



CS 152 Computer Architecture and Engineering

Lecture 13 - Out-of-Order Issue, Register Renaming, & Branch Prediction

Krste Asanovic

Electrical Engineering and Computer Sciences
University of California at Berkeley

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

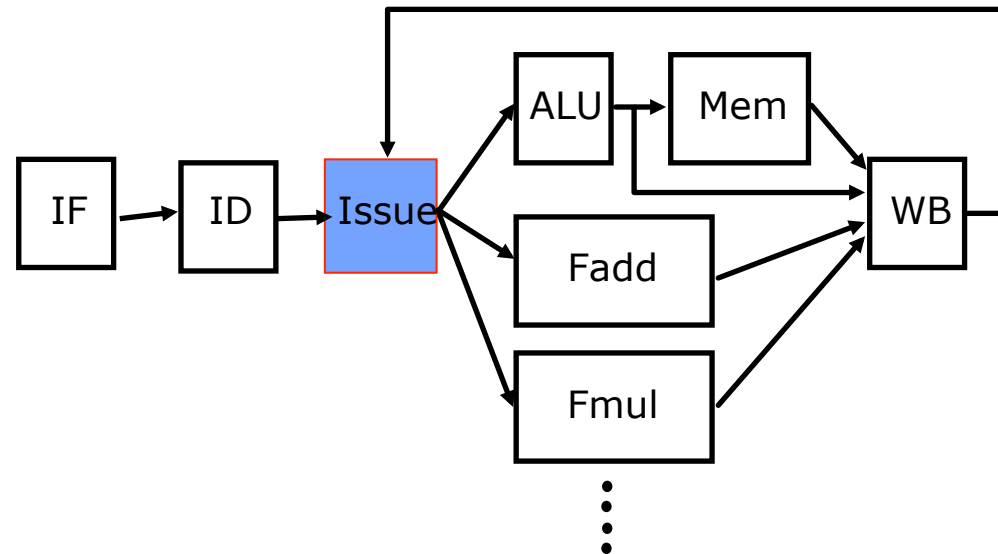


Last time in Lecture 12

- 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 hazards.



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.



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*



Instruction-level Parallelism via 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

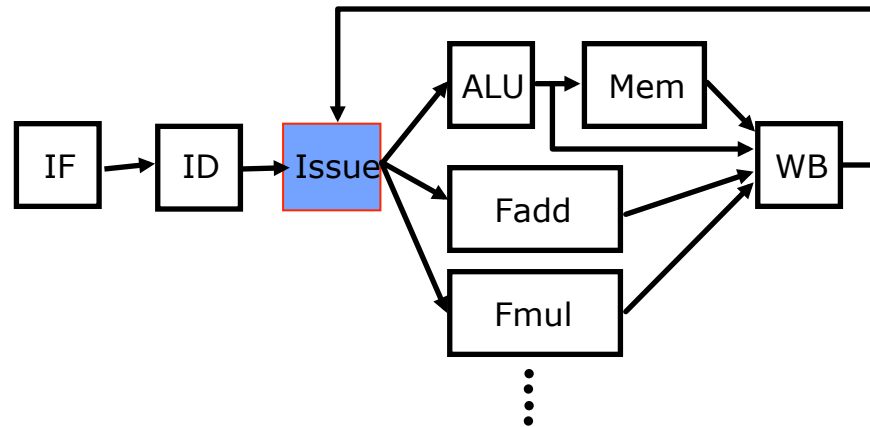


In-order: 1 (2,1) 2 3 4 4 3 5 5 6 6
 Out-of-order: 1 (2,1) 4 4 5 2 (3,5) 3 6 6

*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



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)



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	LD	F2,	34(R2)
2	LD	F4,	45(R3)
3	MULTD	F6,	F4, F2
4	SUBD	F8,	F2, F2
5	DIVD	F4,	F2, F8
6	ADDD	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




Data-Driven Execution

*Renaming
table &
reg file*

*Reorder
buffer*

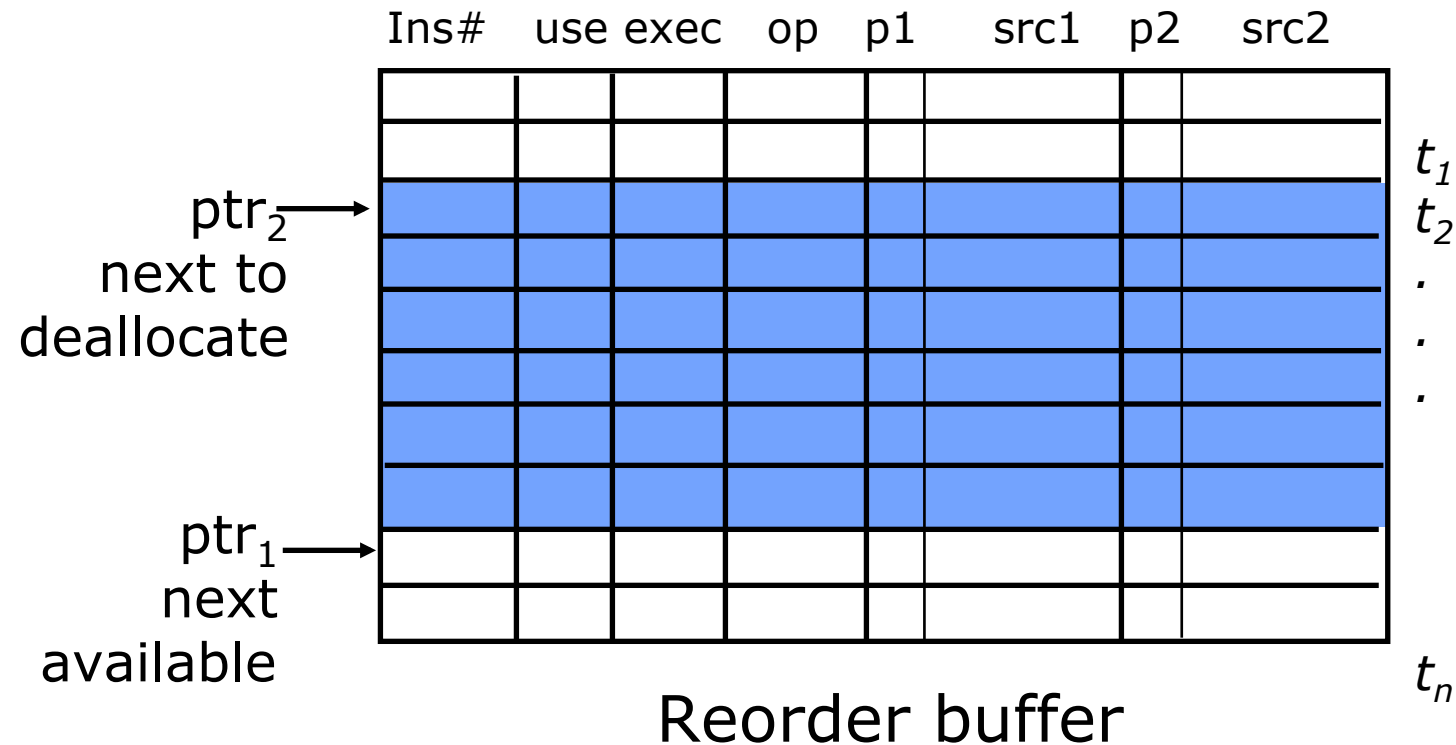
Replacing the
tag by its value
is an expensive
operation



