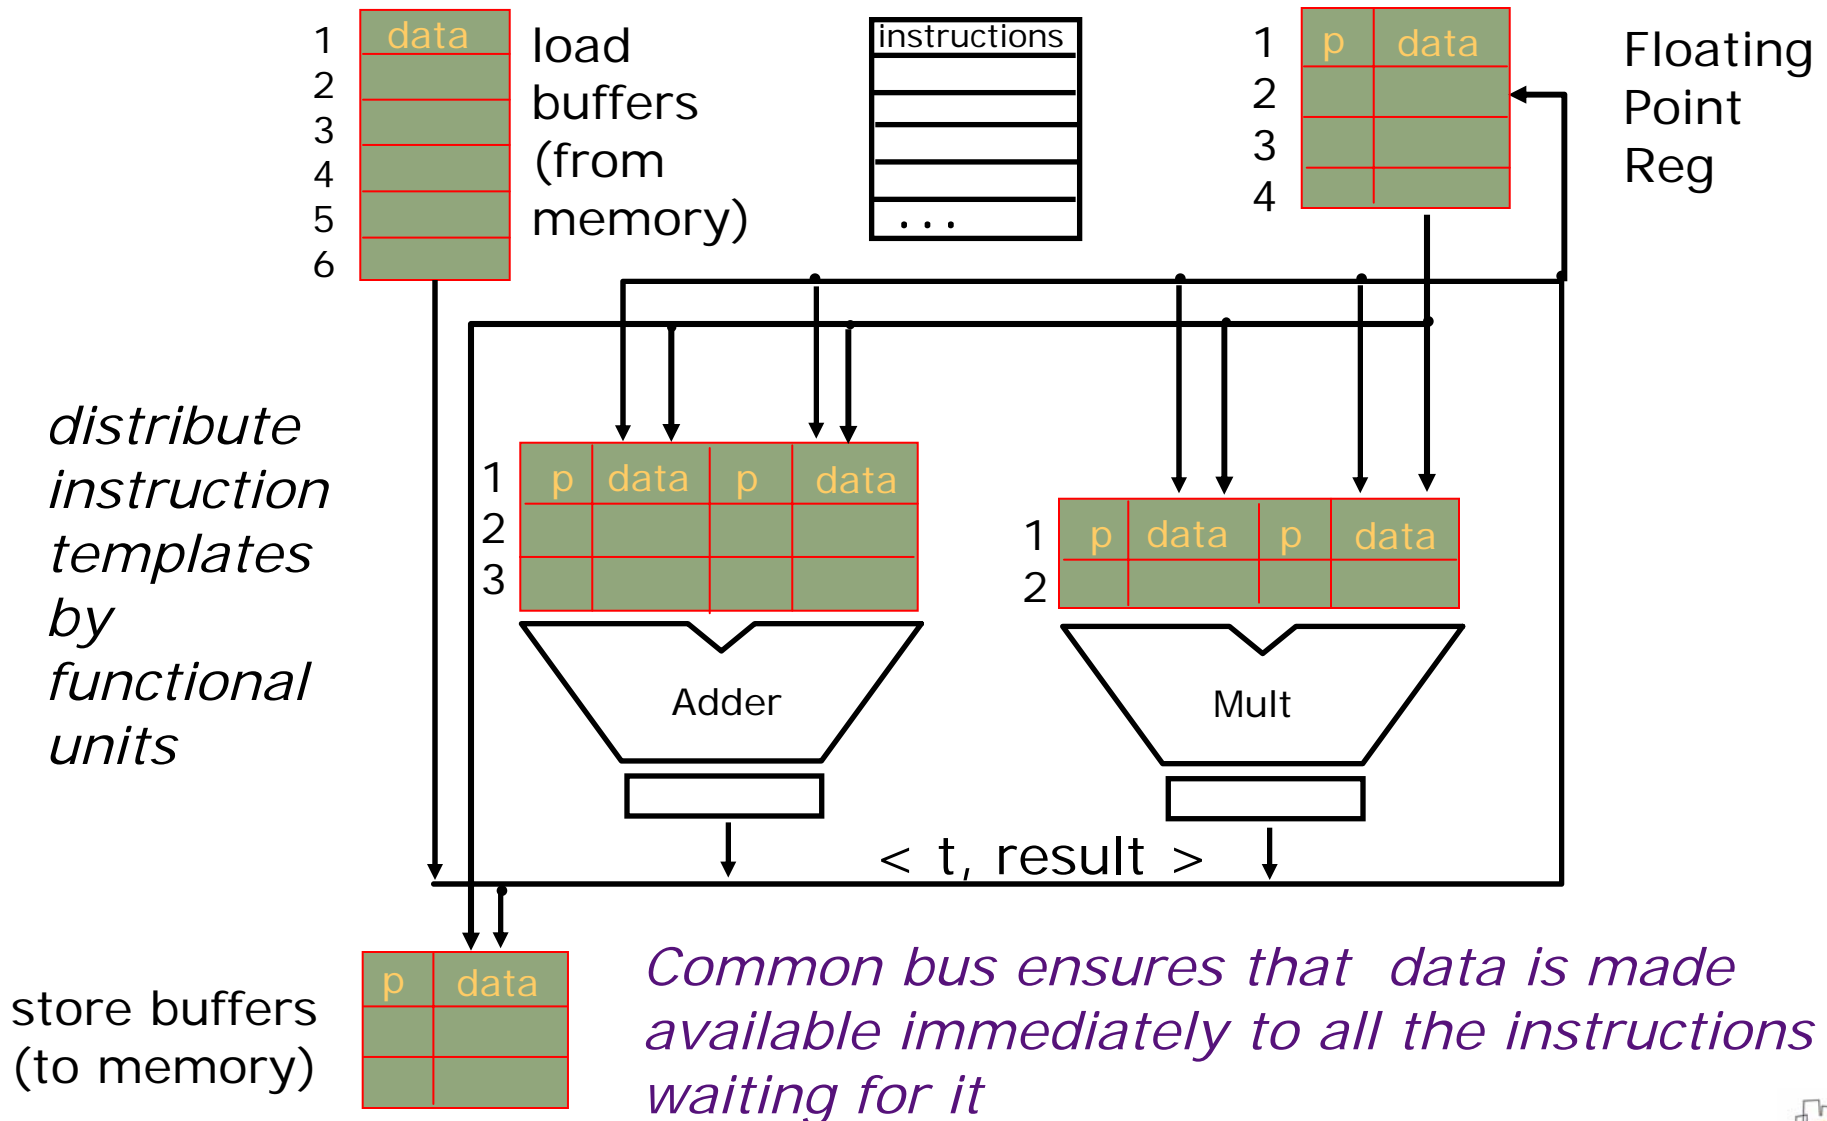


Instruction buffer is managed circularly

- “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 did not show up in the subsequent models until mid-Nineties.

Why ?

