

CS252 Graduate Computer Architecture Fall 2015

Lecture 6: Modern Out-of-Order Processors

Krste Asanovic

`krste@eecs.berkeley.edu`

`http://inst.eecs.berkeley.edu/~cs252/fa15`



Supercomputers

Definitions of a supercomputer:

- Fastest machine in world at given task
 - A device to turn a compute-bound problem into an I/O bound problem
 - Any machine costing \$30M+
 - Any machine designed by Seymour Cray
-
- CDC6600 (Cray, 1964) regarded as first supercomputer

CDC 6600 *Seymour Cray, 1963*



- A fast pipelined machine with 60-bit words
 - 128 Kword main memory capacity, 32 banks
- Ten functional units (parallel, unpipelined)
 - Floating Point: adder, 2 multipliers, divider
 - Integer: adder, 2 incrementers, ...
- Hardwired control (no microcoding)
- *Scoreboard* for dynamic scheduling of instructions
- Ten Peripheral Processors for Input/Output
 - a fast multi-threaded 12-bit integer ALU
- Very fast clock, 10 MHz (FP add in 4 clocks)
- >400,000 transistors, 750 sq. ft., 5 tons, 150 kW, novel freon-based technology for cooling
- Fastest machine in world for 5 years (until 7600)
 - over 100 sold (\$7-10M each)



CDC 6600: A Load/Store Architecture

- Separate instructions to manipulate three types of reg.
 - 8x60-bit data registers (X)
 - 8x18-bit address registers (A)
 - 8x18-bit index registers (B)

- All arithmetic and logic instructions are register-to-register

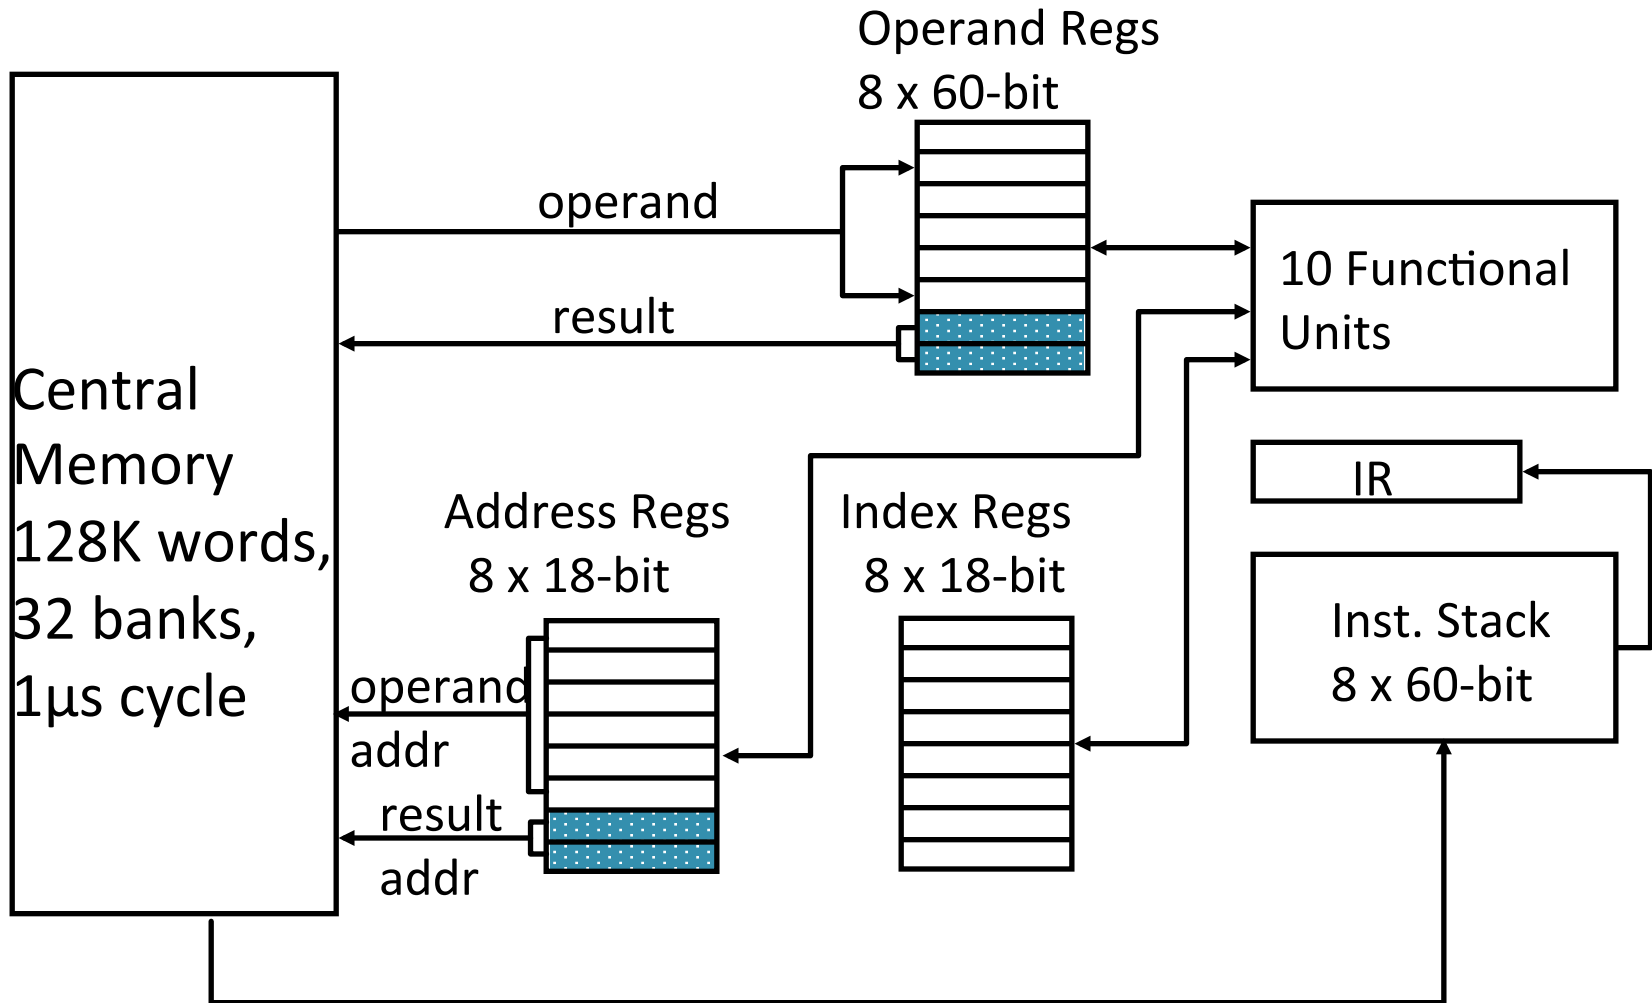


- Only Load and Store instructions refer to memory!



