

CS 152 Computer Architecture and Engineering

Lecture 11 - Out-of-Order Issue, Register Renaming, & Branch Prediction

Dr. George Micheliogiannakis
EECS, University of California at Berkeley
CRD, Lawrence Berkeley National Laboratory

<http://inst.eecs.berkeley.edu/~cs152>

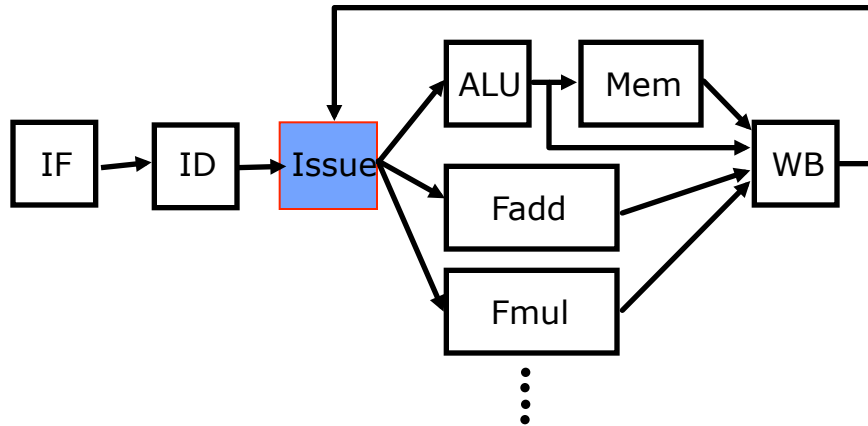
Administrivia

- Lab 2 and PS 2 are due NOW
- Lab 3 release and overview tomorrow
- Pick up PS 1
 - If you can't find your submission talk to me
- Quiz on module 2 next Monday (March 7th)
 - Be on time

Last time in Lecture 10

- Pipelining is complicated by multiple and/or variable latency functional units
- Out-of-order and/or pipelined execution requires tracking of dependencies
 - RAW
 - WAR
 - WAW
- Dynamic issue logic can support out-of-order execution to improve performance
 - Last time, looked at simple scoreboard to track out-of-order completion
- Hardware register renaming can further improve performance by removing WAW and WAR hazards.

Register Renaming



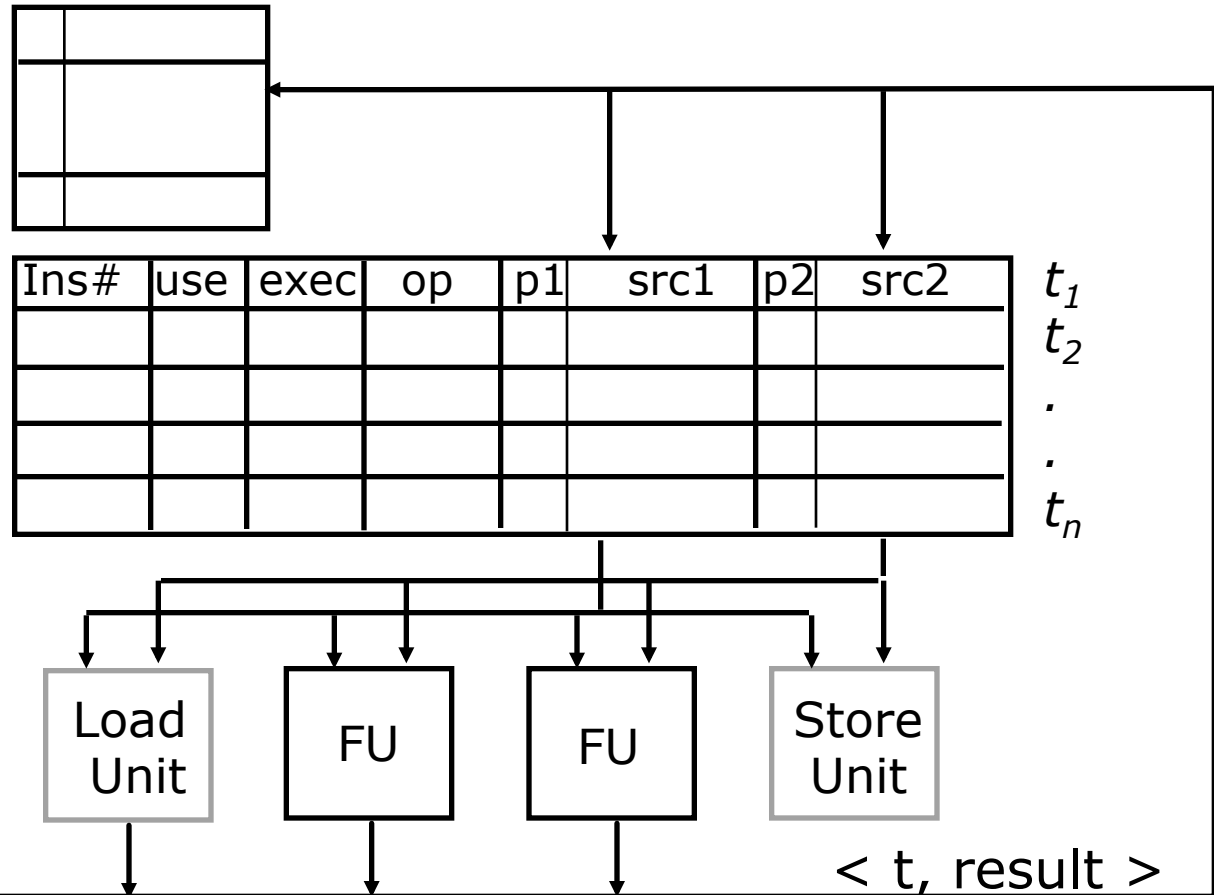
- Decode does register renaming and adds instructions to the issue-stage instruction reorder buffer (ROB)
 - ⇒ renaming makes WAR or WAW hazards impossible
- Any instruction in ROB whose RAW hazards have been satisfied can be issued.
 - ⇒ Out-of-order or dataflow execution

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

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

t_1
 t_2
 t_3
 t_4
 t_5
.
.

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

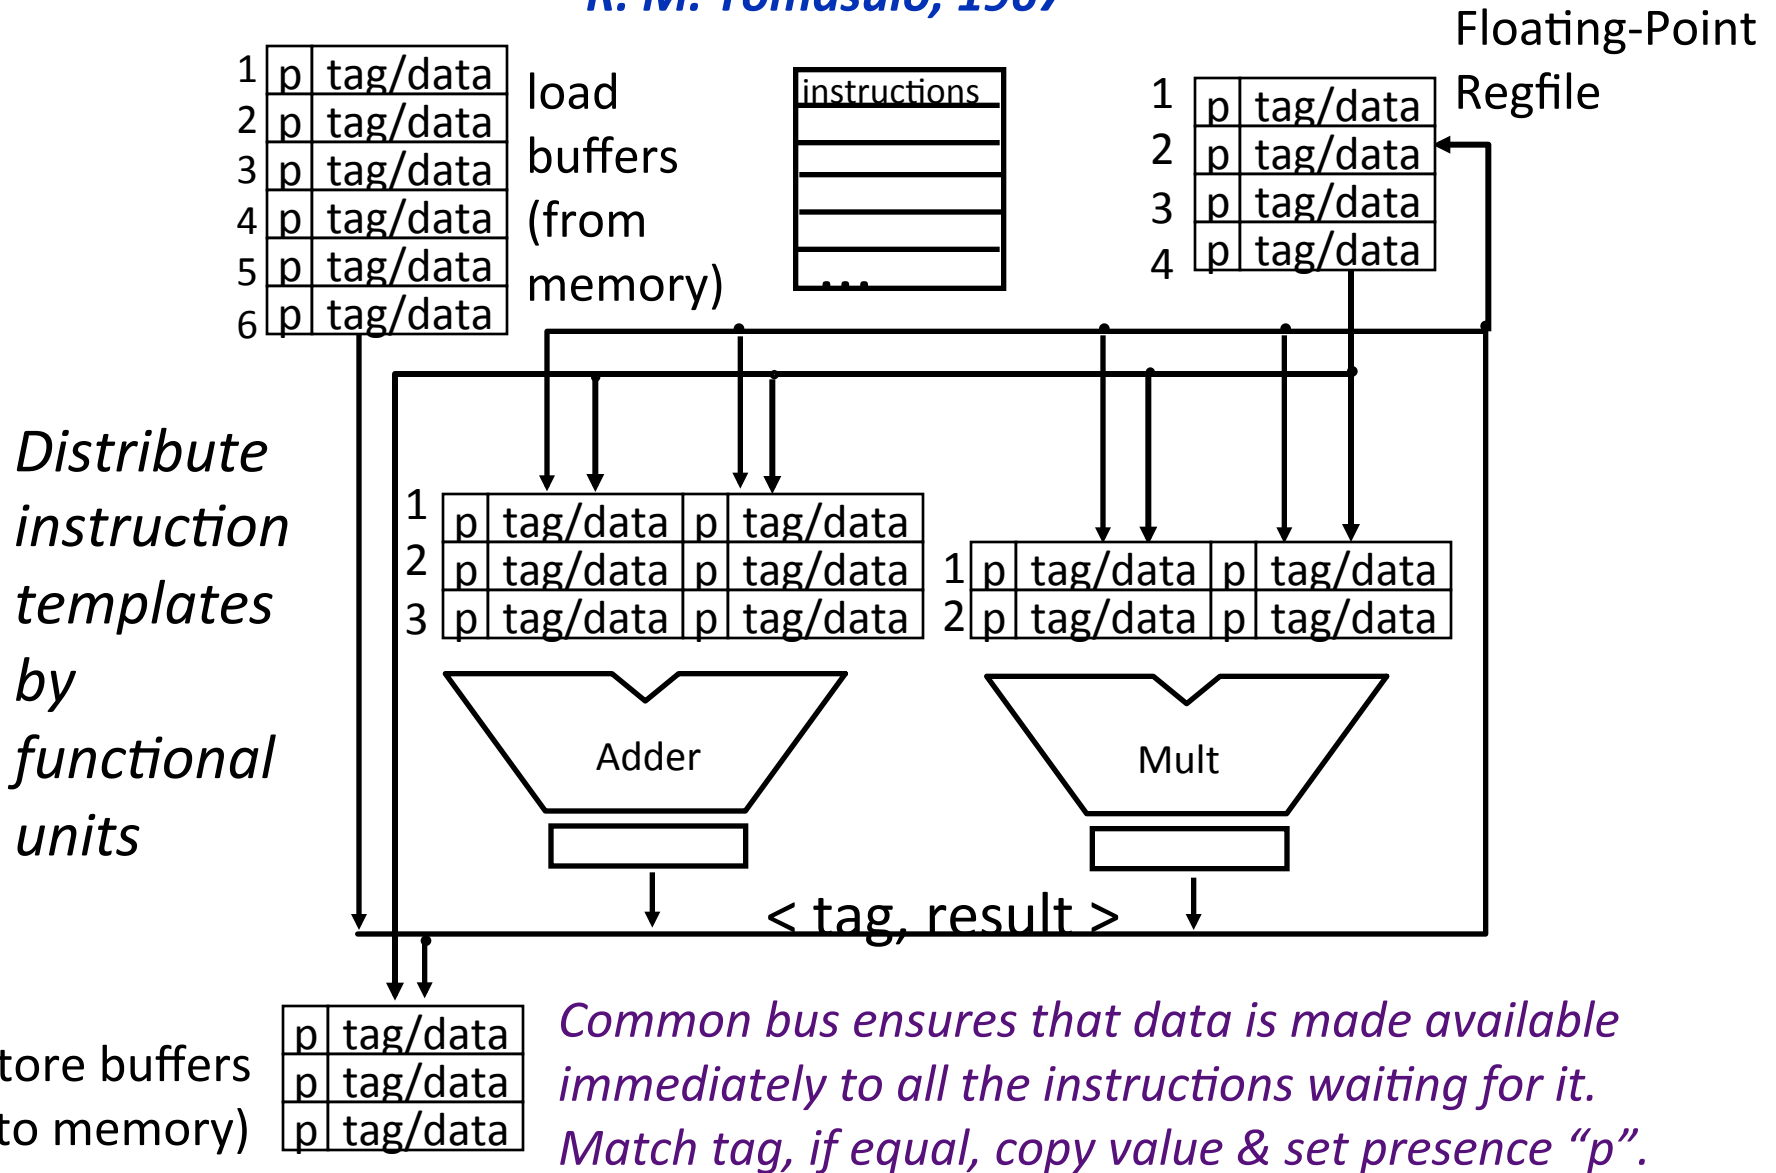
- When are tags in sources replaced by data?
Whenever an FU produces data
- When can a name be reused?
Whenever an instruction completes

In Summary

- Register indexes the compiler emits are used to detect data dependencies
- Tags are then used much like variable names, to denote that the value is the same (an instruction creates a tag)
- When an instruction writes a register, that updates the tag if there was one before
 - WAW was a reason for register renaming (see previous lecture)

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. Effective on a very small class of programs
2. Memory latency a much bigger problem
3. Exceptions not precise!