Reasons

1. Exceptions not precise!
2. Effective on a very small class of programs

One more problem needed to be solved

Control transfers



Five-minute break to stretch your legs

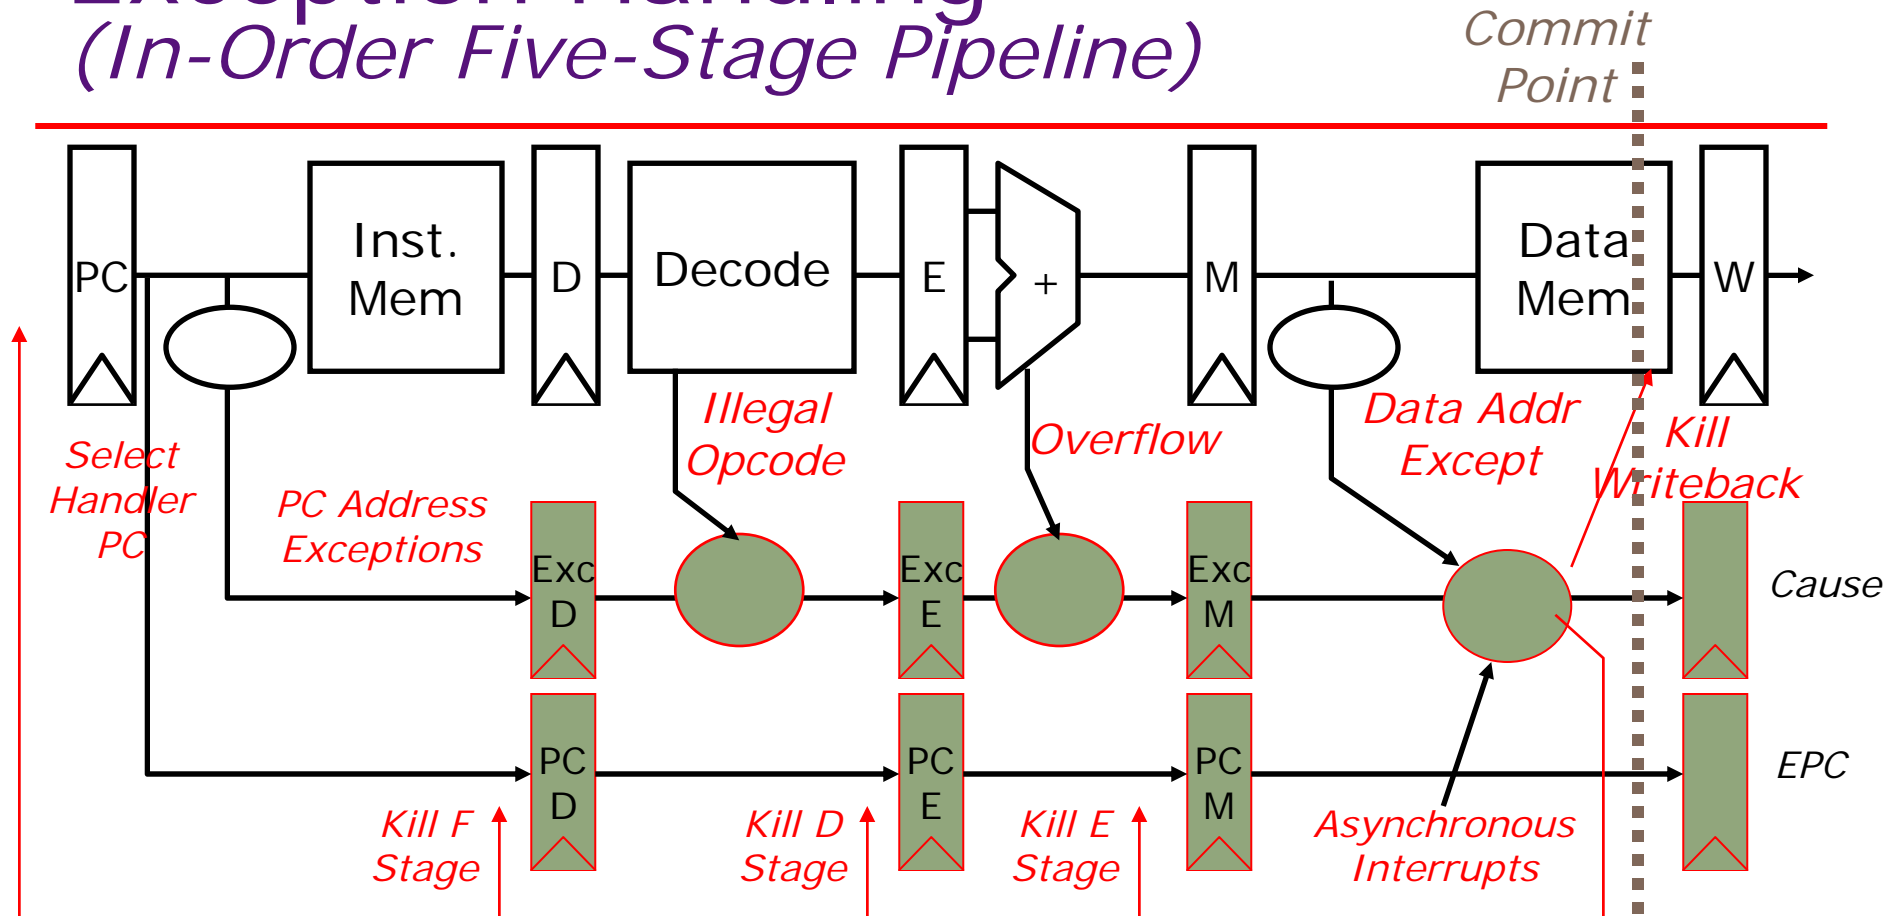
Precise Interrupts

It must appear as if an interrupt is taken between two instructions (say I_i and I_{i+1})