Touching address registers 1 to 5 initiates a load
 6 to 7 initiates a store
 - *very useful for vector operations*

CDC 6600: Datapath



CDC6600 ISA designed to simplify high-performance implementation

- Use of three-address, register-register ALU instructions simplifies pipelined implementation
 - Only 3-bit register specifier fields checked for dependencies
 - No implicit dependencies between inputs and outputs
- Decoupling setting of address register (Ar) from retrieving value from data register (Xr) simplifies providing multiple outstanding memory accesses
 - Software can schedule load of address register before use of value
 - Can interleave independent instructions inbetween
- CDC6600 has multiple parallel but unpipelined functional units
 - E.g., 2 separate multipliers
- Follow-on machine CDC7600 used pipelined functional units
 - Foreshadows later RISC designs

CDC6600: Vector Addition

```
      B0 ← - n
loop: JZE  B0, exit
      A0 ← B0 + a0      load X0
      A1 ← B0 + b0      load X1
      X6 ← X0 + X1
      A6 ← B0 + c0      store X6
      B0 ← B0 + 1
      jump loop
```

A_i = address register

B_i = index register

X_i = data register

CDC6600 Scoreboard

- Instructions dispatched in-order to functional units provided no structural hazard or WAW
 - Stall on structural hazard, no functional units available
 - Only one pending write to any register
- Instructions wait for input operands (RAW hazards) before execution
 - Can execute out-of-order
- Instructions wait for output register to be read by preceding instructions (WAR)
 - Result held in functional unit until register free

MEMORANDUM

August 28, 1963

Memorandum To: Messrs. A. L. Williams
T. V. Learson
H. W. Miller, Jr.
E. R. Piore
O. M. Scott
M. B. Smith
A. K. Watson

Last week CDC had a press conference during which they officially announced their 6600 system. I understand that in the laboratory developing this system there are only 34 people, "including the janitor." Of these, 14 are engineers and 4 are programmers, and only one person has a Ph. D., a relatively junior programmer. To the outsider, the laboratory appeared to be cost conscious, hard working and highly motivated.

Contrasting this modest effort with our own vast development activities, I fail to understand why we have lost our industry leadership position by letting someone else offer the world's most powerful computer. At Jenny Lake, I think top priority should be given to a discussion as to what we are doing wrong and how we should go about changing it immediately.

TJW, Jr:jmc

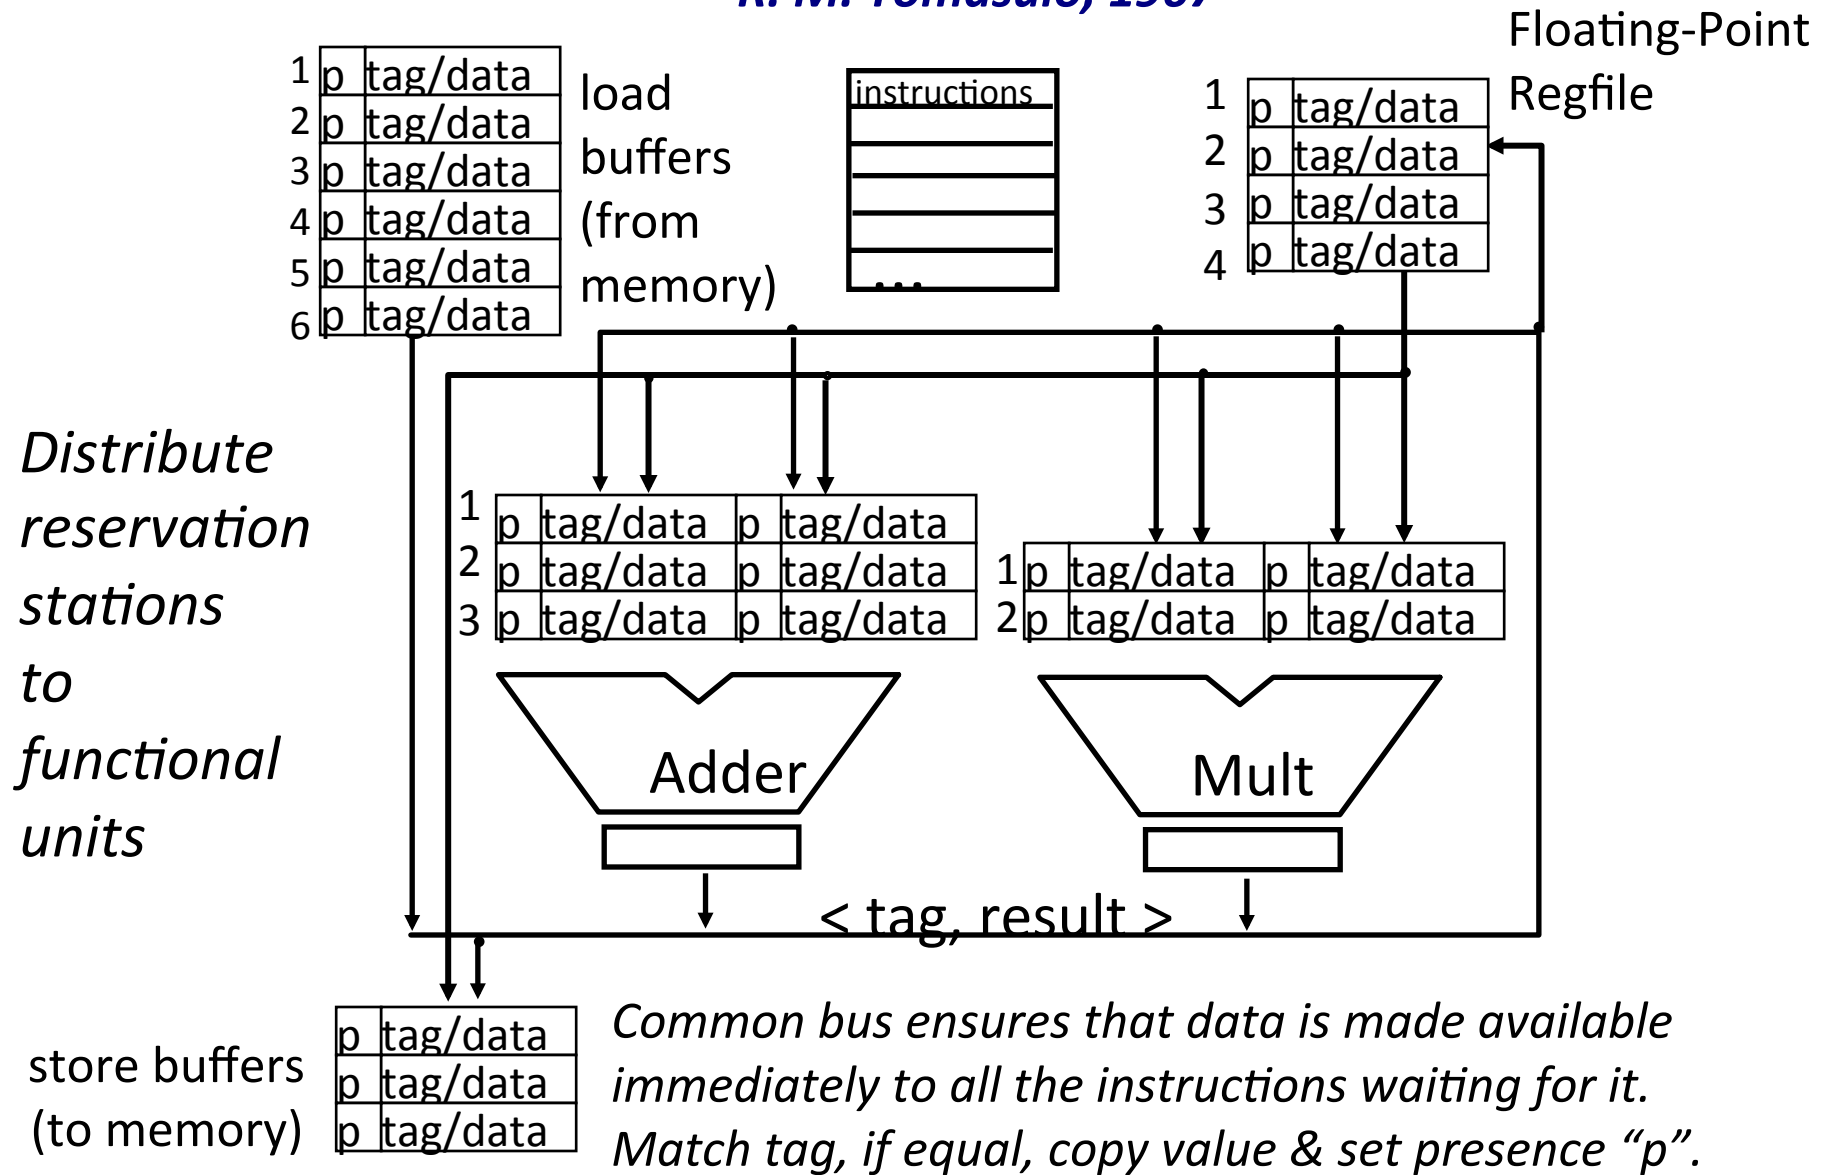
T. J. Watson, Jr.