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



Simplifying Allocation/Deallocation



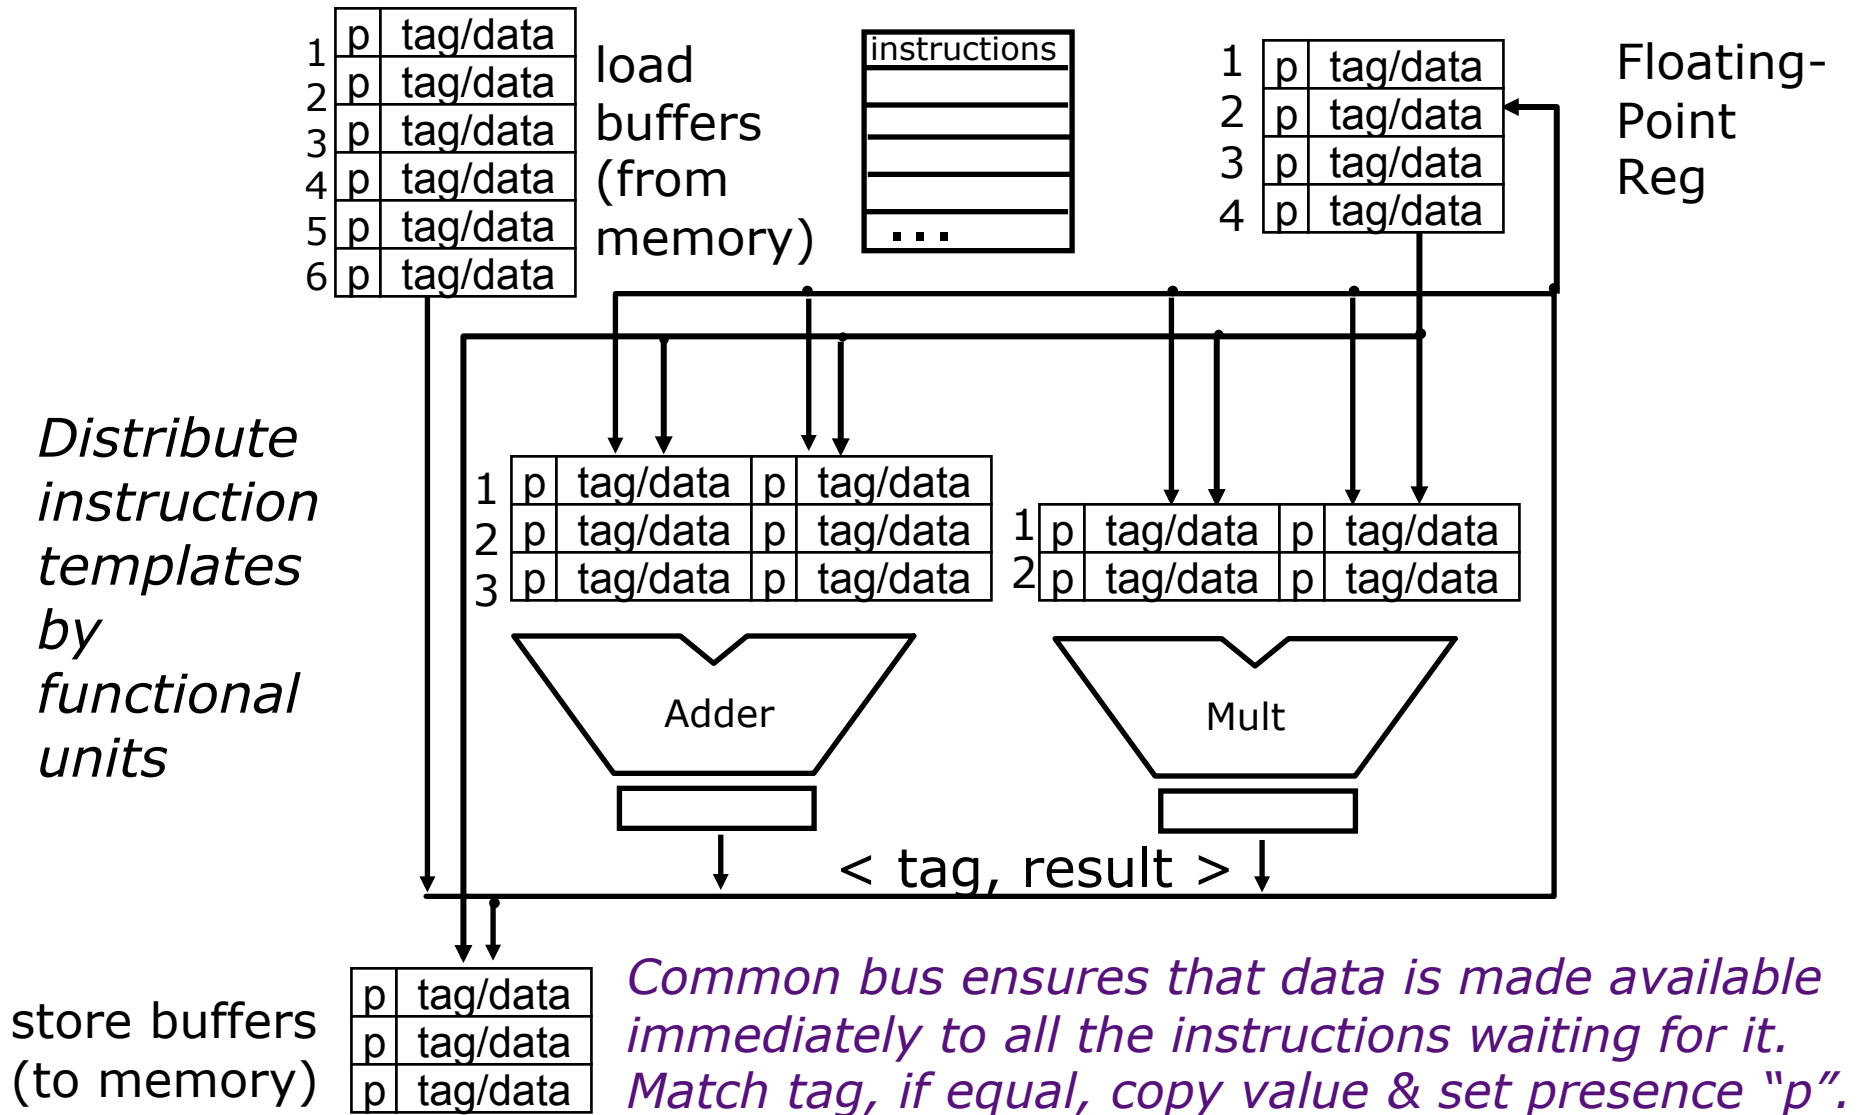
Instruction buffer is managed circularly

- "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 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} .