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

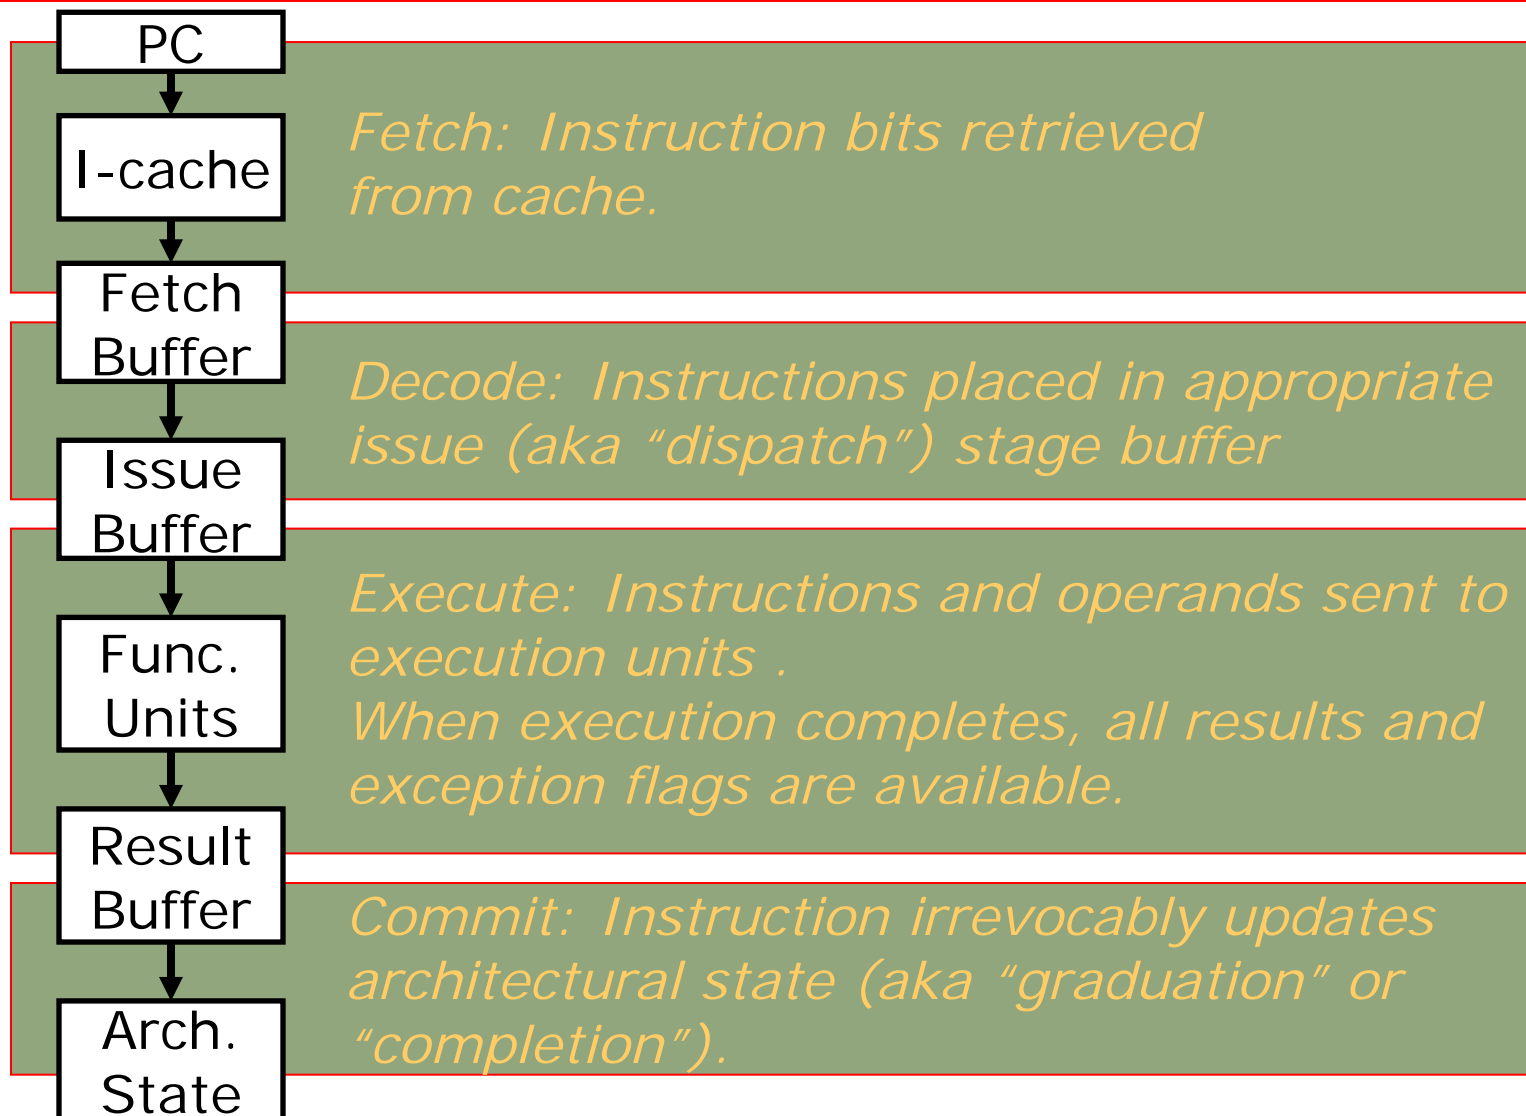
The interrupt handler either aborts the program or restarts it at I_{i+1} .

Exception Handling (In-Order Five-Stage Pipeline)




- Hold exception flags in pipeline until commit point (M stage)
- Exceptions in earlier pipe stages override later exceptions
- Inject external interrupts at commit point (override others)
- If exception at commit: update Cause and EPC registers, kill all stages, inject handler PC into fetch stage