cc: Mr. W. B. McWhirter

[© IBM]

IBM 360/91 Floating-Point Unit

R. M. Tomasulo, 1967



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

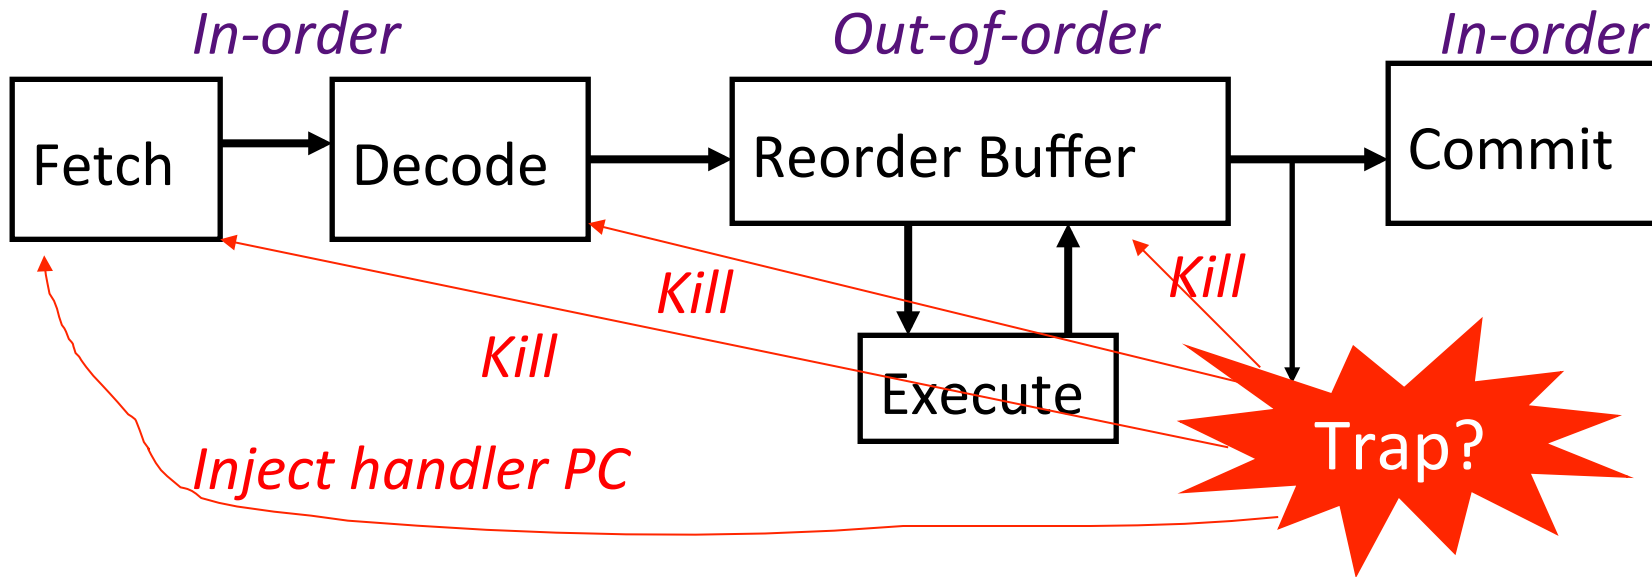
Also, simpler machine designs in new technology beat complicated machines in old technology