One more problem needed to be solved

Control transfers

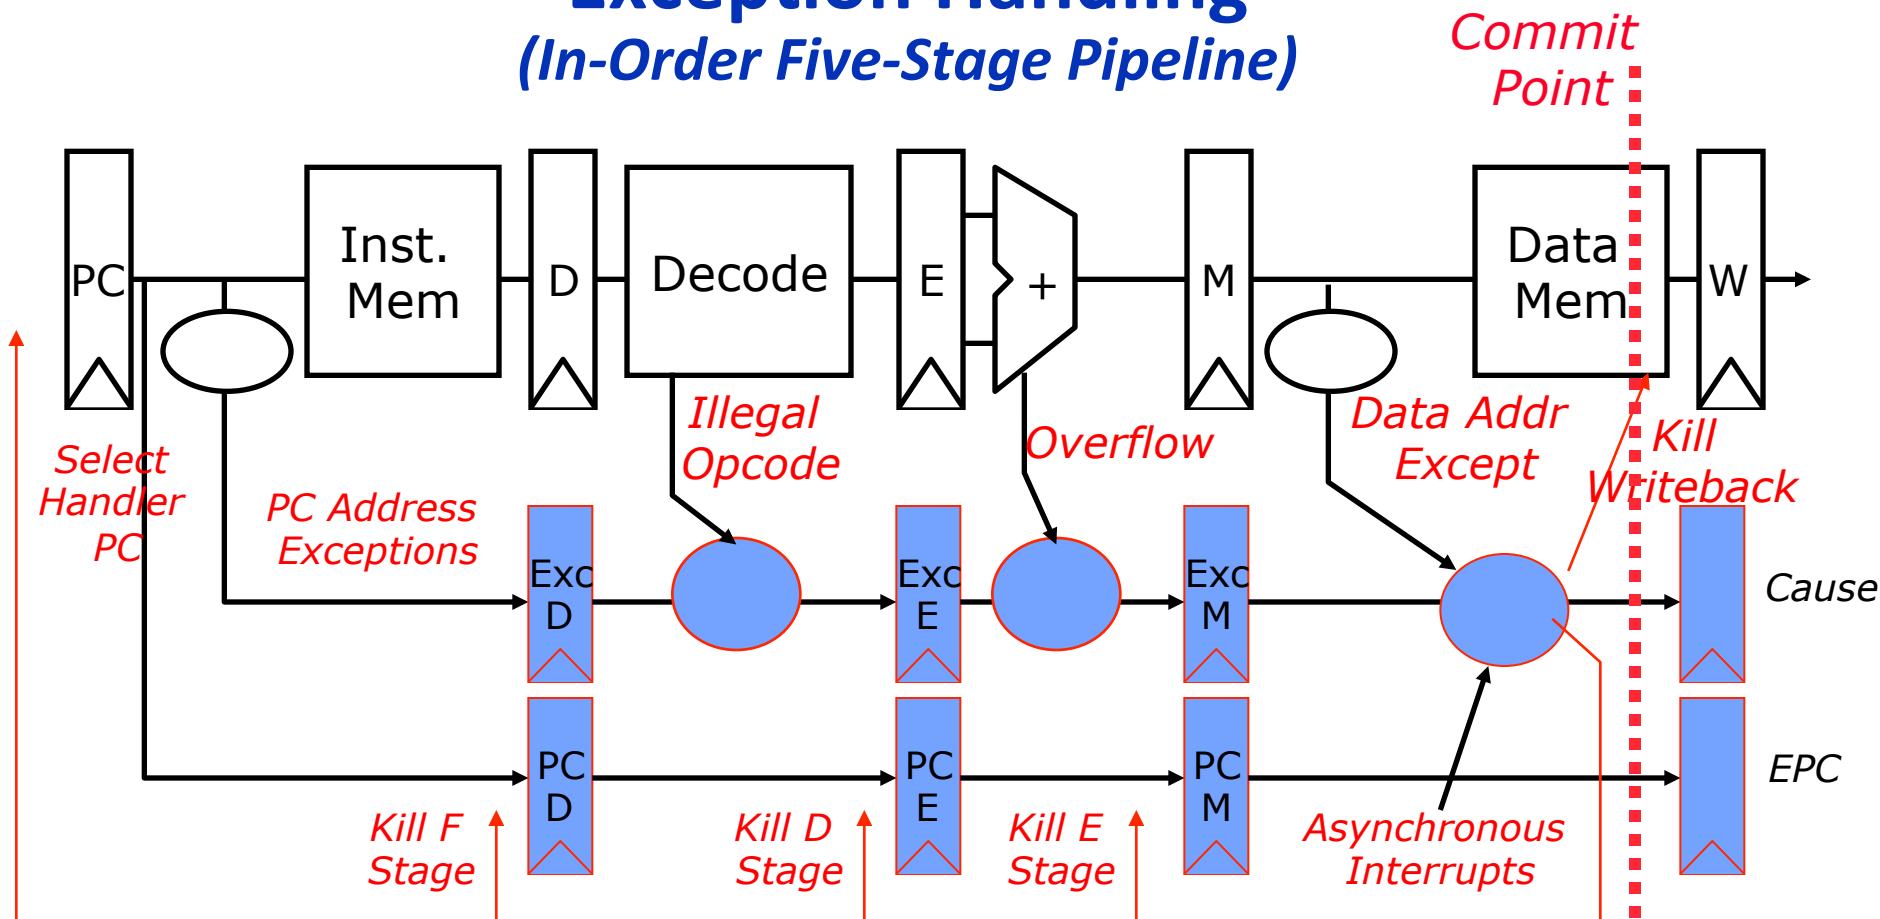
Precise Interrupts

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

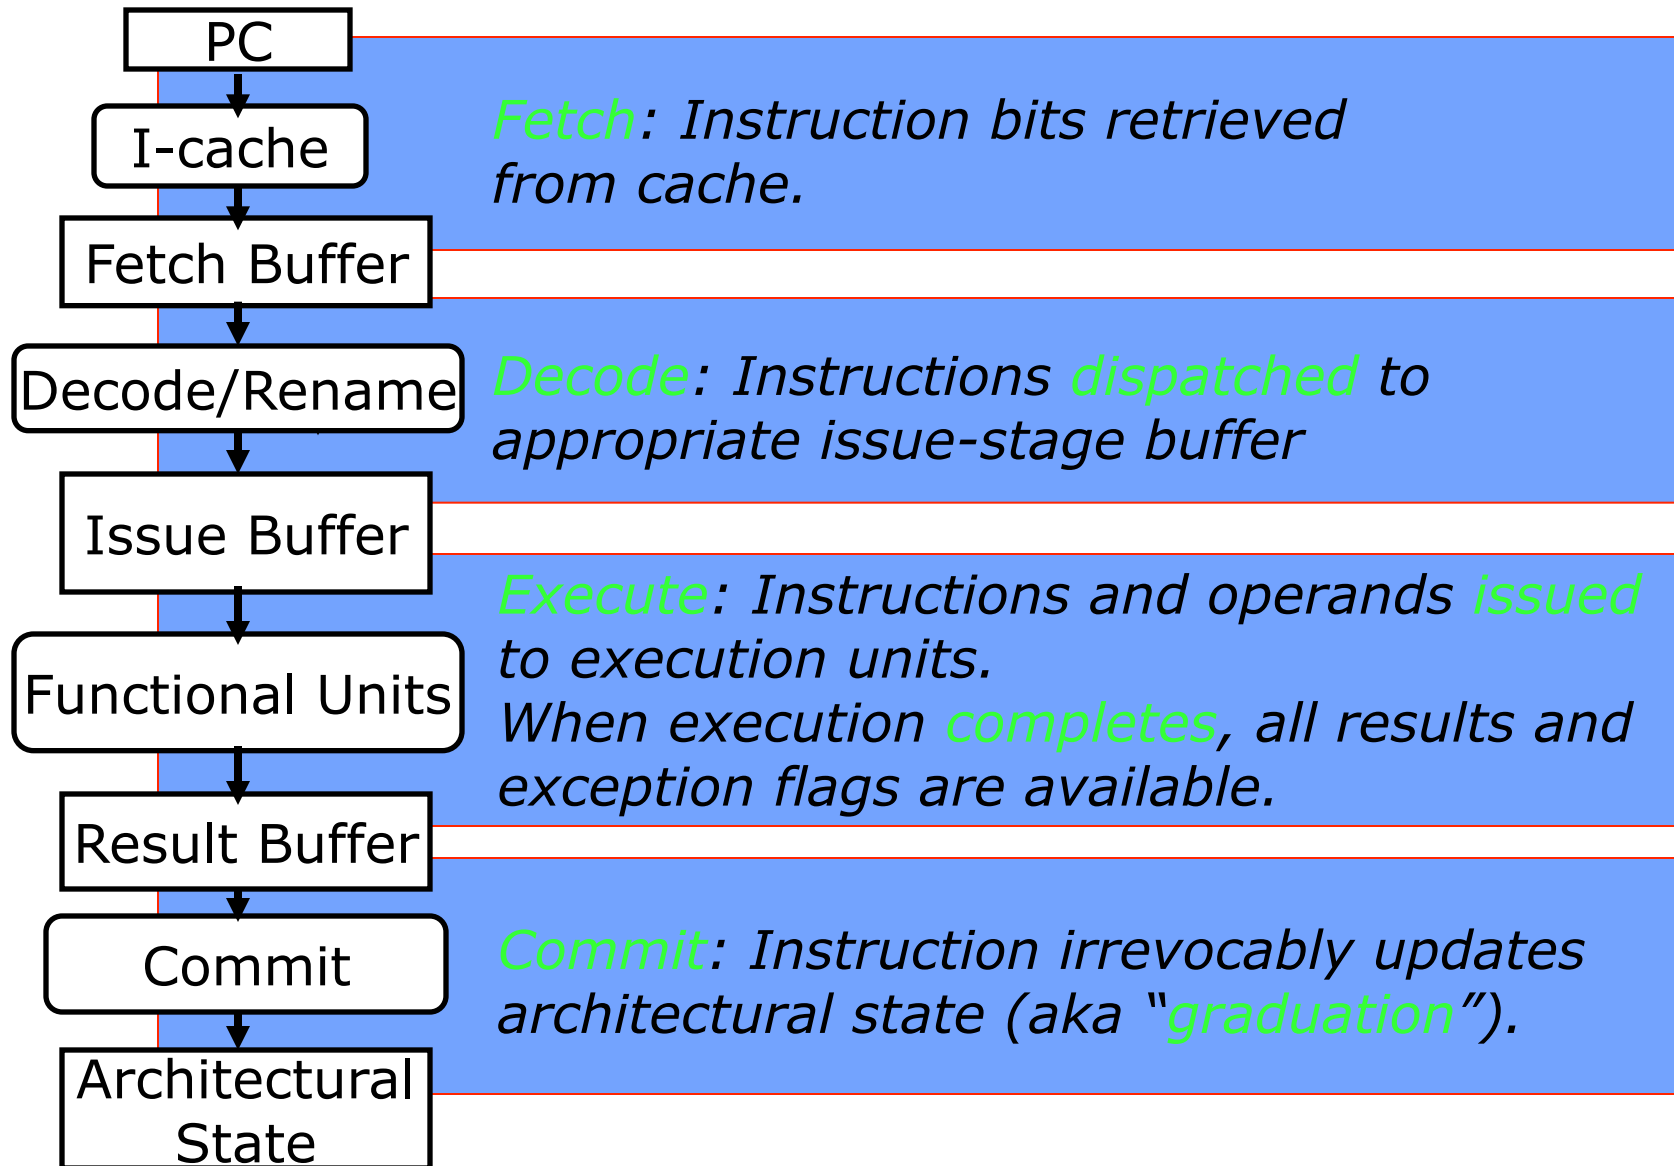
The interrupt handler either aborts the program or restarts it at I_{i+1} .