Effect on Interrupts

Out-of-order Completion

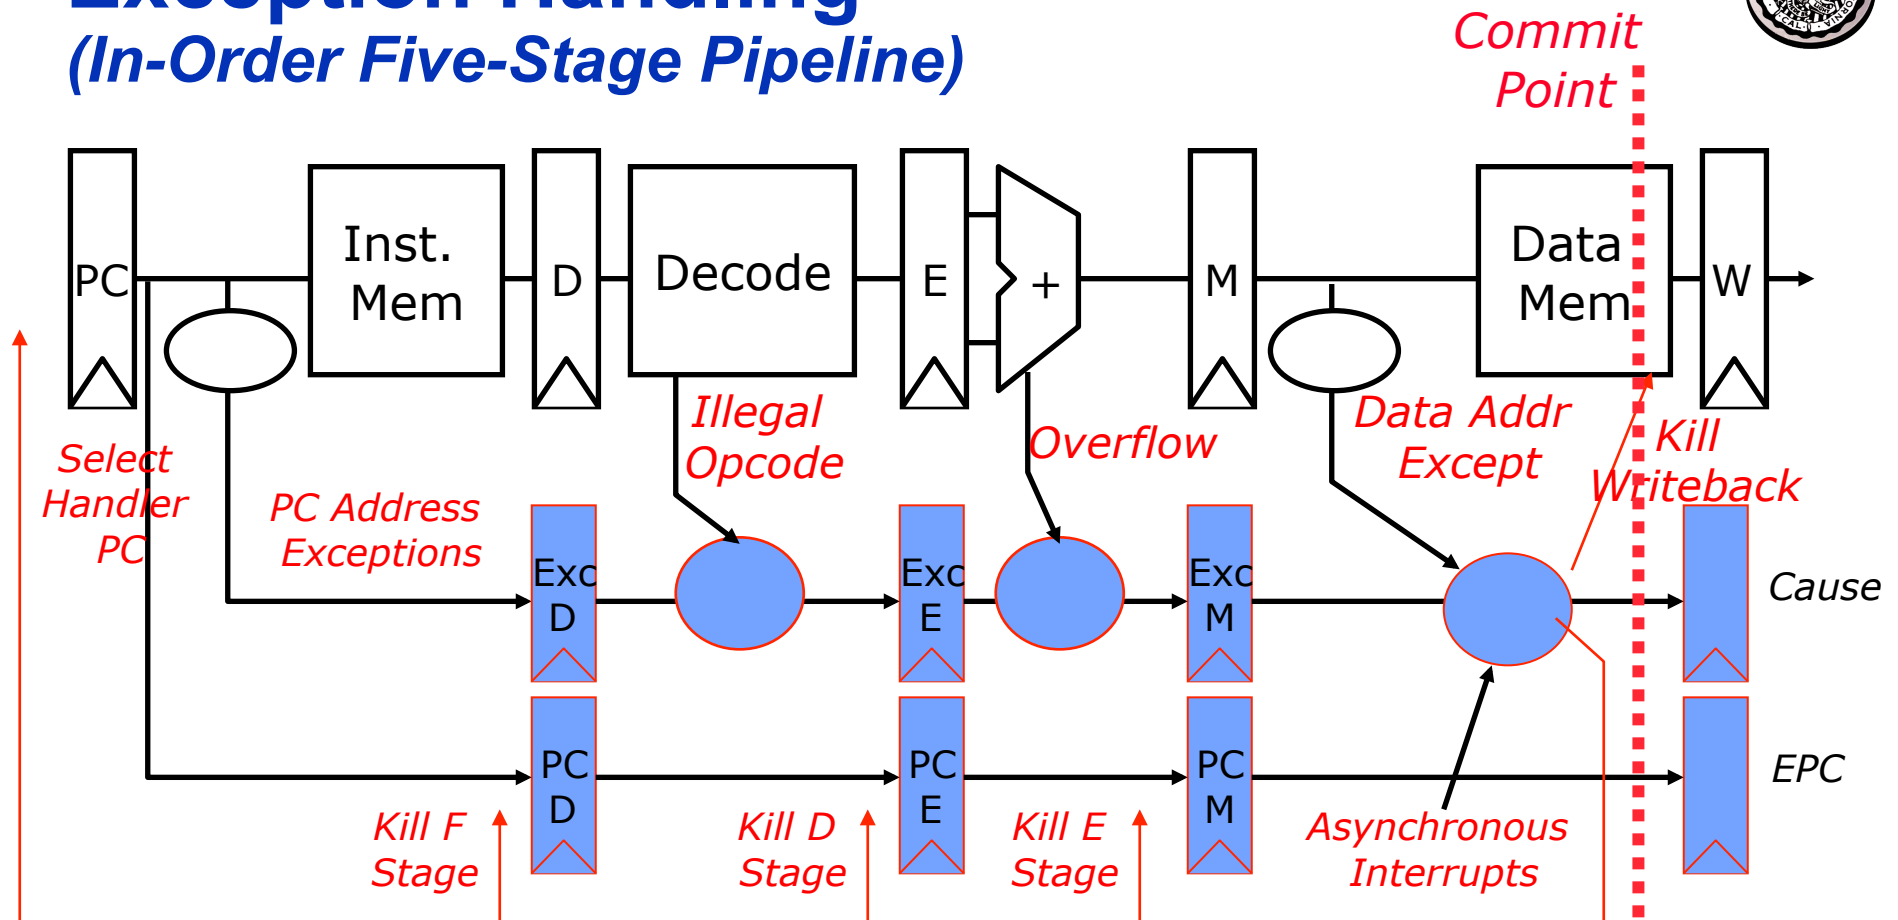
I_1	DIVD	f6,	f6,	f4
I_2	LD	f2,	45(r3)	
I_3	MULTD	f0,	f2,	f4
I_4	DIVD	f8,	f6,	f2
I_5	SUBD	f10,	f0,	f6
I_6	ADDD	f6,	f8,	f2



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



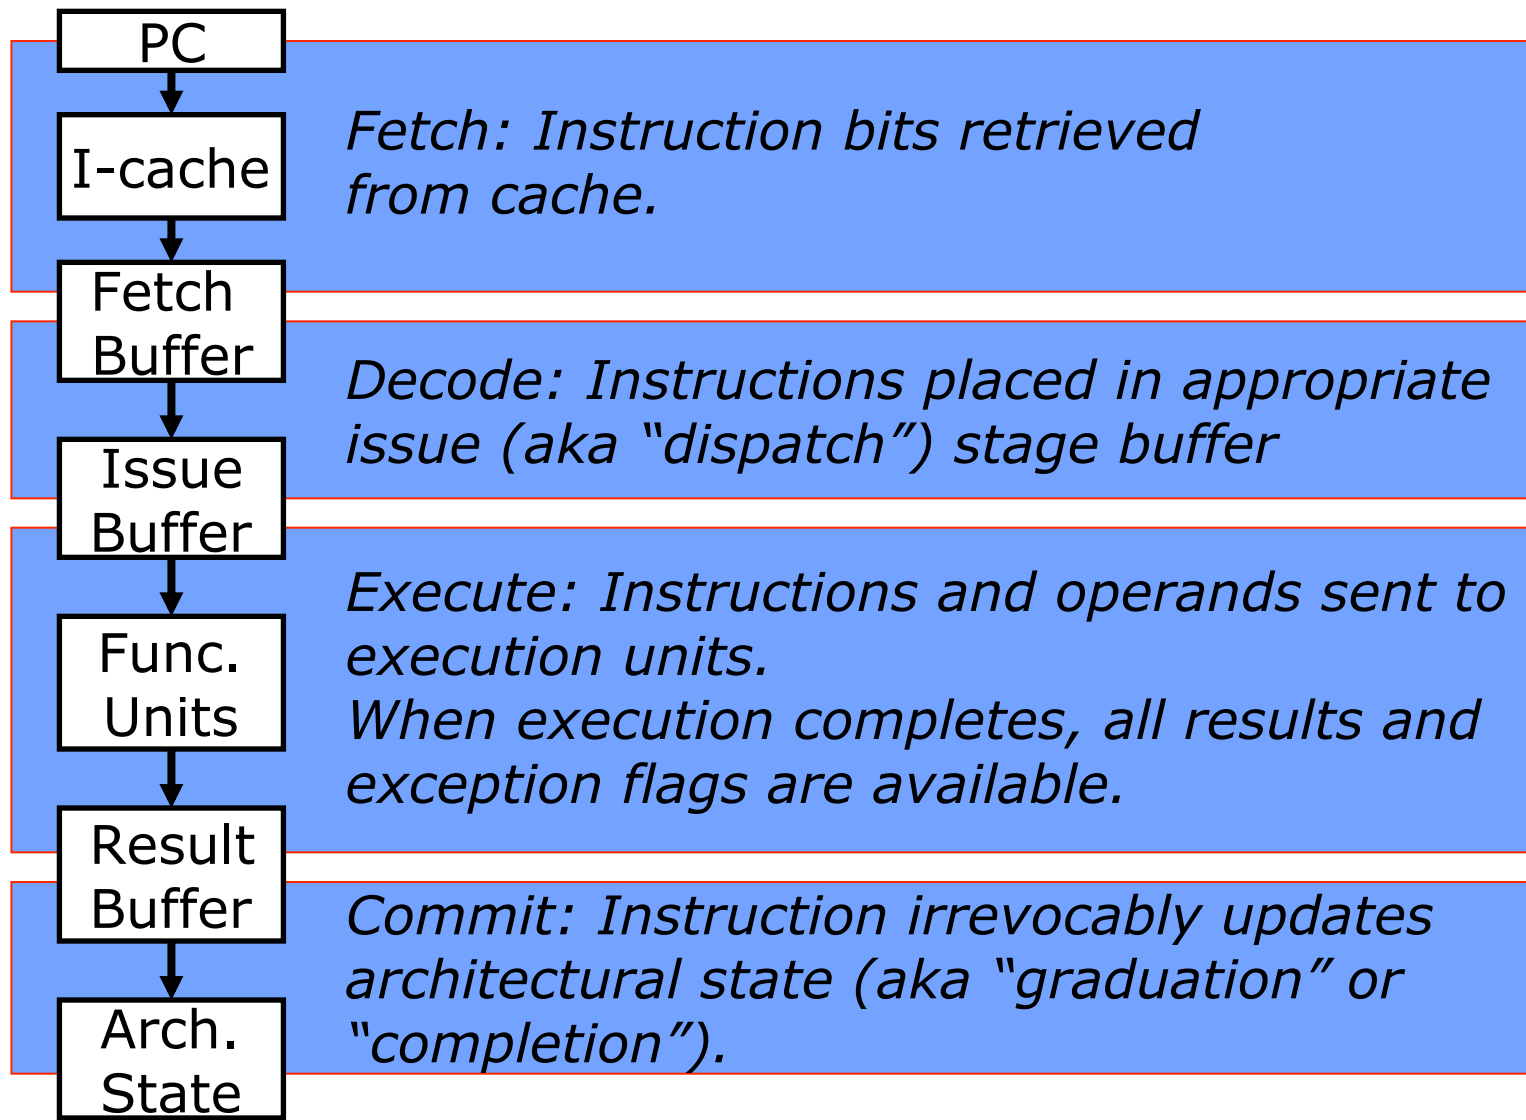
Exception Handling (In-Order Five-Stage Pipeline)