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

Separating Completion from Commit

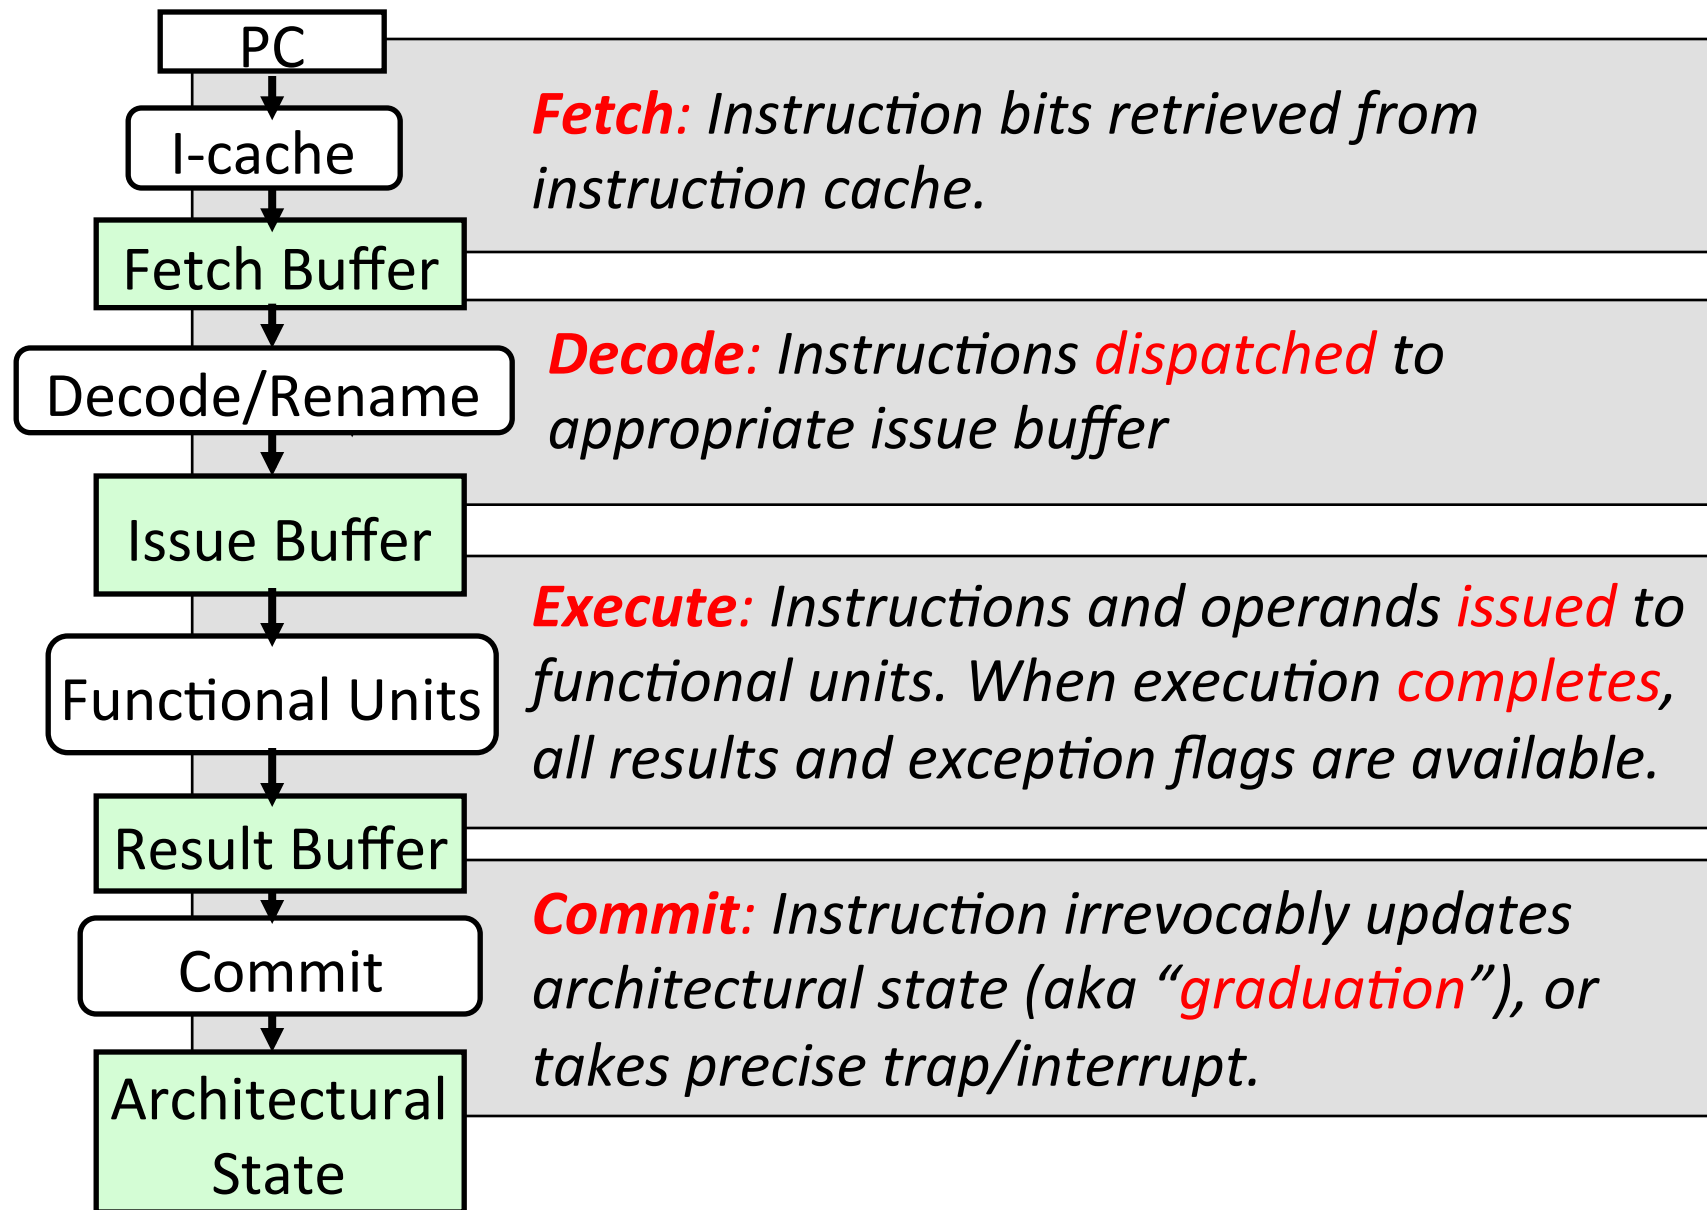
- Re-order buffer holds register results from completion until commit
 - Entries allocated in program order during decode
 - Buffers completed values and exception state until in-order commit point
 - Completed values can be used by dependents before committed (bypassing)
 - Each entry holds program 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

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

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

In-Order versus Out-of-Order Commit

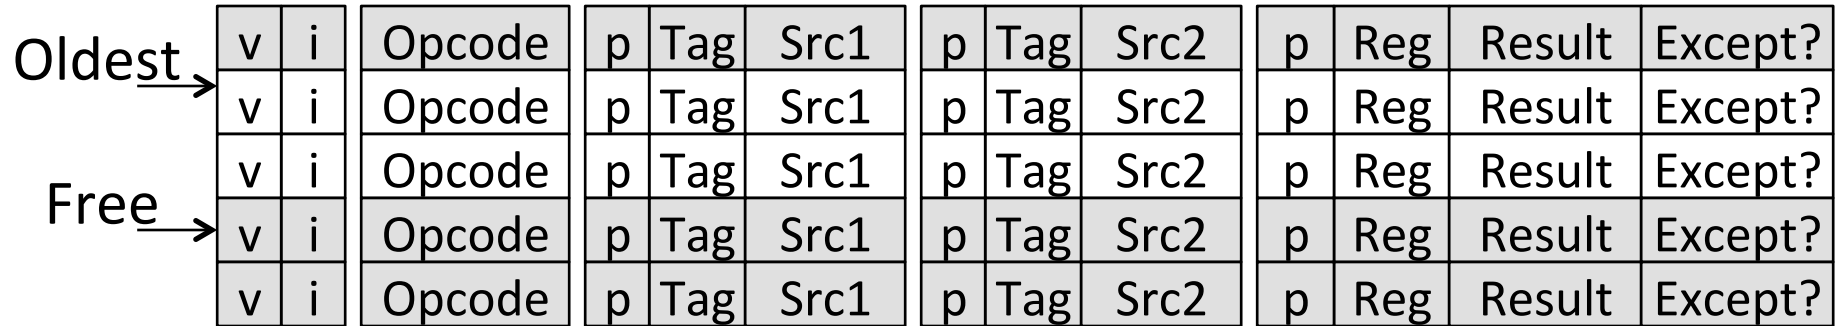
- In-order commit supports precise traps, standard today
 - Some proposals to reduce the cost of in-order commit by retiring some instructions early to compact reorder buffer, but this is just an optimized in-order commit
- Out-of-order commit was effectively what early OoO machines implemented (imprecise traps) as completion irrevocably changed machine state
 - i.e., complete == commit in these machines

OoO Design Choices

- Where are reservation stations?
 - Part of reorder buffer, or in separate issue window?
 - Distributed by functional units, or centralized?
- How is register renaming performed?
 - Tags and data held in reservation stations, with separate architectural register file
 - Tags only in reservation stations, data held in unified physical register file

“Data-in-ROB” Design

(HP PA8000, Pentium Pro, Core2Duo, Nehalem)



- Managed as circular buffer in program order, new instructions dispatched to free slots, oldest instruction committed/reclaimed when done (“p” bit set on result)
- Tag is given by index in ROB (Free pointer value)
- In dispatch, non-busy source operands read from architectural register file and copied to Src1 and Src2 with presence bit “p” set. Busy operands copy tag of producer and clear “p” bit.
- Set valid bit “v” on dispatch, set issued bit “i” on issue
- On completion, search source tags, set “p” bit and copy data into src on tag match. Write result and exception flags to ROB.
- On commit, check exception status, and copy result into architectural register file if no trap.
- On trap, flush machine and ROB, set free=oldest, jump to handler