Exception Handling (In-Order Five-Stage Pipeline)

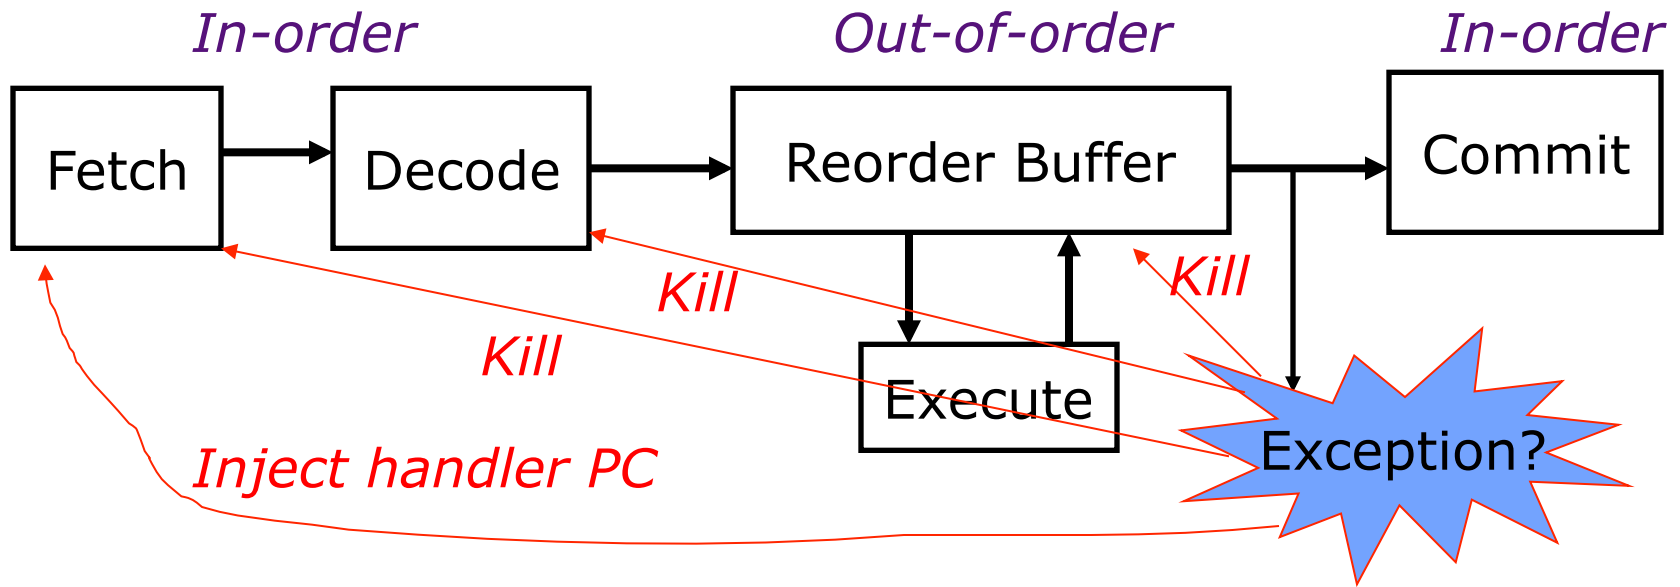


- Hold exception flags in pipeline until commit point (M stage)
- Exceptions in earlier pipe stages override later exceptions
- Inject external interrupts at commit point (override others)
- If exception at commit: update Cause and EPC registers, kill all stages, inject handler PC into fetch stage

Phases of Instruction Execution



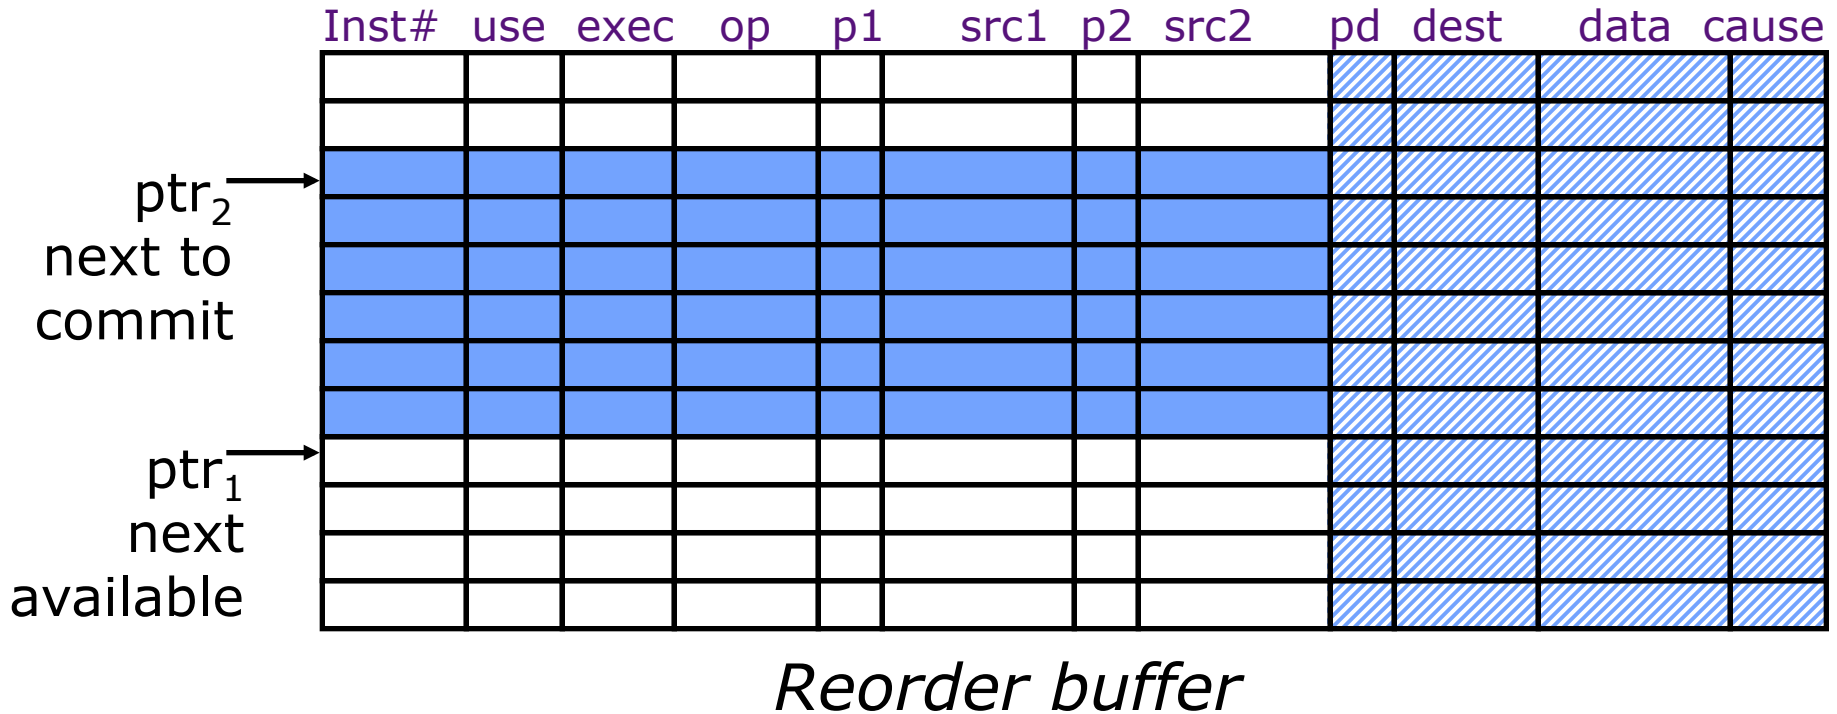
In-Order Commit for Precise Exceptions



- 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
(now holds only committed state)



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


Rename Table



Register File



Reorder buffer



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?

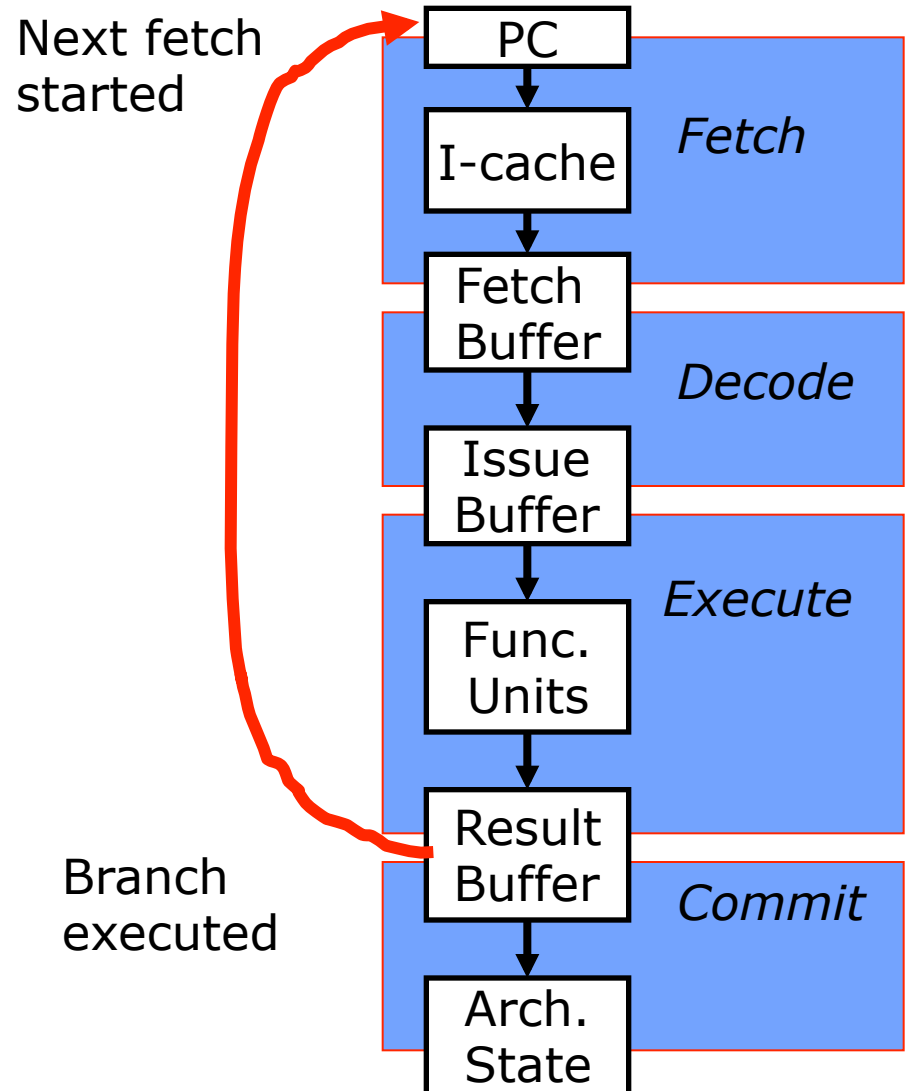
Control transfers

Control Flow Penalty

*Modern processors may have
> 10 pipeline stages between
next PC calculation and branch
resolution !*