- 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

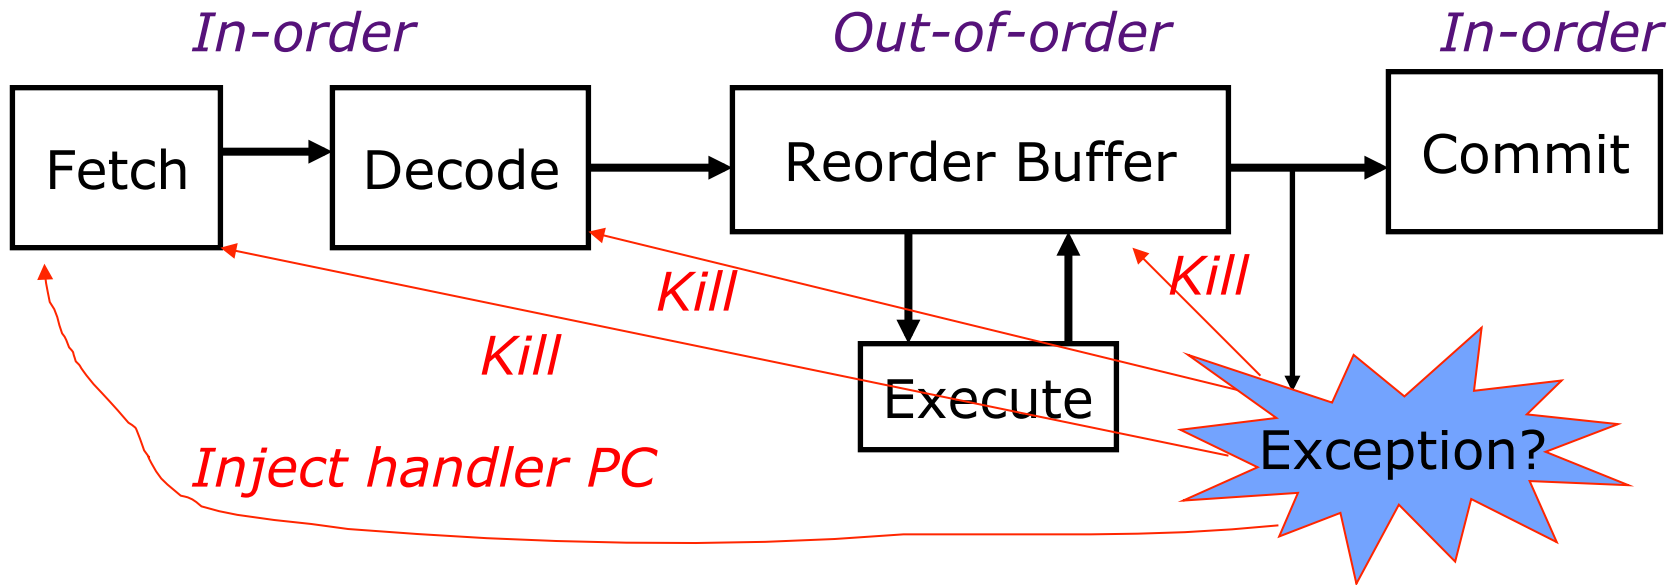


Phases of Instruction Execution





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)