Phases of Instruction Execution



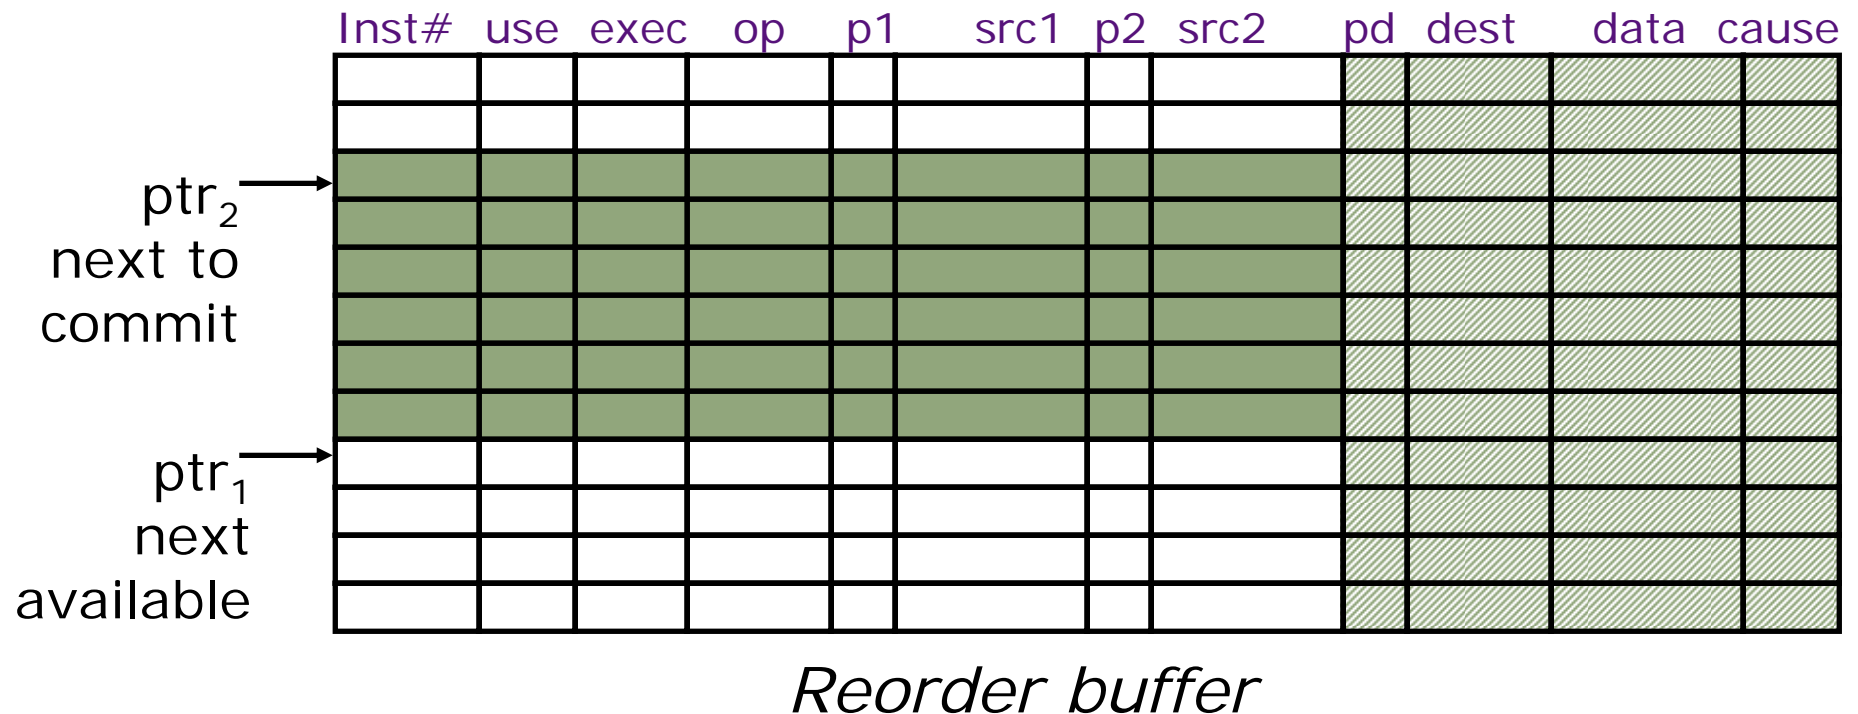
In-Order Commit for Precise Exceptions



- 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)