*How much work is lost if
pipeline doesn't follow
correct instruction flow?*

~ Loop length x pipeline width



Mispredict Recovery

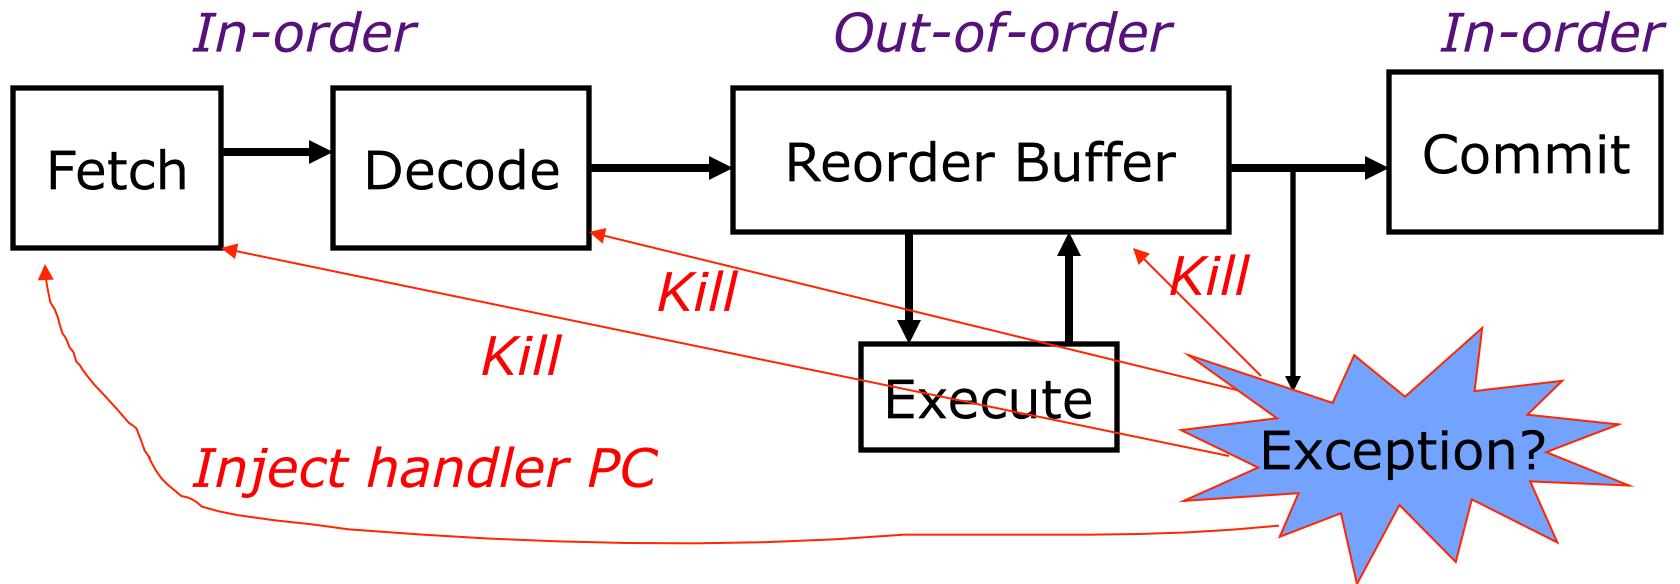
In-order execution machines:

- Assume no instruction issued after branch can write-back before branch resolves
- Kill all instructions in pipeline behind mispredicted branch

Out-of-order execution?

- Multiple instructions following branch in program order can complete before branch resolves

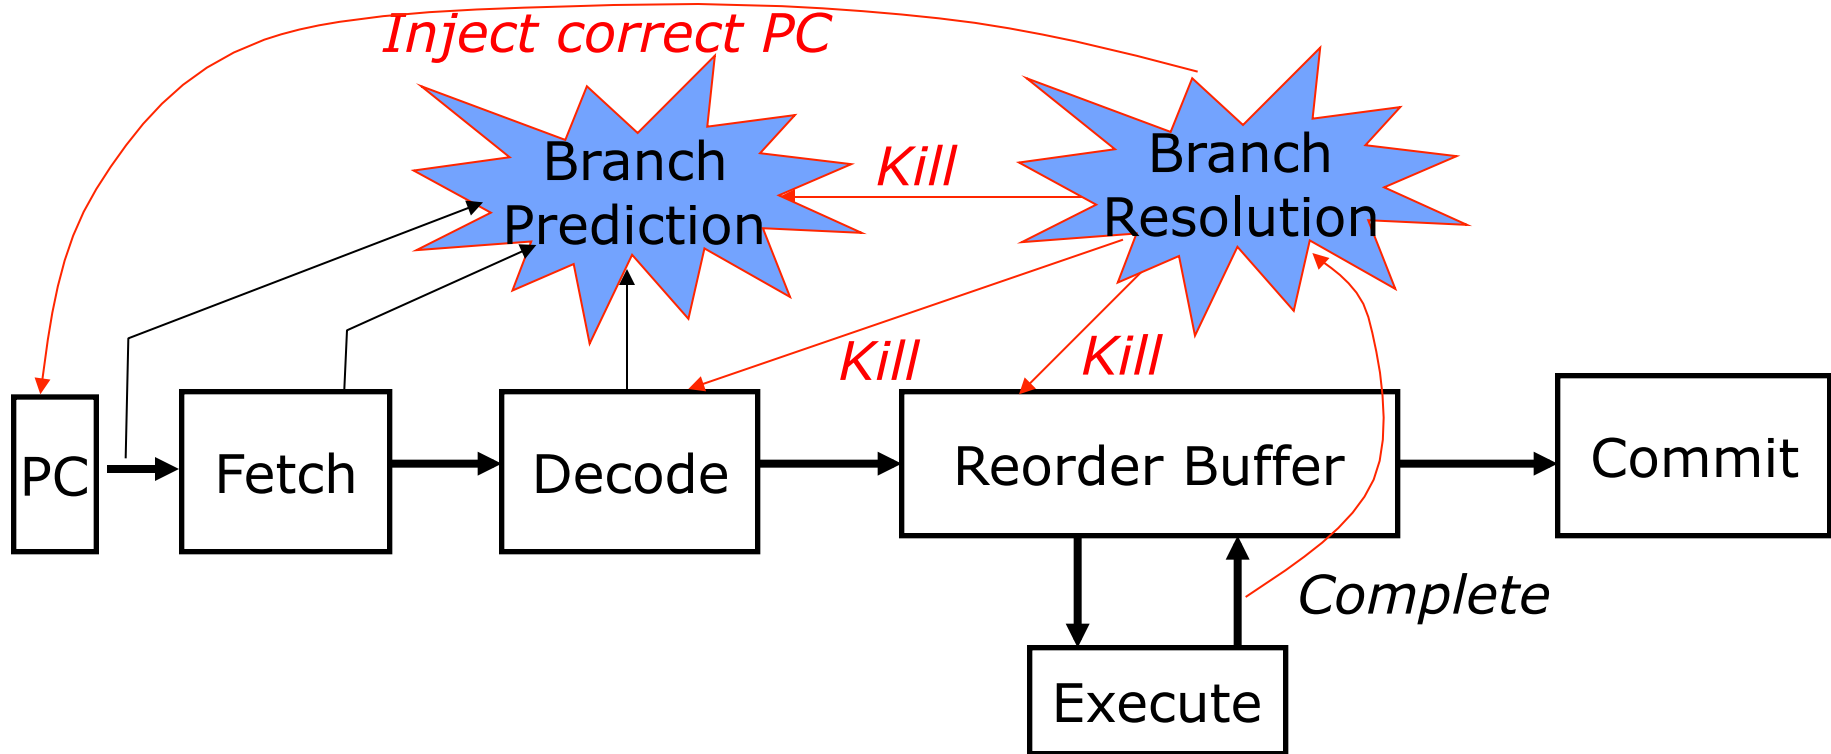
In-Order Commit for Precise Exceptions



- 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 in ROB to hold results before commit

Branch Misprediction in Pipeline




- Can have multiple unresolved branches in ROB
- Can resolve branches out-of-order by killing all the instructions in ROB that follow a mispredicted branch

“Data-in-ROB” Design


(HP PA8000, Pentium Pro, Core2Duo, Nehalem)

Register File
holds only
committed state



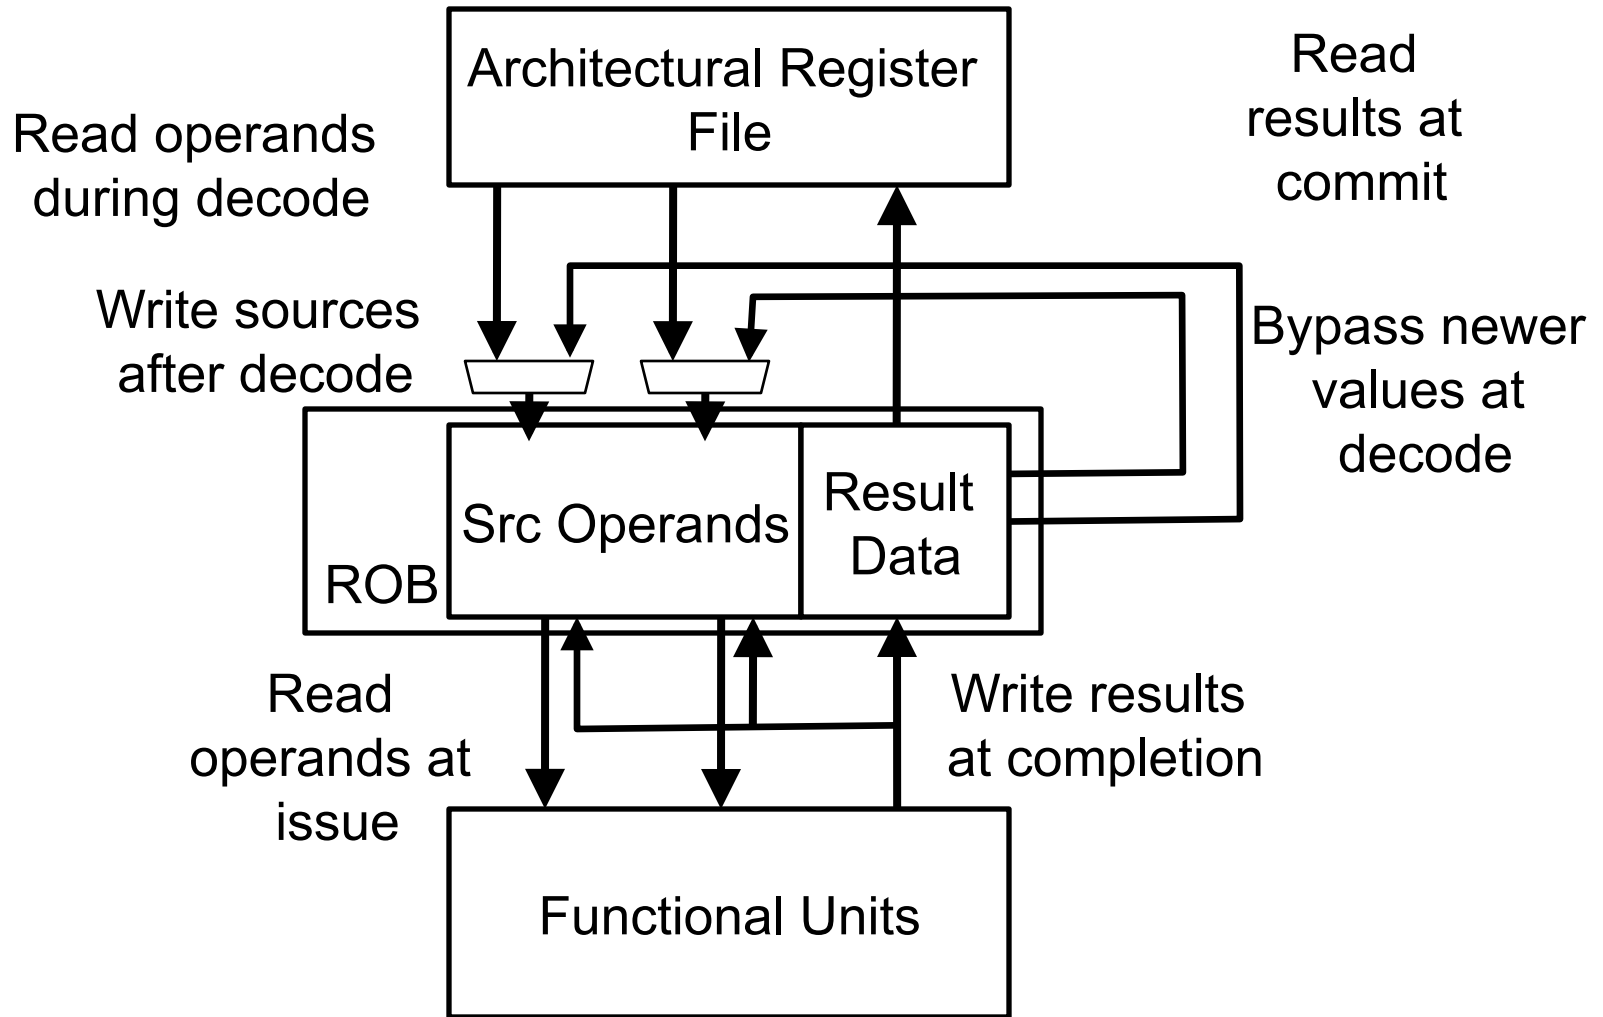
Reorder
buffer

Ins#	use	exec	op	p1	src1	p2	src2	pd	dest	data	
											t_1
											t_2
											⋮
											t_n