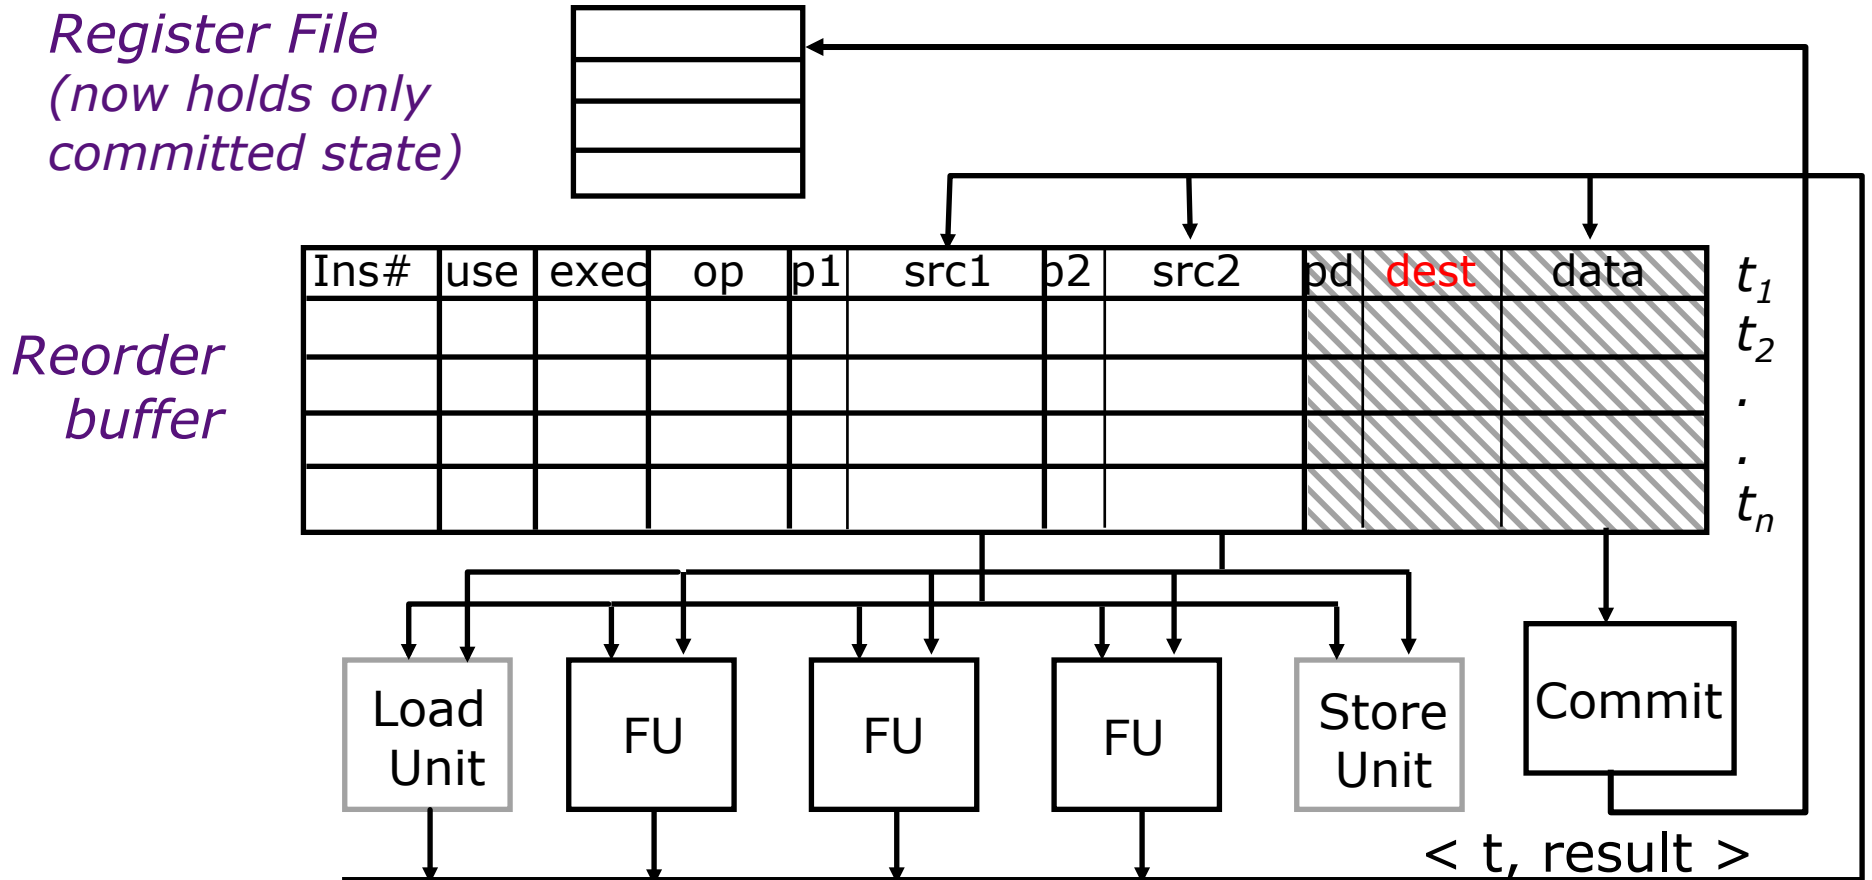
Extensions for Precise Exceptions



- add <pd, dest, data, cause> fields in the instruction template
- commit instructions to reg file and memory in program order ⇒ buffers can be maintained circularly
- on exception, clear reorder buffer by resetting ptr₁=ptr₂
(stores must wait for commit before updating memory)