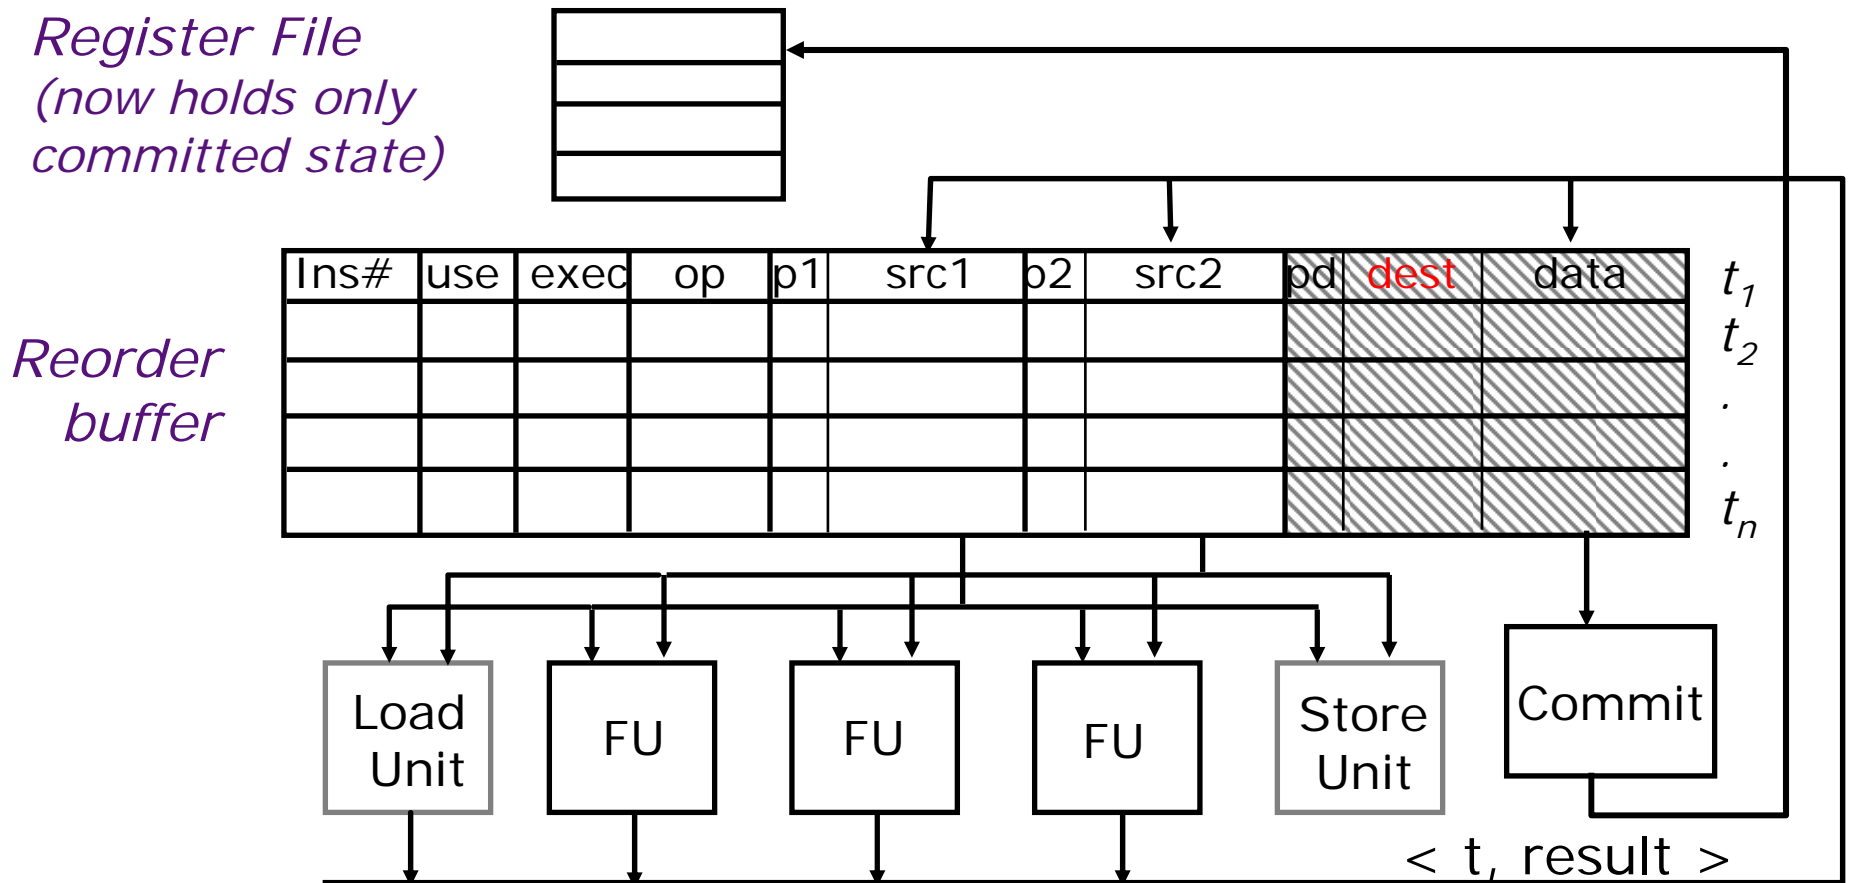
Temporary storage needed to hold results before commit (shadow registers and store buffers)