- On dispatch into ROB, ready sources can be in regfile or in ROB dest (copied into src1/src2 if ready before dispatch)
- On completion, write to dest field and broadcast to src fields.
- On issue, read from ROB src fields

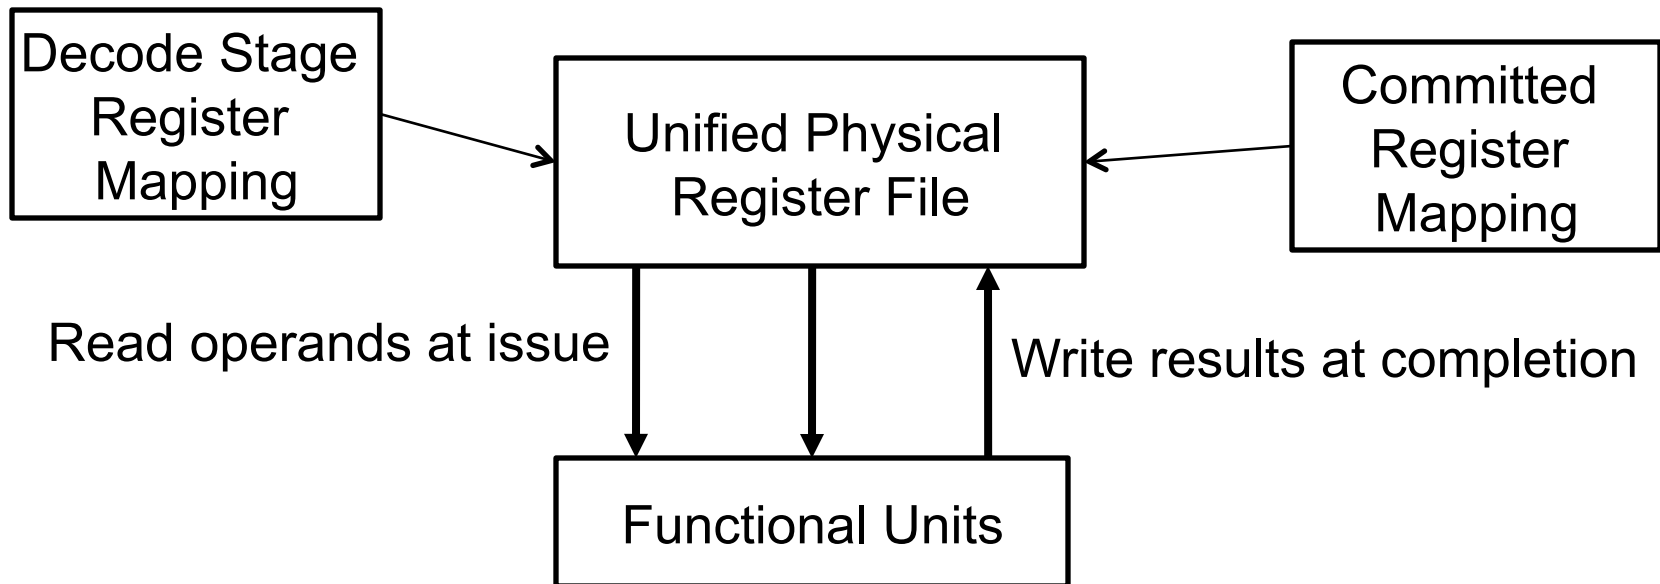
Data Movement in Data-in-ROB Design



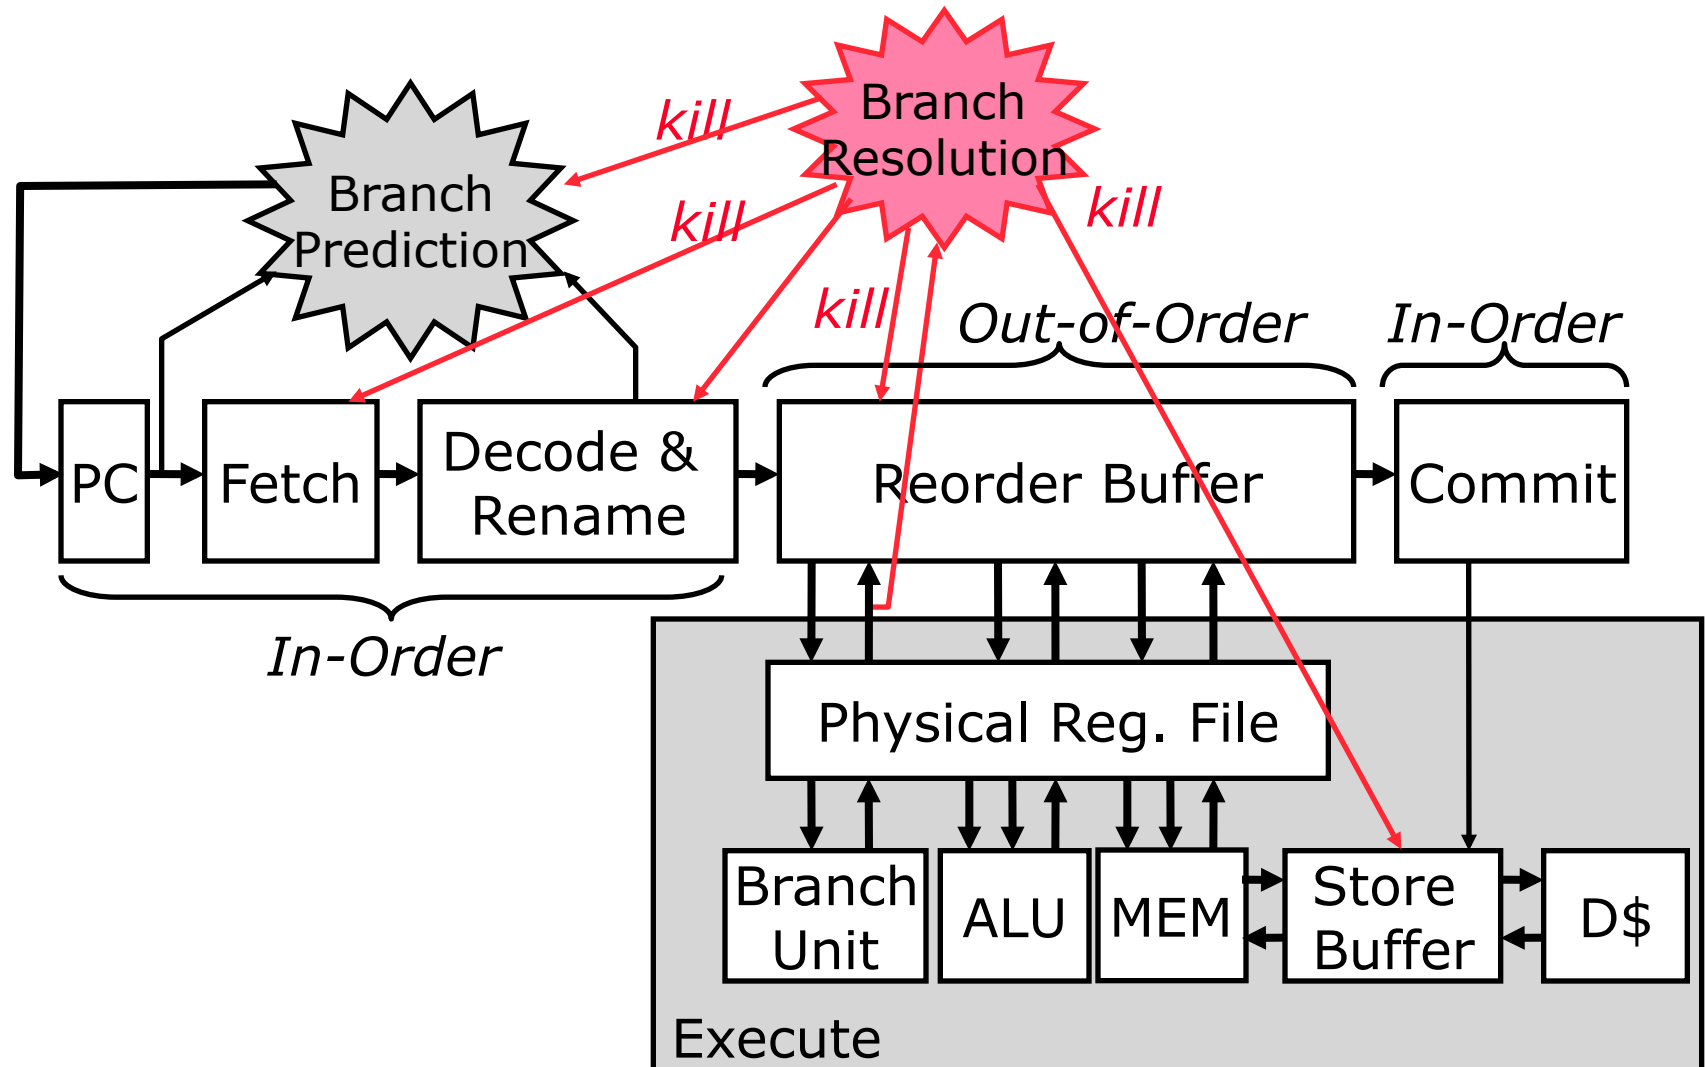
Unified Physical Register File

(MIPS R10K, Alpha 21264, Intel Pentium 4 & Sandy Bridge)

- Rename all architectural registers into a single *physical* register file during decode, no register values read
 - x1 -> P1
- Functional units read and write from single unified register file holding committed and temporary registers in execute
- Commit only updates mapping of architectural register to physical register, no data movement




Pipeline Design with Physical Regfile



Lifetime of Physical Registers

- Physical regfile holds committed and speculative values
- Physical registers decoupled from ROB entries (*no data in ROB*)

```
ld x1, (x3)
addi x3, x1, #4
sub x6, x7, x9
add x3, x3, x6
ld x6, (x1)
add x6, x6, x3
sd x6, (x1)
ld x6, (x11)
```