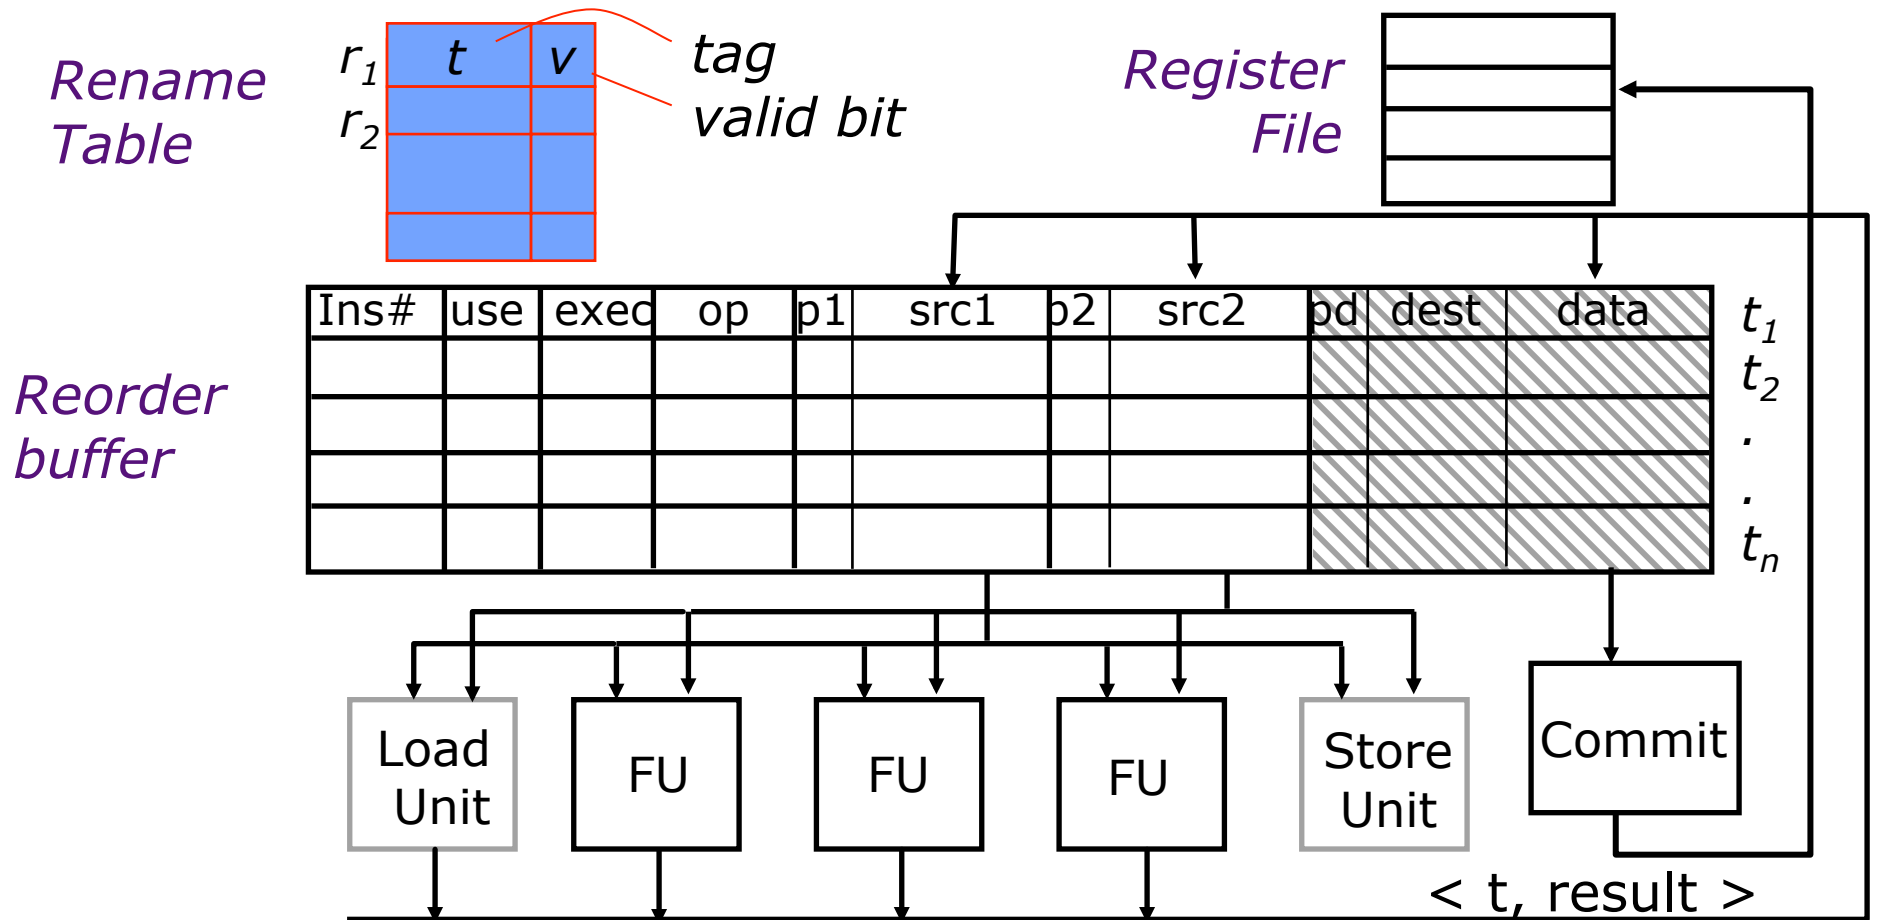
Rollback and Renaming



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*



CS152 Administrivia

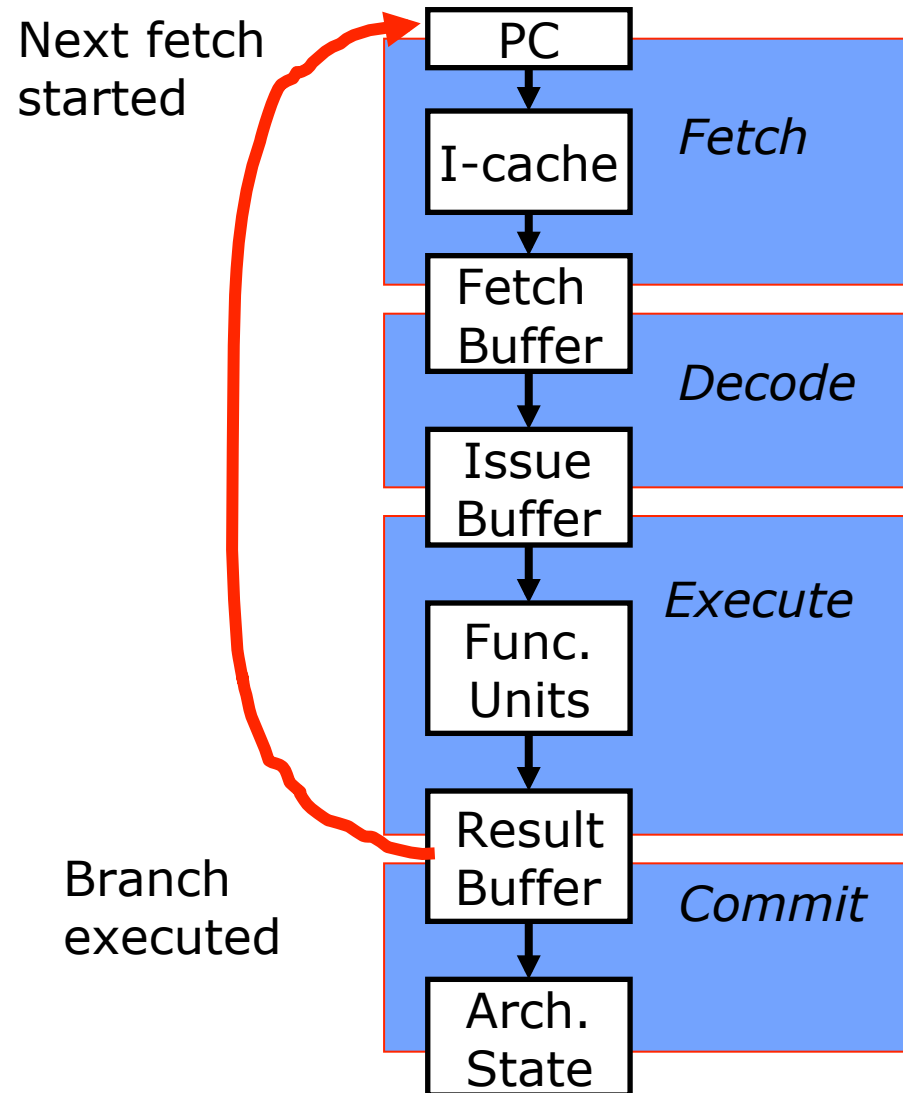


Control Flow Penalty

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





MIPS Branches and Jumps

Each instruction fetch depends on one or two pieces of information from the preceding instruction:

- 1) Is the preceding instruction a taken branch?
- 2) If so, what is the target address?

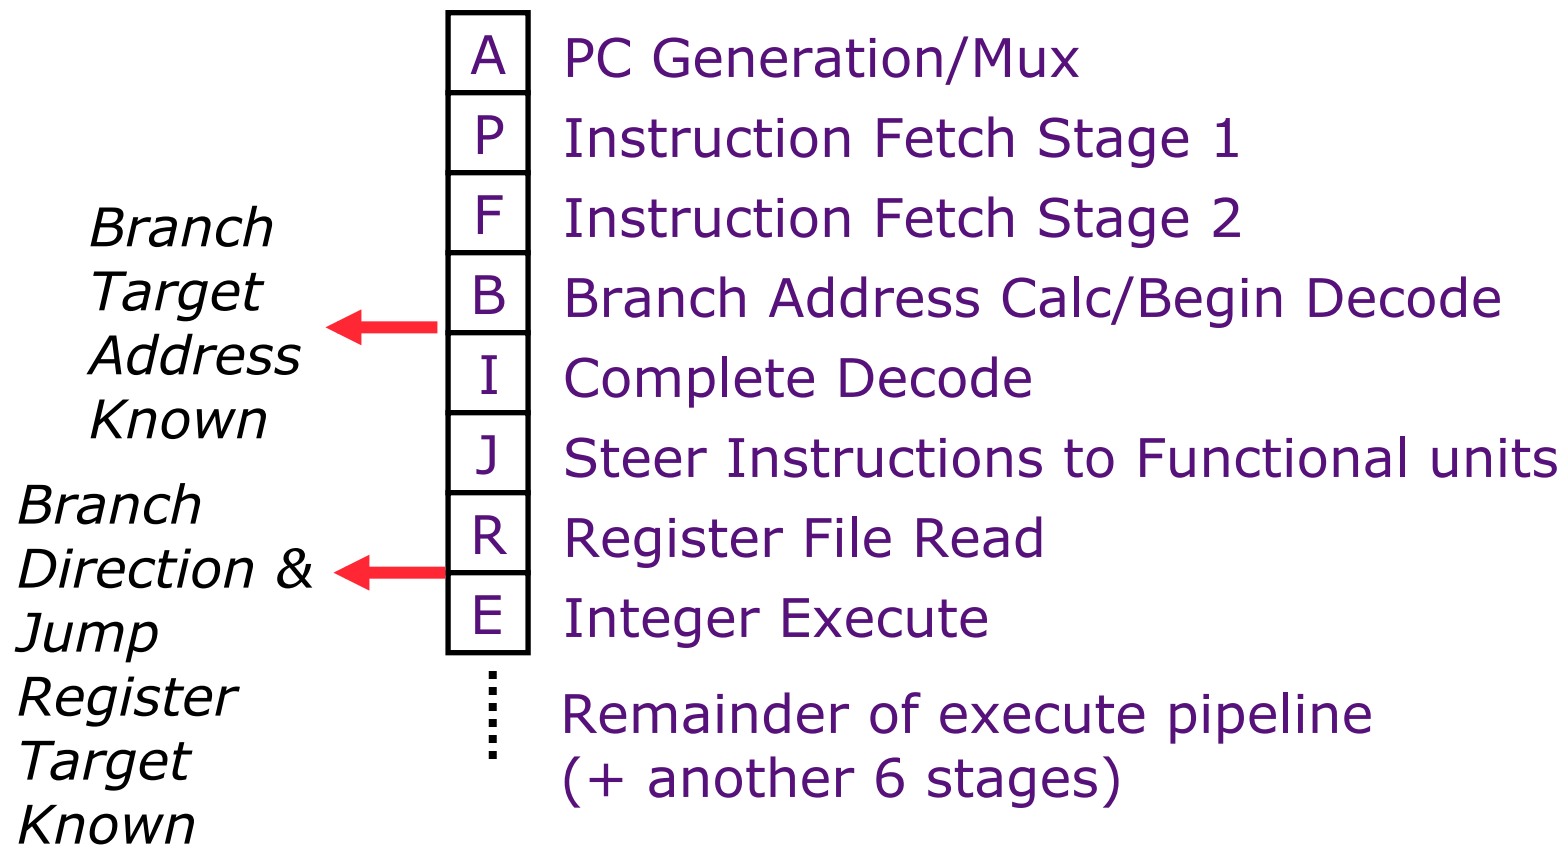
<i>Instruction</i>	<i>Taken known?</i>	<i>Target known?</i>
J	After Inst. Decode	After Inst. Decode
JR	After Inst. Decode	After Reg. Fetch
BEQZ/BNEZ	After Reg. Fetch*	After Inst. Decode