Managing Rename for Data-in-ROB

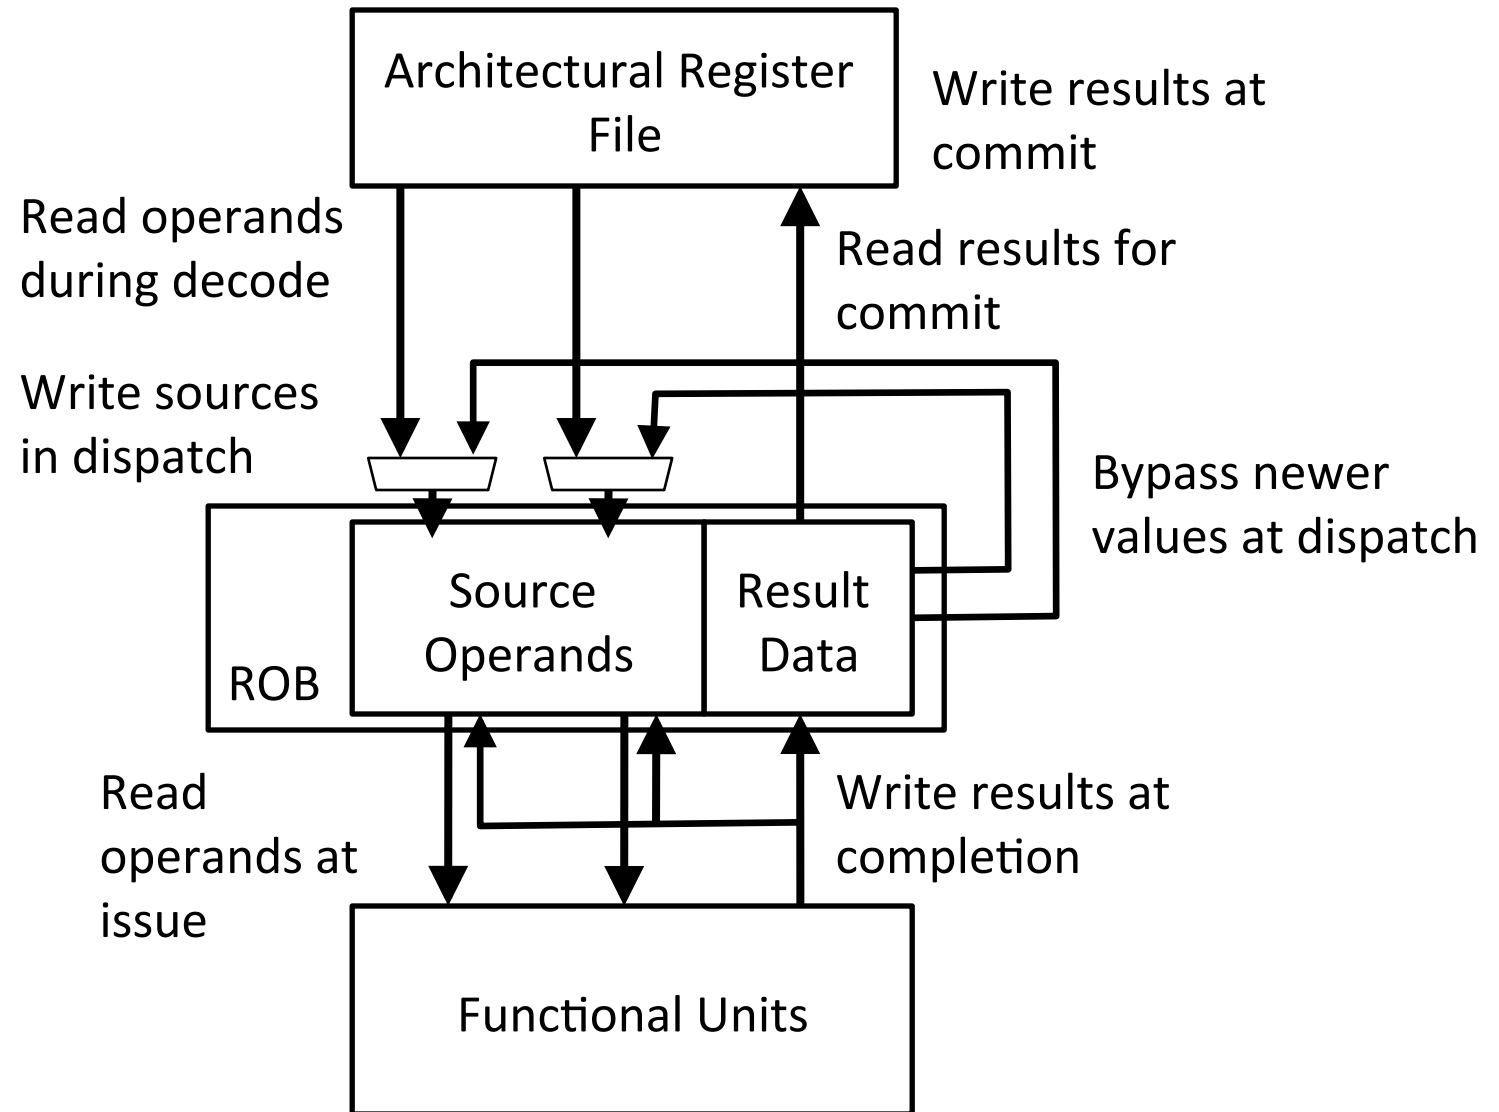
Rename table associated with architectural registers, managed in decode/dispatch

p	Tag	Value
p	Tag	Value
p	Tag	Value
p	Tag	Value

One entry per arch. register

- If “p” bit set, then use value in architectural register file
- Else, tag field indicates instruction that will/has produced value
- For dispatch, read source operands $\langle p, \text{tag}, \text{value} \rangle$ from arch. regfile, and also read $\langle p, \text{result} \rangle$ from producing instruction in ROB, bypassing as needed. Copy to ROB
- Write destination arch. register entry with $\langle 0, \text{Free}, _ \rangle$, to assign tag to ROB index of this instruction
- On commit, update arch. regfile with $\langle 1, _, \text{Result} \rangle$
- On trap, reset table (All $p=1$)