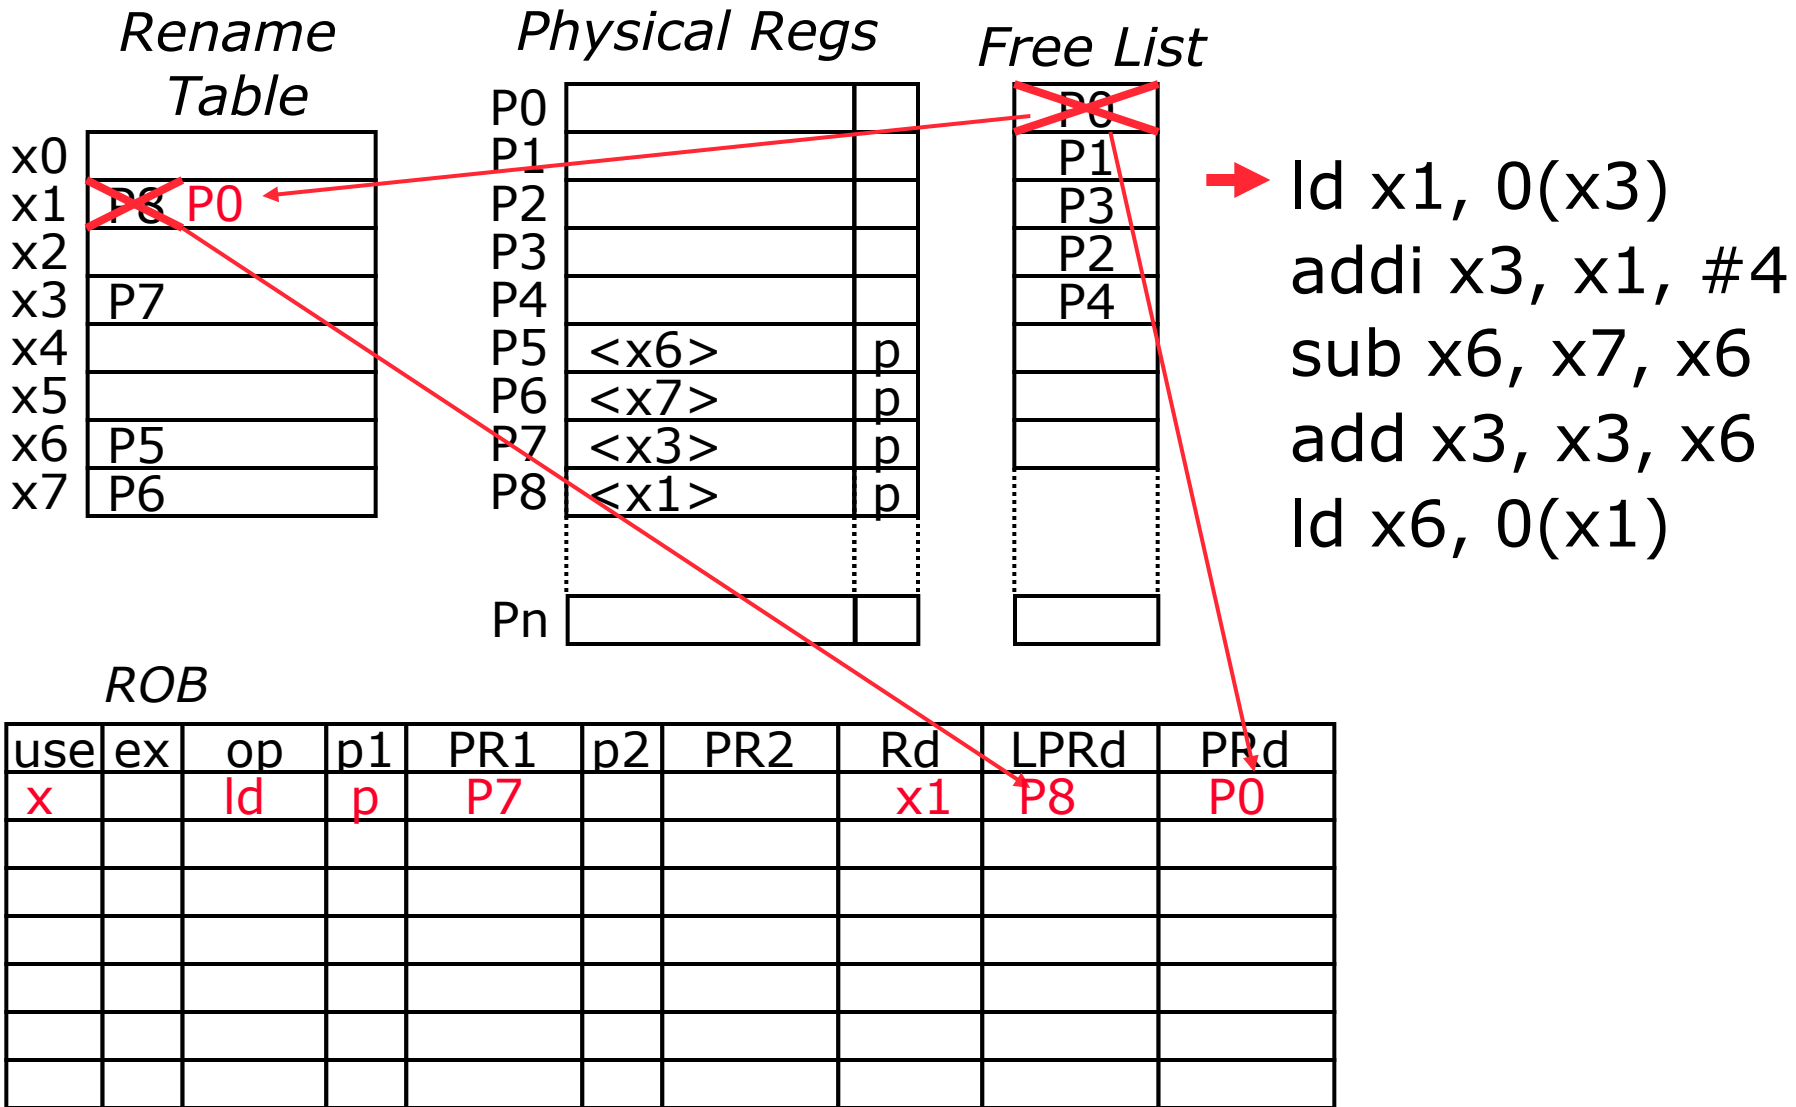


```
ld P1, (Px)
addi P2, P1, #4
sub P3, Py, Pz
add P4, P2, P3
ld P5, (P1)
add P6, P5, P4
sd P6, (P1)
ld P7, (Pw)
```

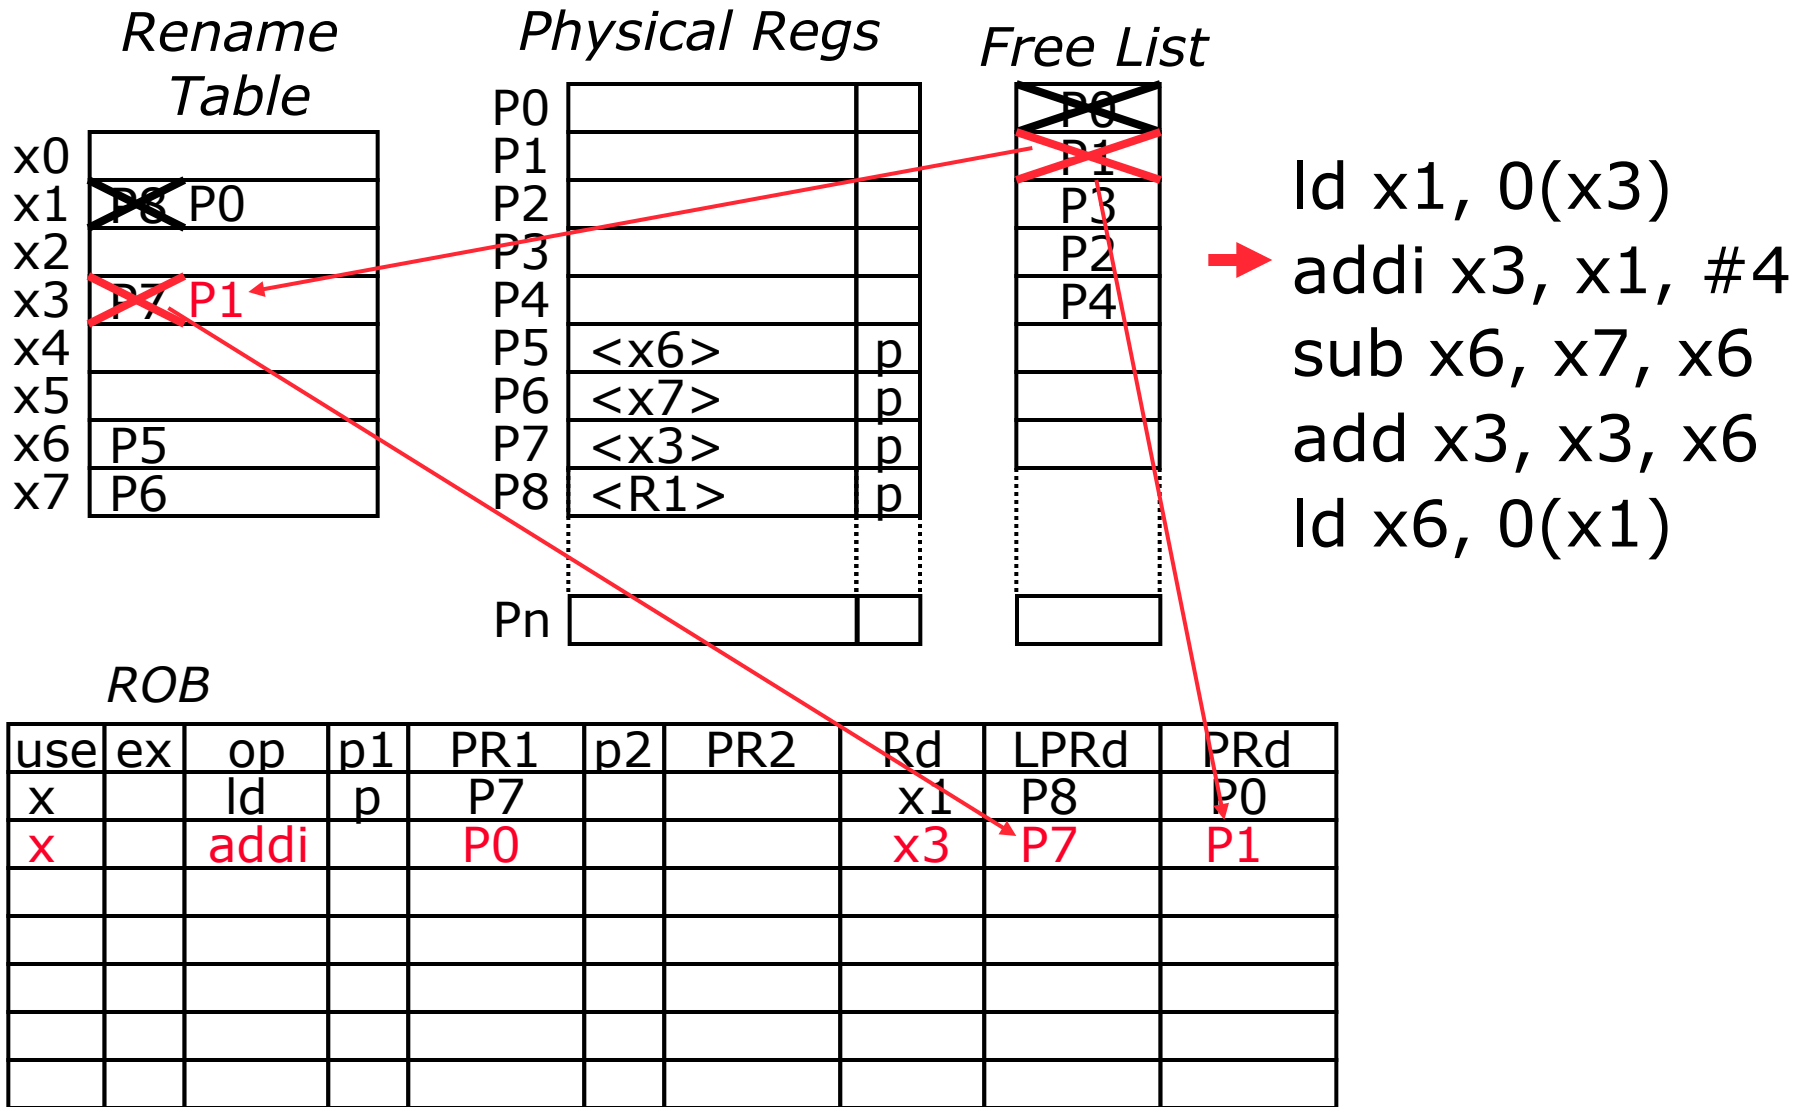
When can we reuse a physical register?