*Assuming zero detect on register read



Branch Penalties in Modern Pipelines

UltraSPARC-III instruction fetch pipeline stages
(in-order issue, 4-way superscalar, 750MHz, 2000)





Reducing Control Flow Penalty

Software solutions

- *Eliminate branches - loop unrolling*
Increases the run length
- *Reduce resolution time - instruction scheduling*
Compute the branch condition as early as possible (of limited value)

Hardware solutions

- Find something else to do - *delay slots*
Replaces pipeline bubbles with useful work (requires software cooperation)
- *Speculate - branch prediction*
Speculative execution of instructions beyond the branch



Branch Prediction

Motivation:

Branch penalties limit performance of deeply pipelined processors

Modern branch predictors have high accuracy (>95%) and can reduce branch penalties significantly

Required hardware support:

Prediction structures:

- Branch history tables, branch target buffers, etc.


Mispredict recovery mechanisms:

- Keep result computation separate from commit
- Kill instructions following branch in pipeline
- Restore state to state following branch



Static Branch Prediction

Overall probability a branch is taken is $\sim 60-70\%$ but:



ISA can attach preferred direction semantics to branches,
e.g., Motorola MC88110

`bne0` (*preferred taken*) `beq0` (*not taken*)

ISA can allow arbitrary choice of statically predicted direction,
e.g., HP PA-RISC, Intel IA-64

typically reported as $\sim 80\%$ accurate



Dynamic Branch Prediction

learning based on past behavior

Temporal correlation

The way a branch resolves may be a good predictor of the way it will resolve at the next execution


Spatial correlation

Several branches may resolve in a highly correlated manner (*a preferred path of execution*)



Branch Prediction Bits

- Assume 2 BP bits per instruction
- Change the prediction after two consecutive mistakes!