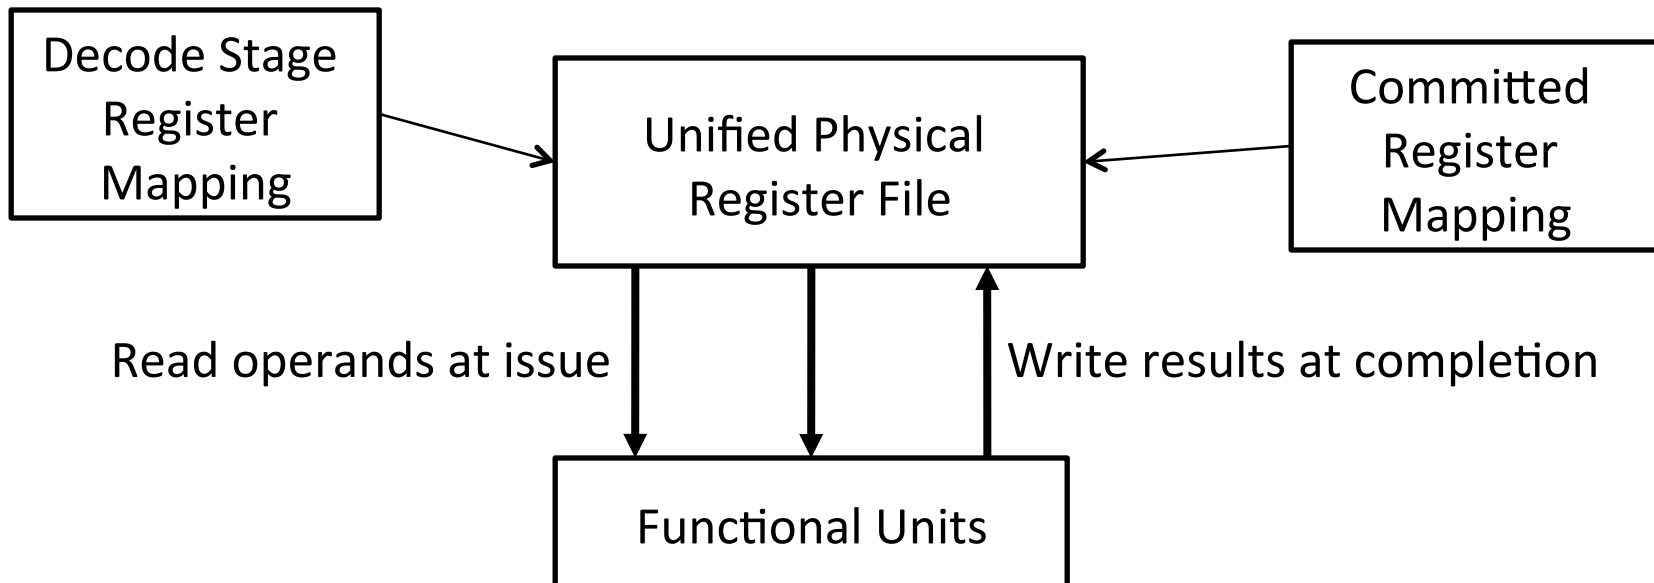
Data Movement in Data-in-ROB Design



Unified Physical Register File

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


- Rename all architectural registers into a single *physical* register file during decode, no register values read
- 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



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)
```

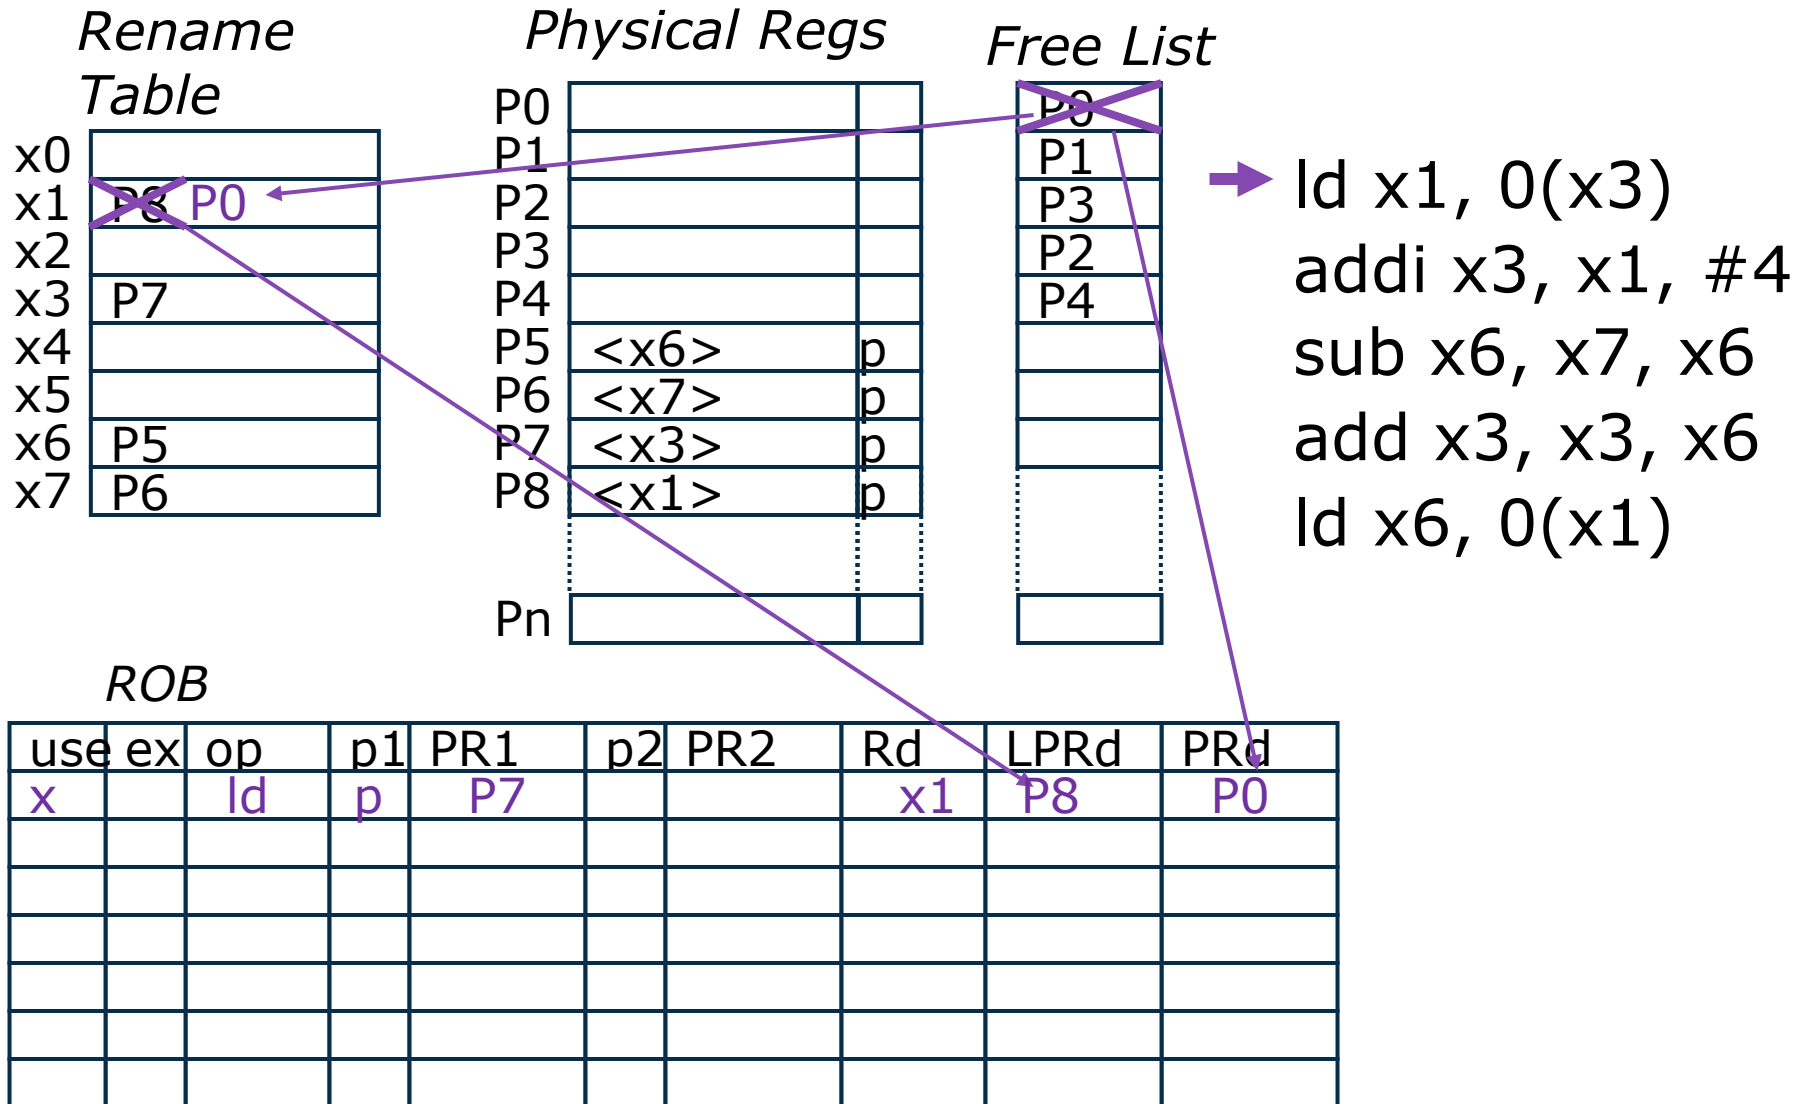