Extensions for Precise Exceptions



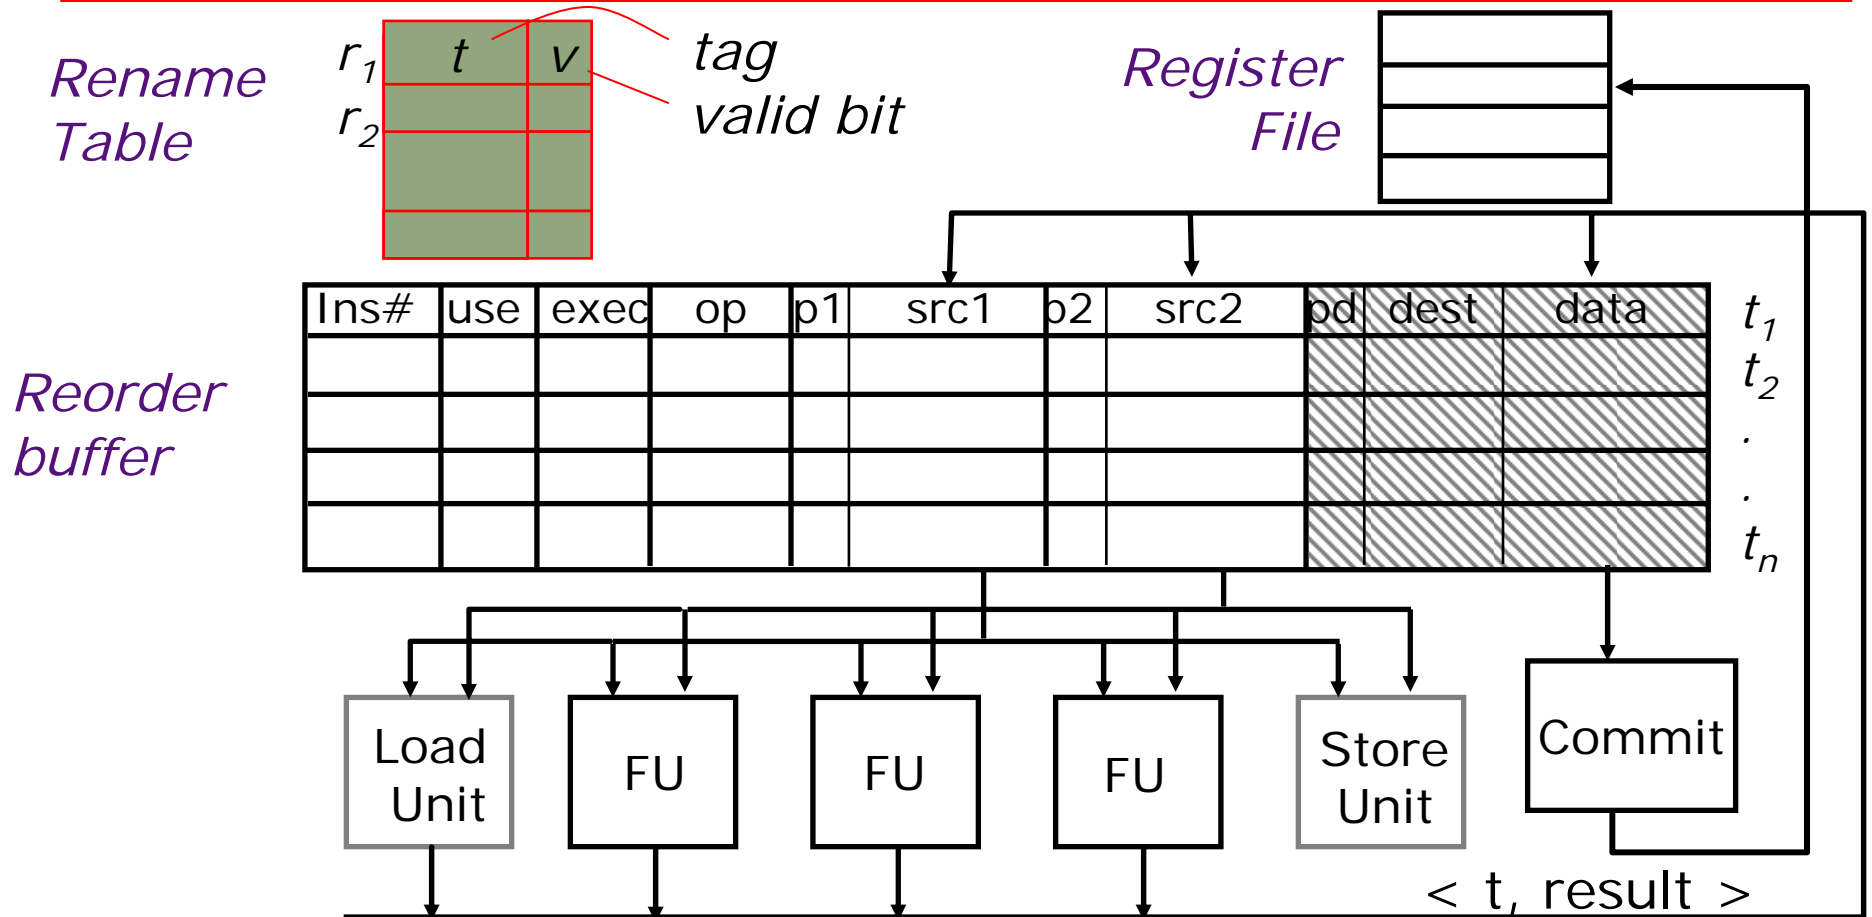
- add $\langle \text{pd}, \text{dest}, \text{data}, \text{cause} \rangle$ fields in the instruction template
- commit instructions to reg file and memory in program order \Rightarrow buffers can be maintained circularly
- on exception, clear reorder buffer by resetting $\text{ptr}_1 = \text{ptr}_2$
(stores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