When next write of same architectural register commits

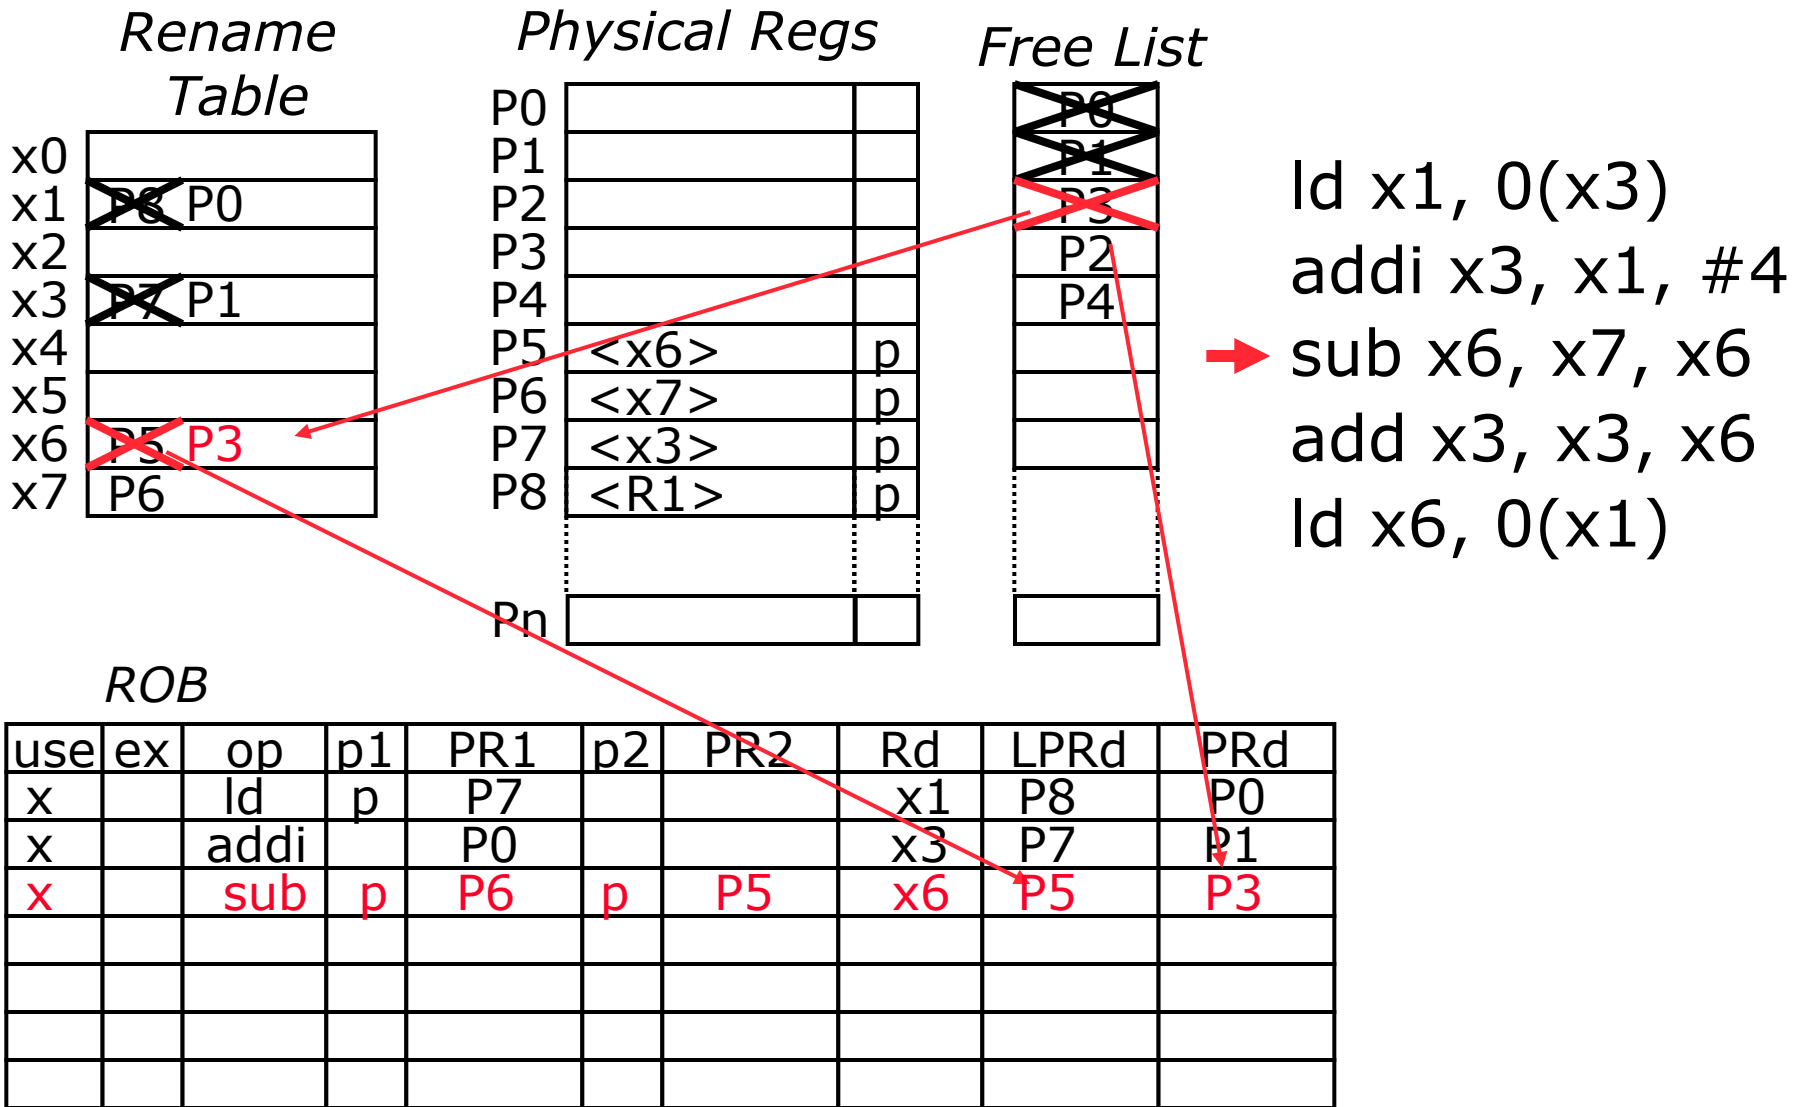
Physical Register Management



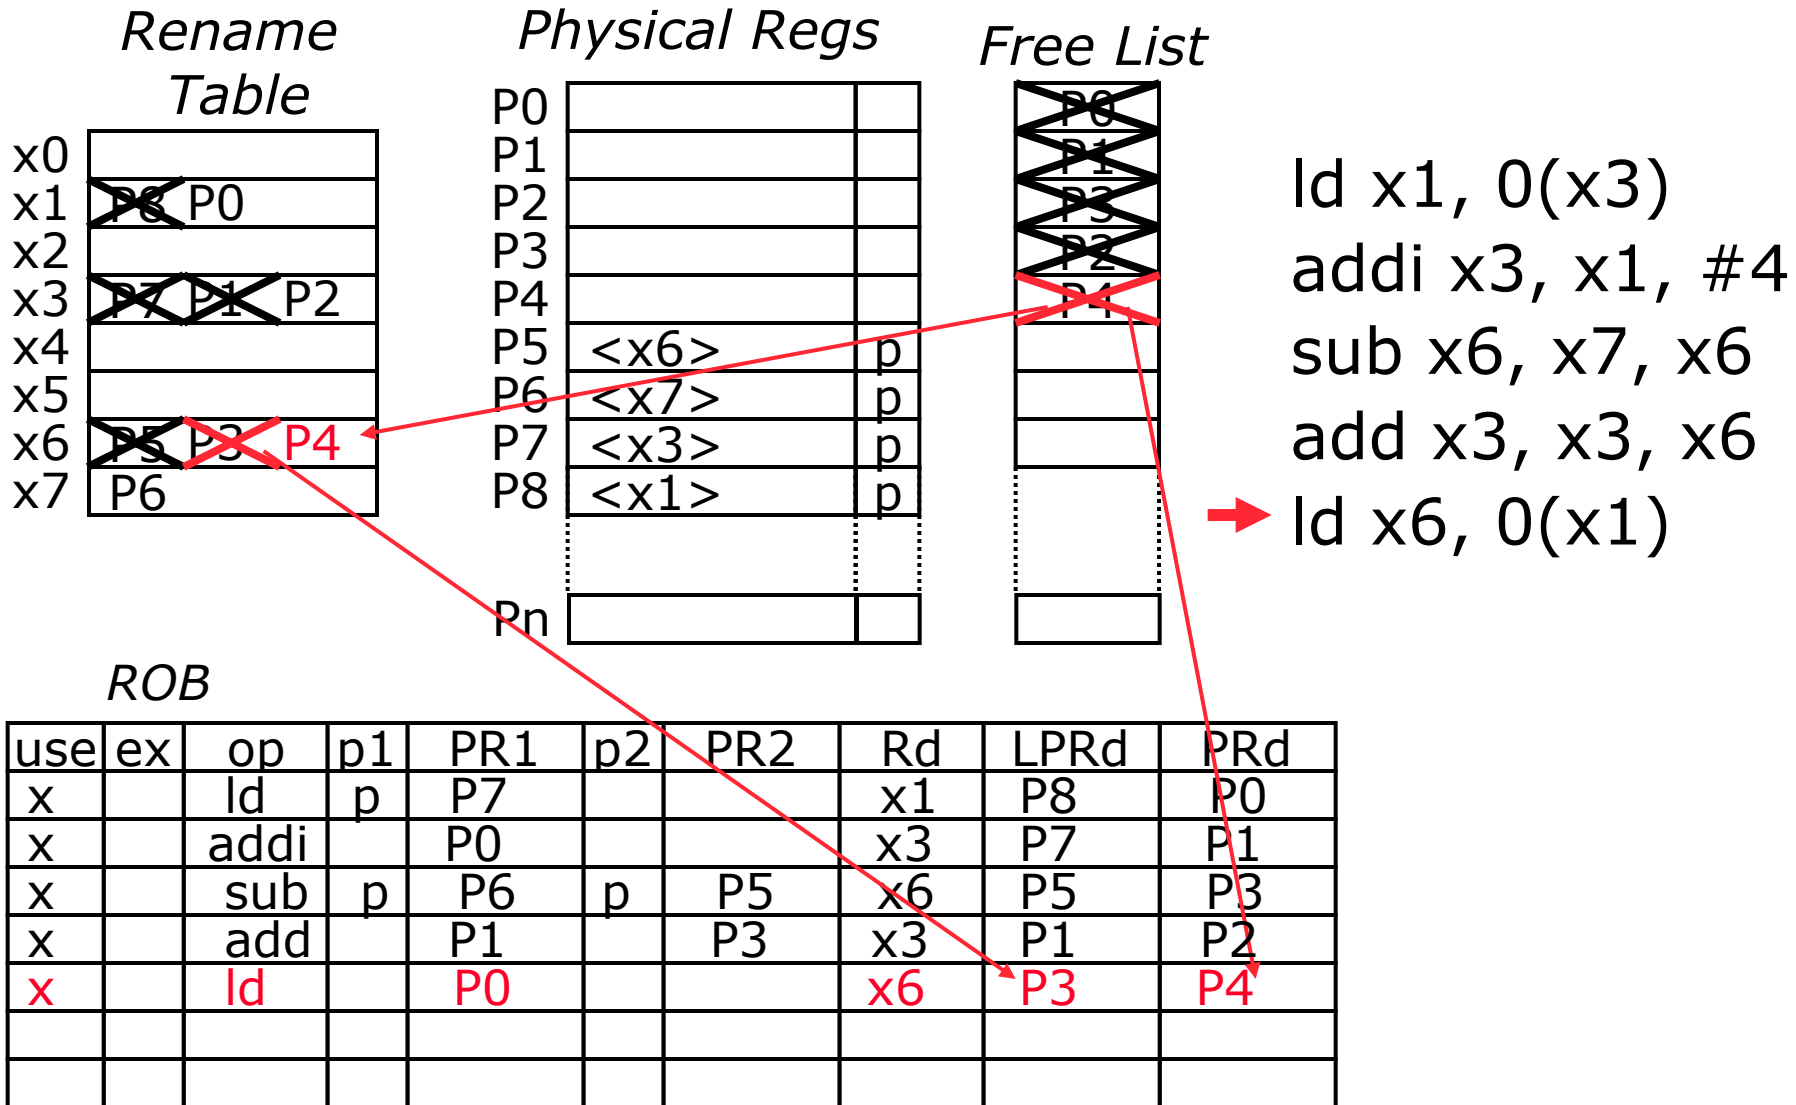
Physical Register Management



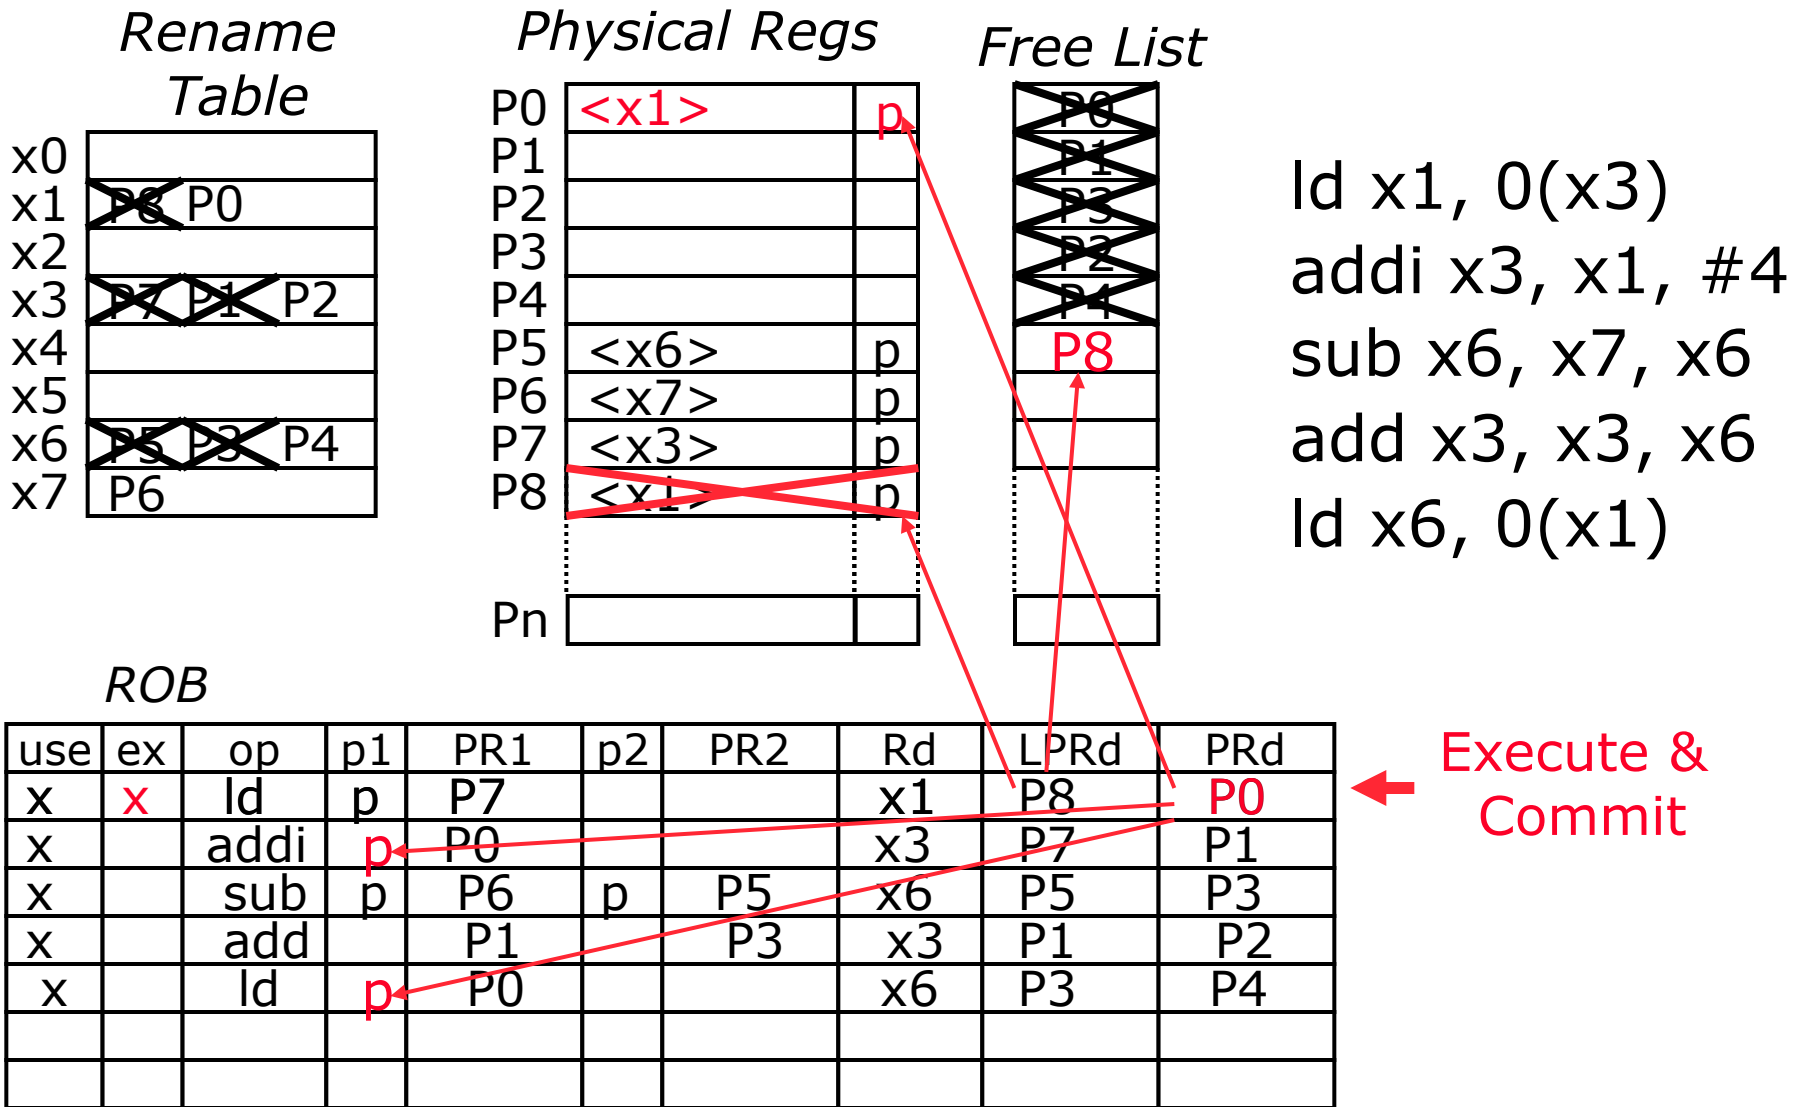
Physical Register Management



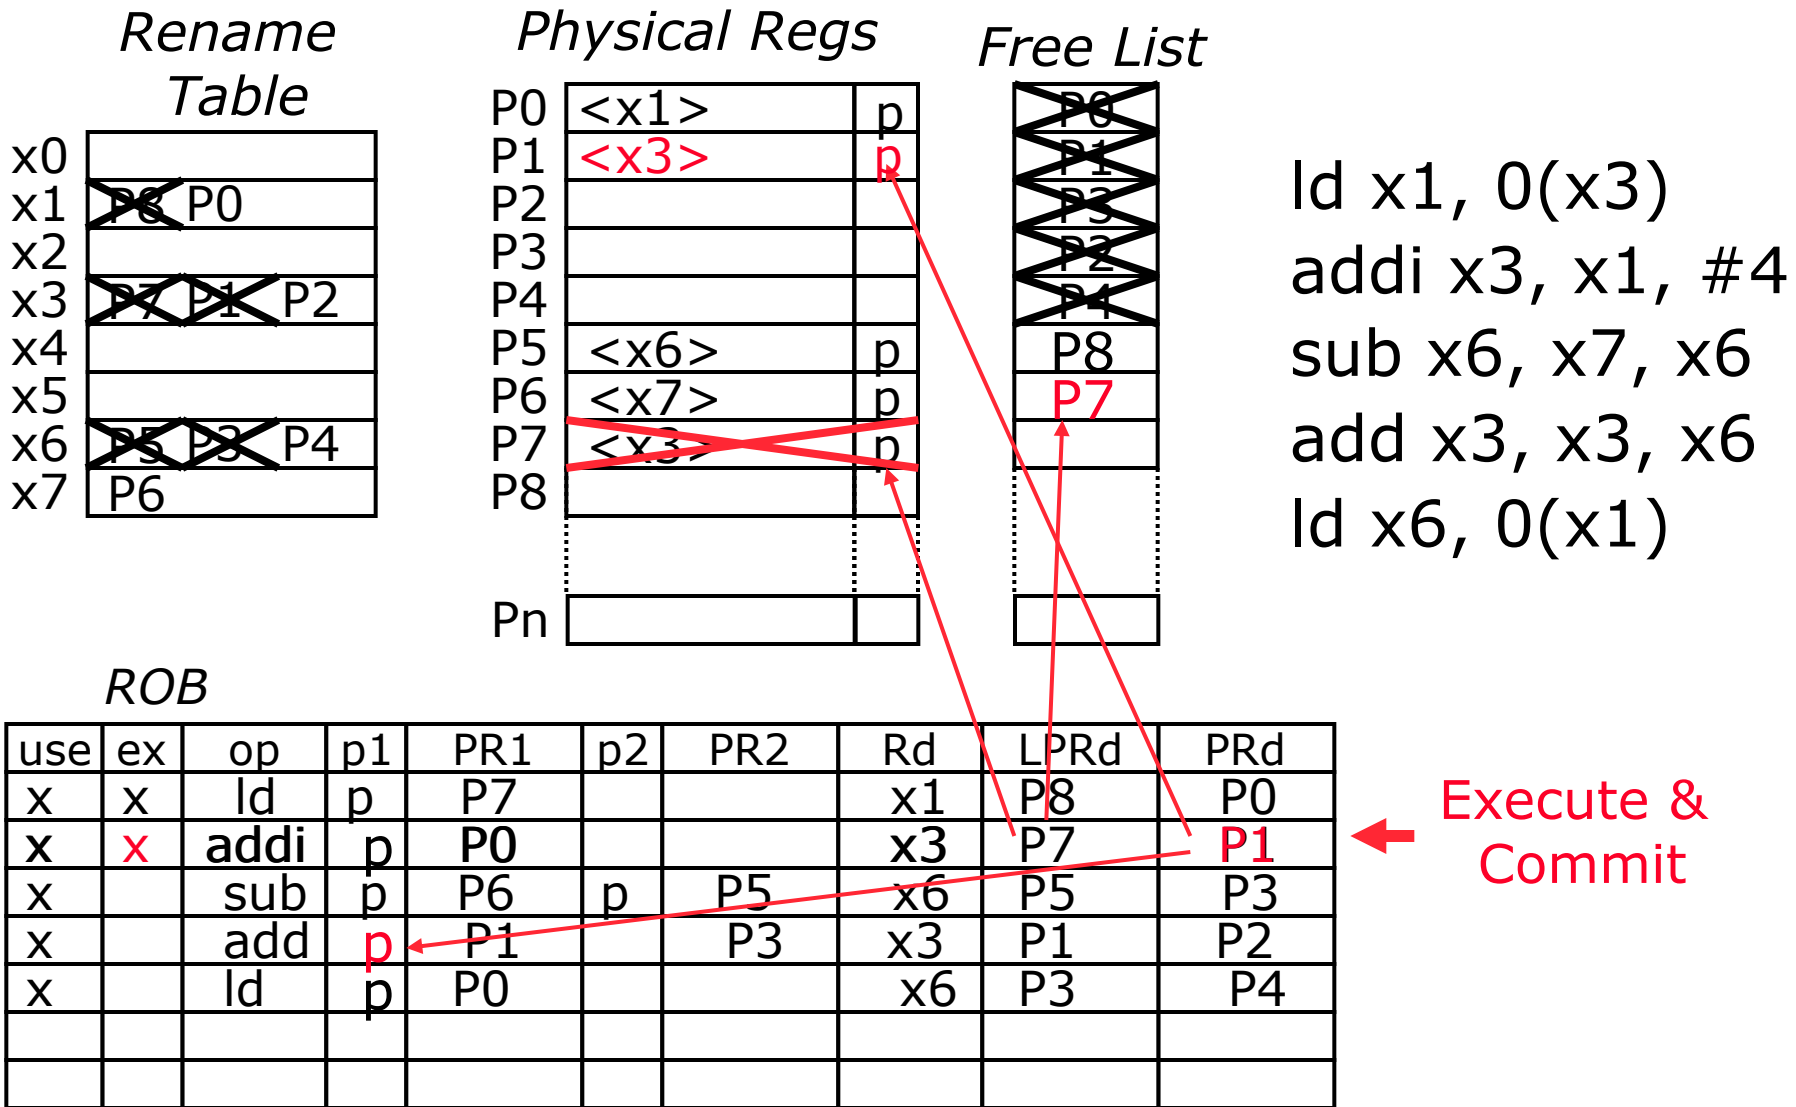
Physical Register Management



Physical Register Management



Physical Register Management



Acknowledgements

- These slides contain material developed and copyright by:
 - Arvind (MIT)
 - Krste Asanovic (MIT/UCB)
 - Joel Emer (Intel/MIT)
 - James Hoe (CMU)
 - John Kubiatowicz (UCB)
 - David Patterson (UCB)
- MIT material derived from course 6.823
- UCB material derived from course CS252