```
ld P1, (Px)
addi P2, P1, #4
sub P3, Py, Pz
add P4, P2, P3
ld P5, (P1)
add P6, P5, P4
sd P6, (P1)
ld P7, (Pw)
```

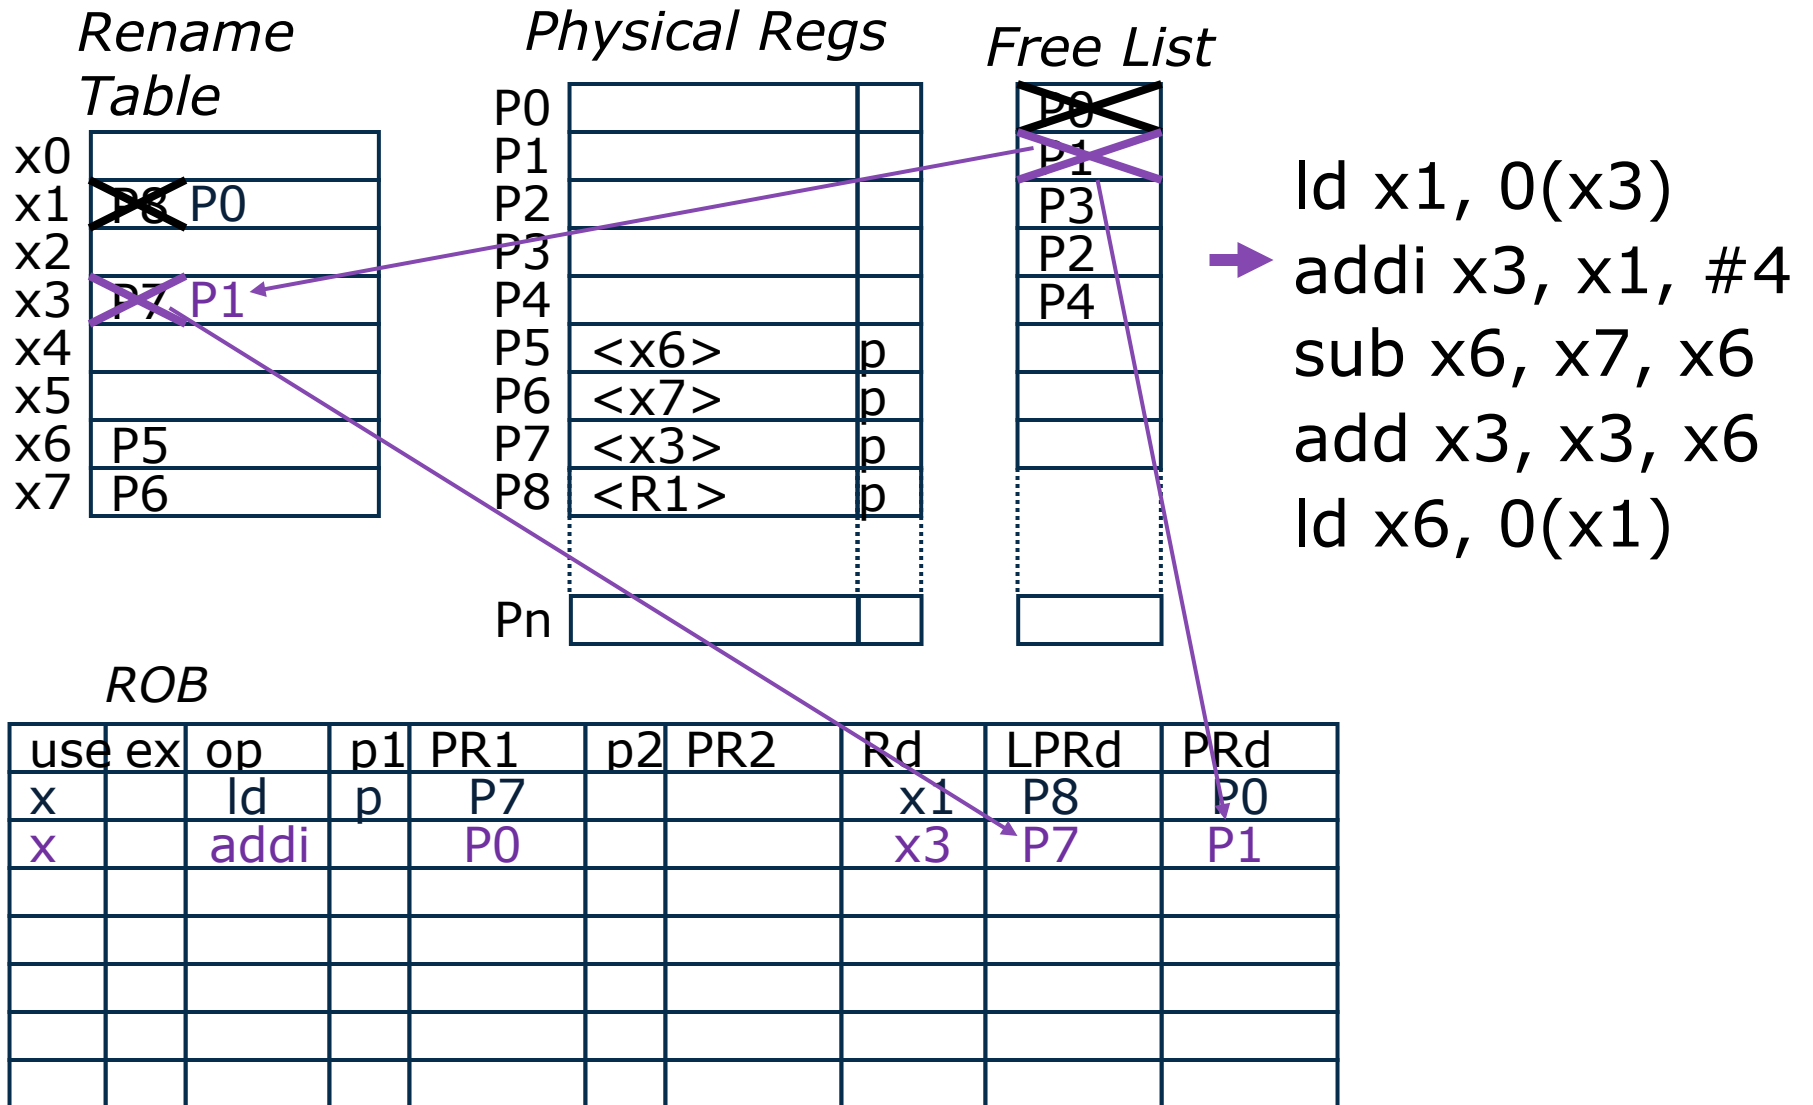
When can we reuse a physical register?

When next writer of same architectural register commits


Physical Register Management




Physical Register Management



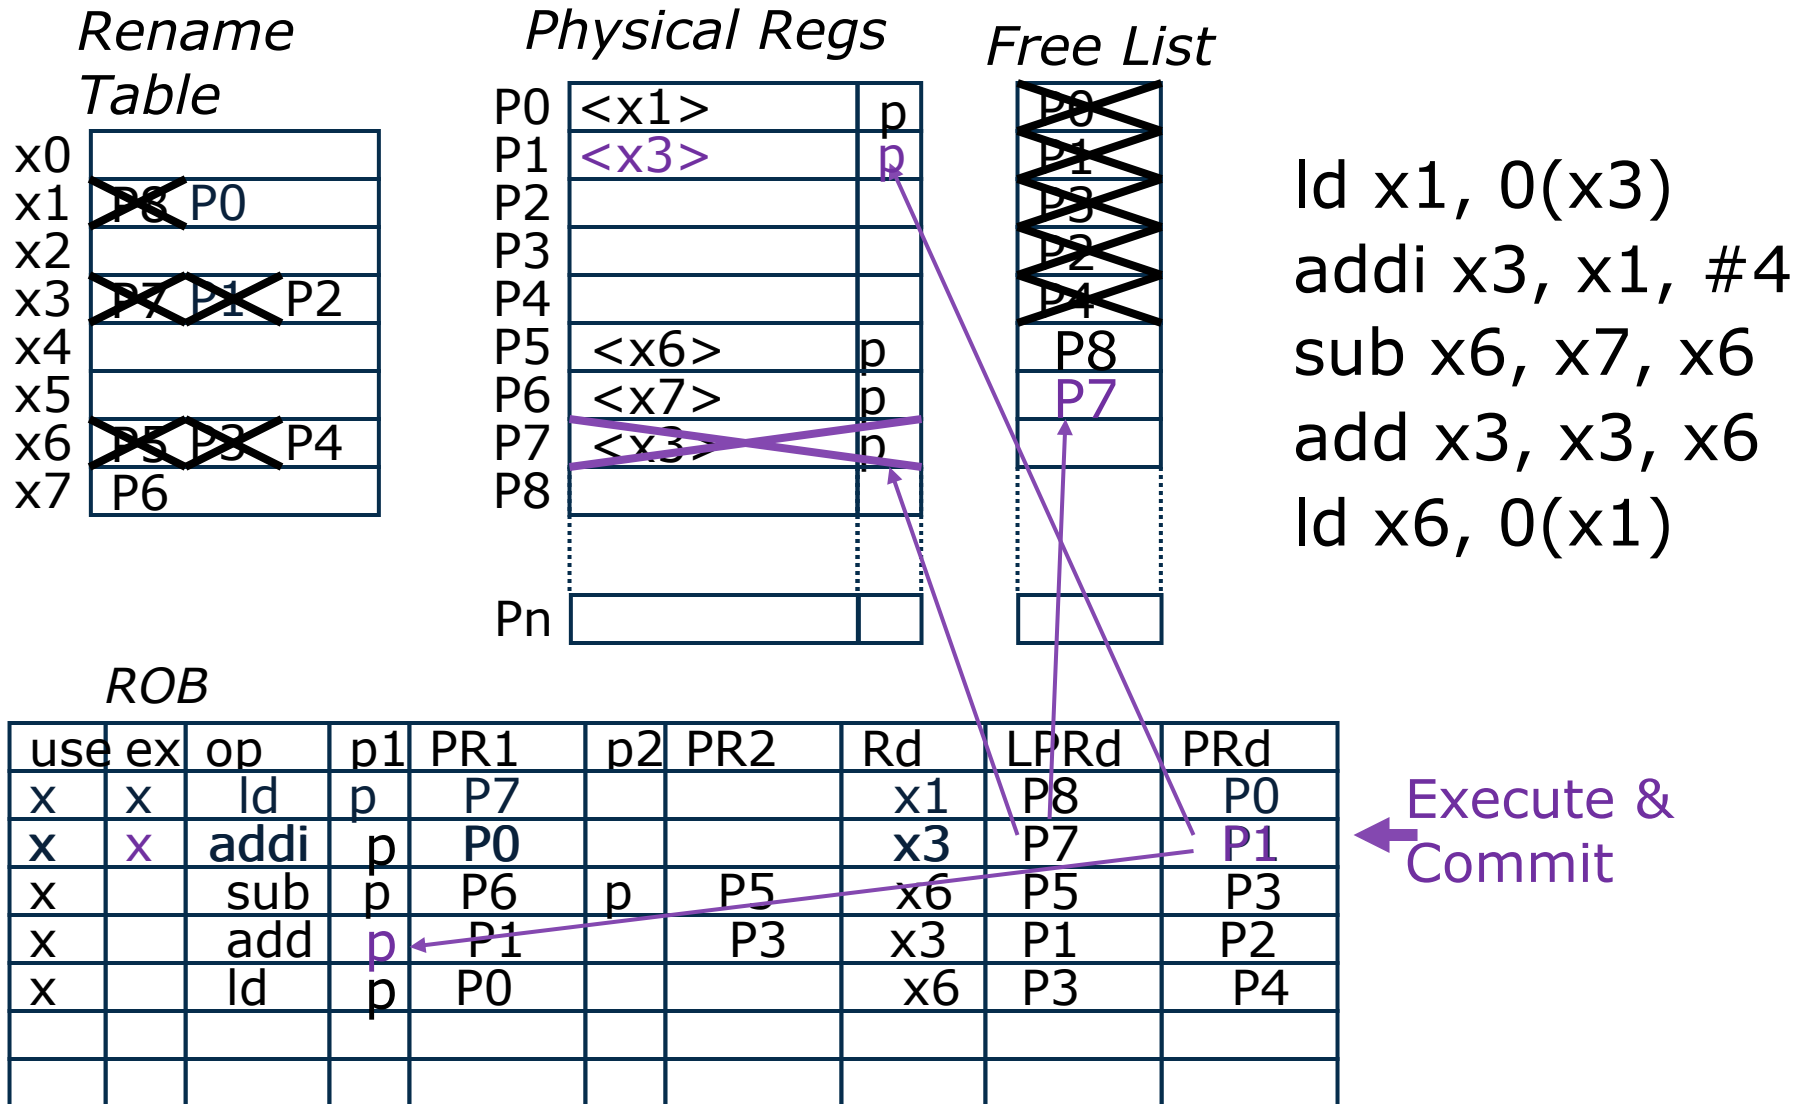
Physical Register Management



Physical Register Management



Physical Register Management



MIPS R10K Trap Handling

- Rename table is repaired by unrenaming instructions in reverse order using the PRd/LPRd fields
- The Alpha 21264 had similar physical register file scheme, but kept complete rename table snapshots for each instruction in ROB (80 snapshots total)
 - Flash copy all bits from snapshot to active table in one cycle

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)