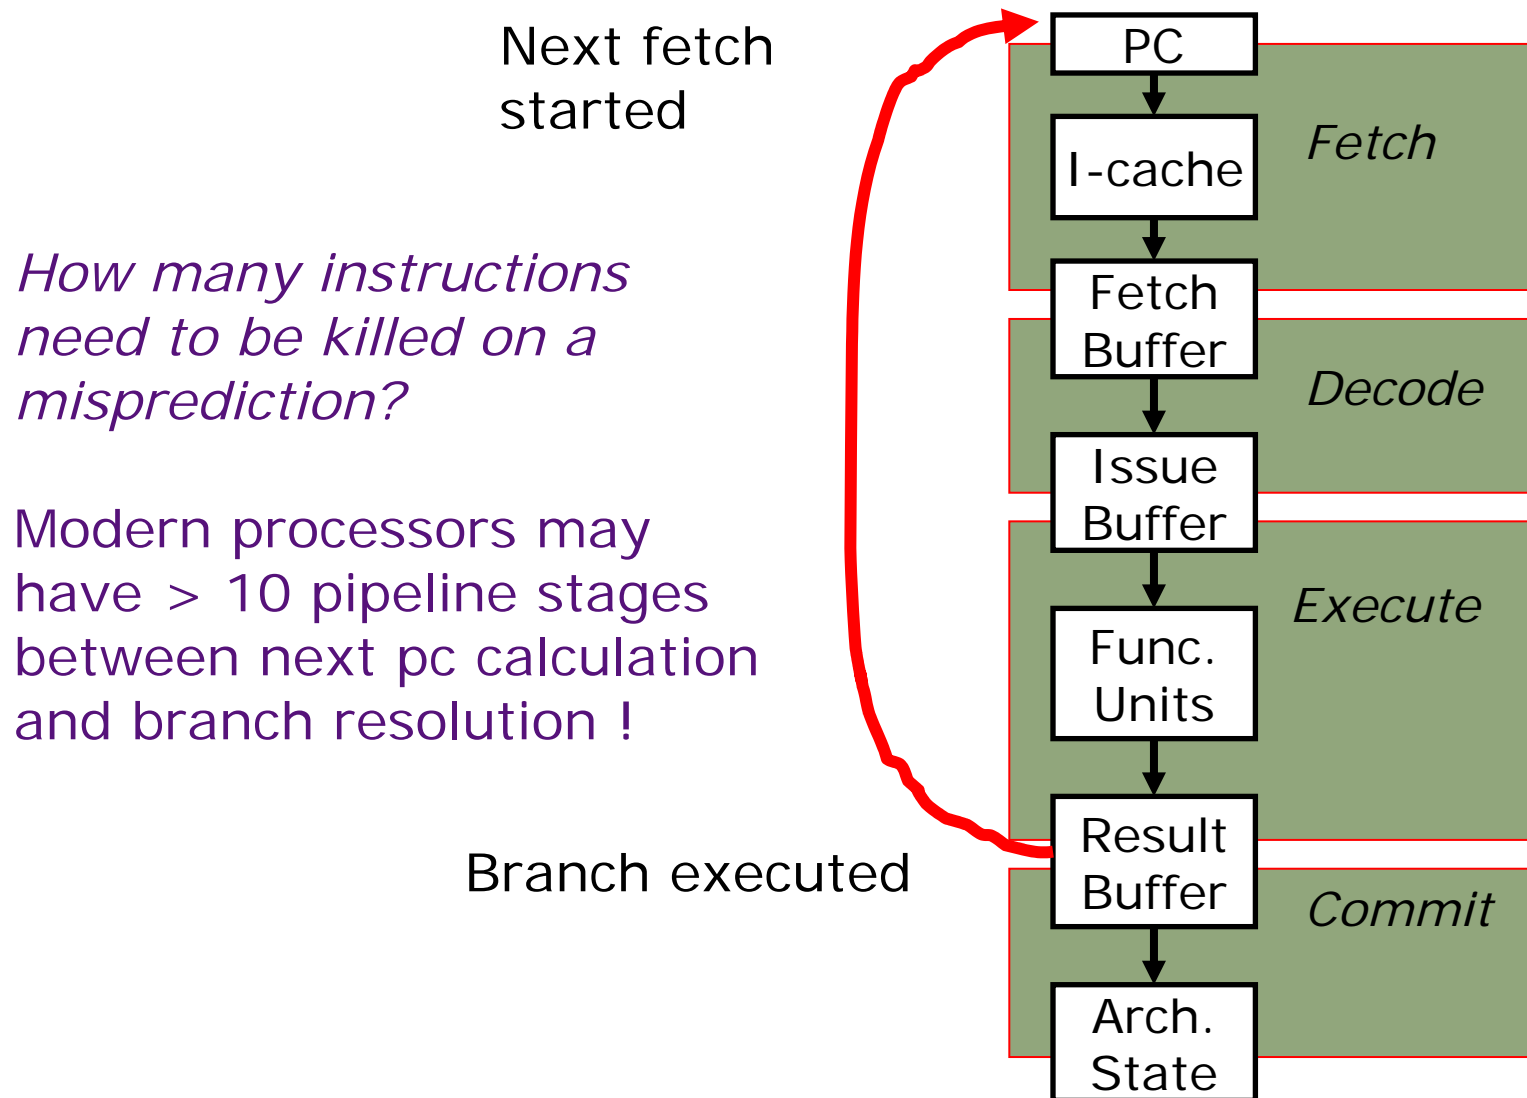


Renaming table is a cache to speed up register name look up.
It needs to be cleared after each exception taken.

When else are valid bits cleared?

Control transfers

Branch Penalty



Average Run-Length between Branches

Average dynamic instruction mix from SPEC92:

	SPECint92	SPECfp92
ALU	39 %	13 %
FPU Add		20 %
FPU Mult		13 %
load	26 %	23 %
store	9 %	9 %
branch	16 %	8 %
other	10 %	12 %

SPECint92: *compress, eqntott, espresso, gcc, li*
SPECfp92: *doduc, ear, hydro2d, mdijdp2, su2cor*

What is the average *run length* between branches

next lecture: Branch prediction & Speculative execution



Thank you !