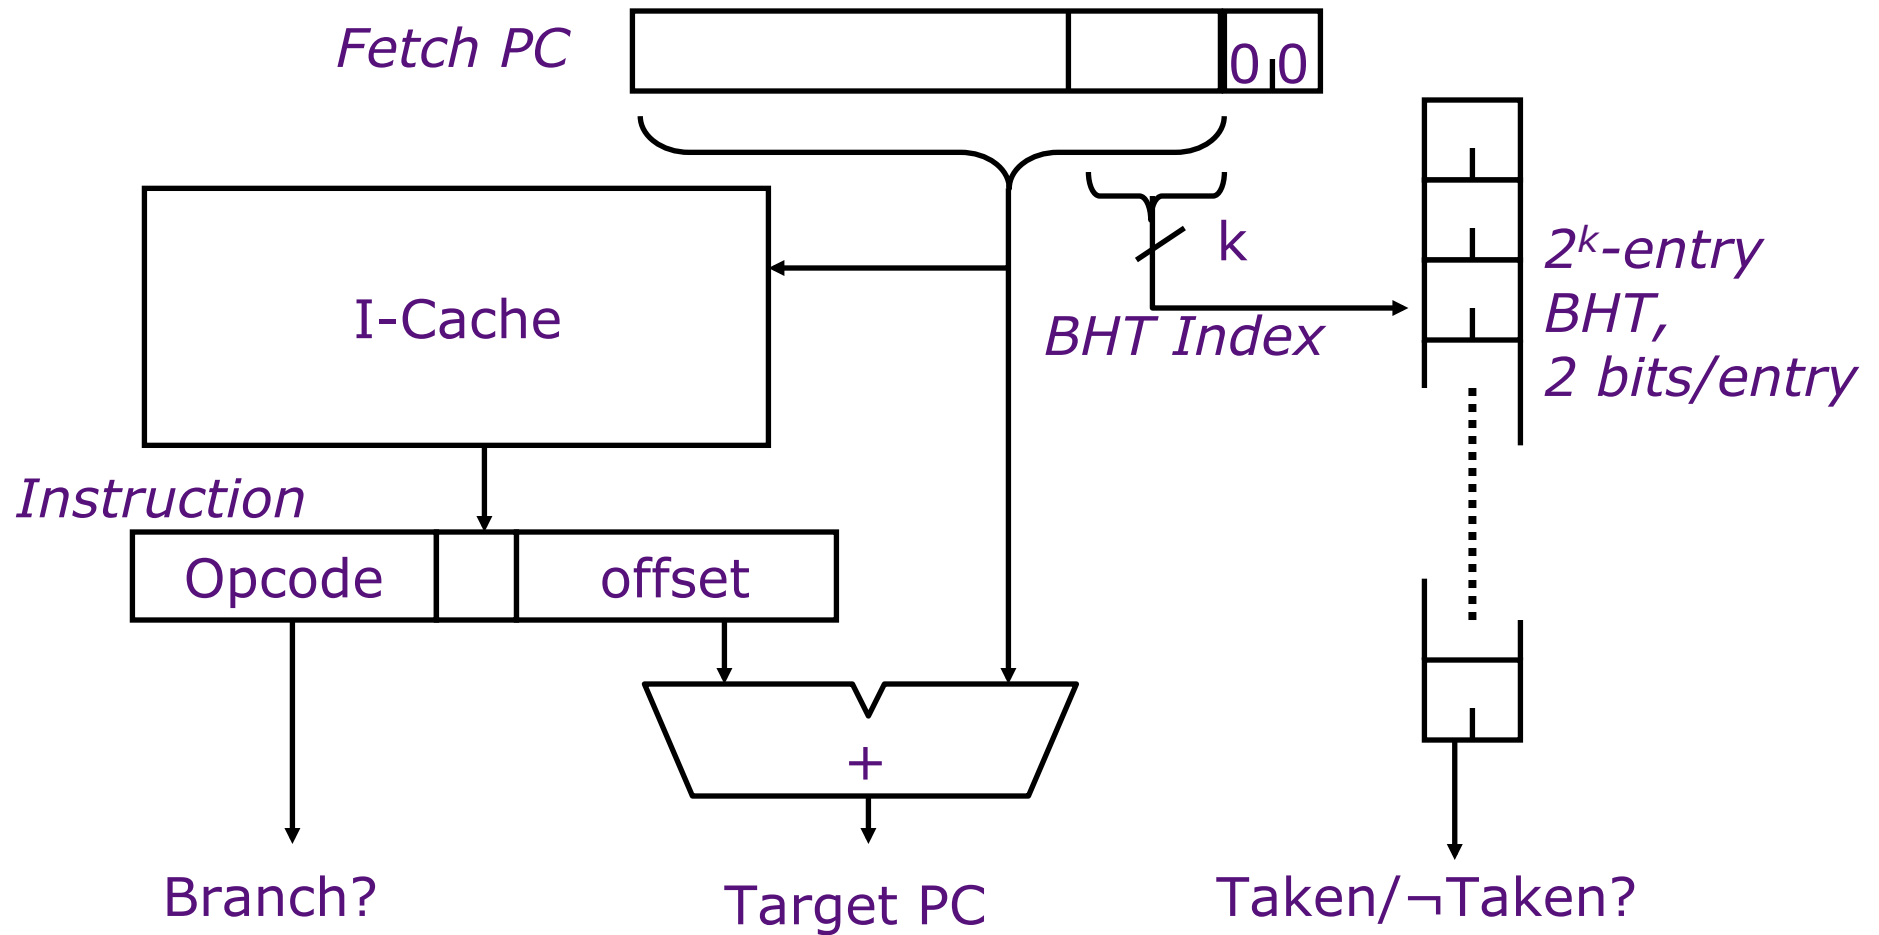


BP state:

(predict take/¬take) x (last prediction right/wrong)



Branch History Table



4K-entry BHT, 2 bits/entry, ~80-90% correct predictions



Exploiting Spatial Correlation

Yeh and Patt, 1992

```
if (x[i] < 7) then
    y += 1;
if (x[i] < 5) then
    c -= 4;
```

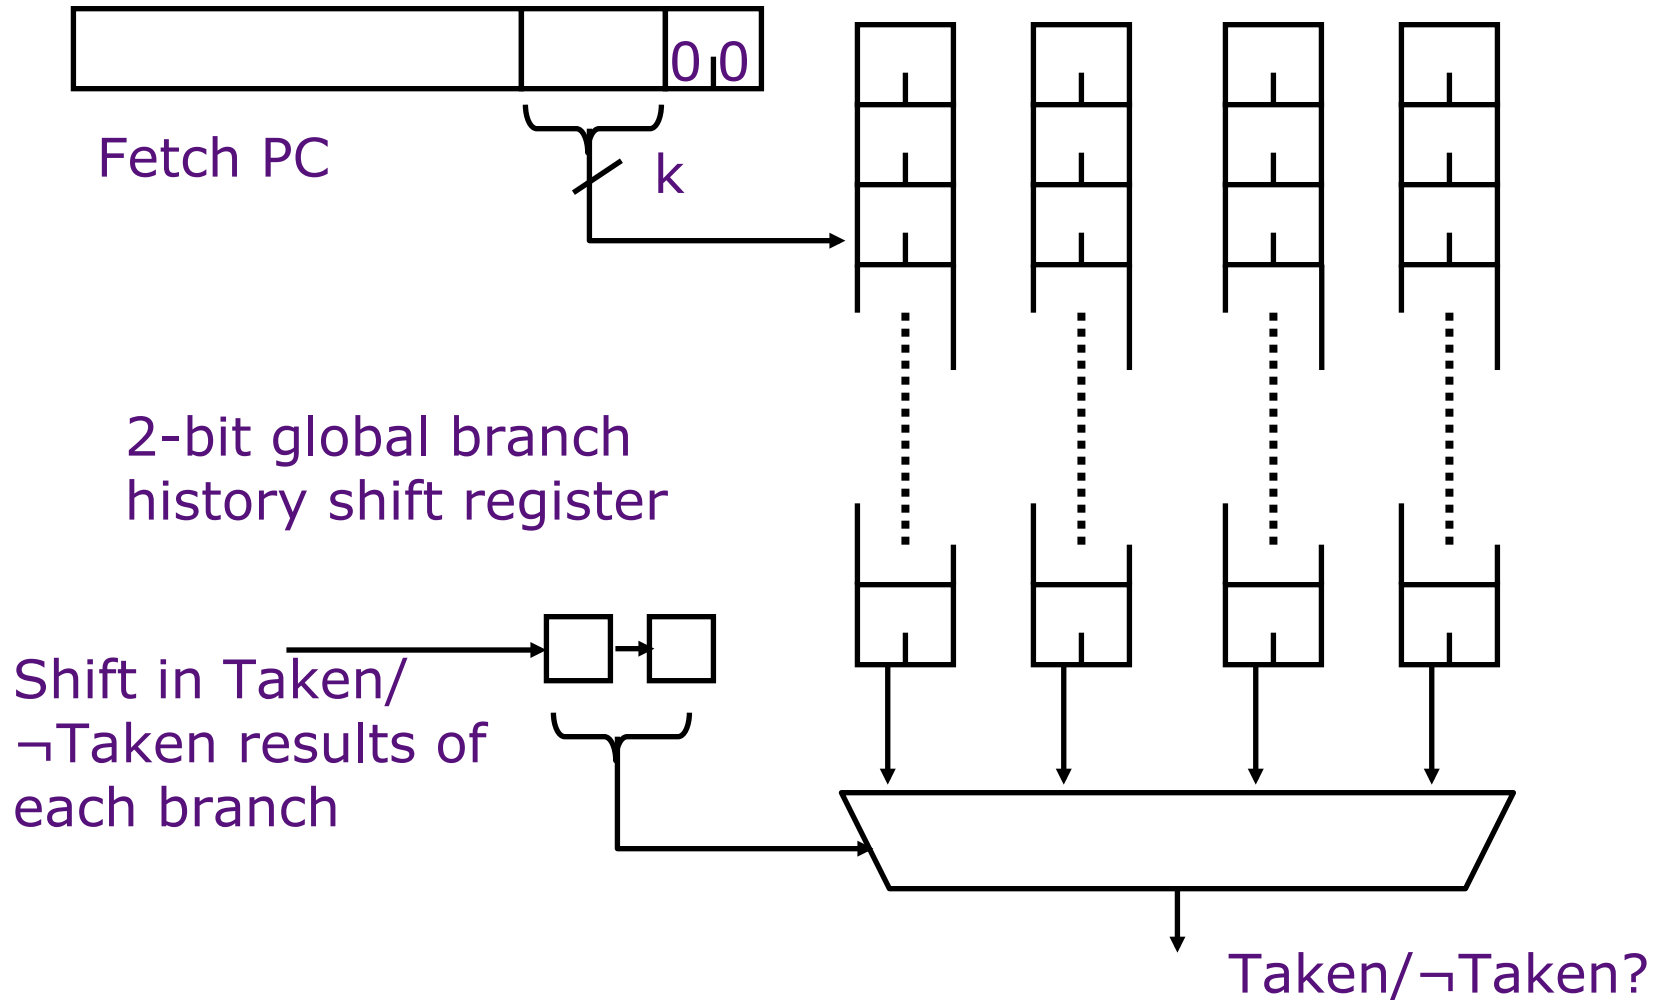
If first condition false, second condition also false

History register, H, records the direction of the last N branches executed by the processor



Two-Level Branch Predictor

Pentium Pro uses the result from the last two branches to select one of the four sets of BHT bits (~95% correct)





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