

CS 152 Computer Architecture and Engineering

Lecture 12 - Advanced Out-of-Order Superscalars

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

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

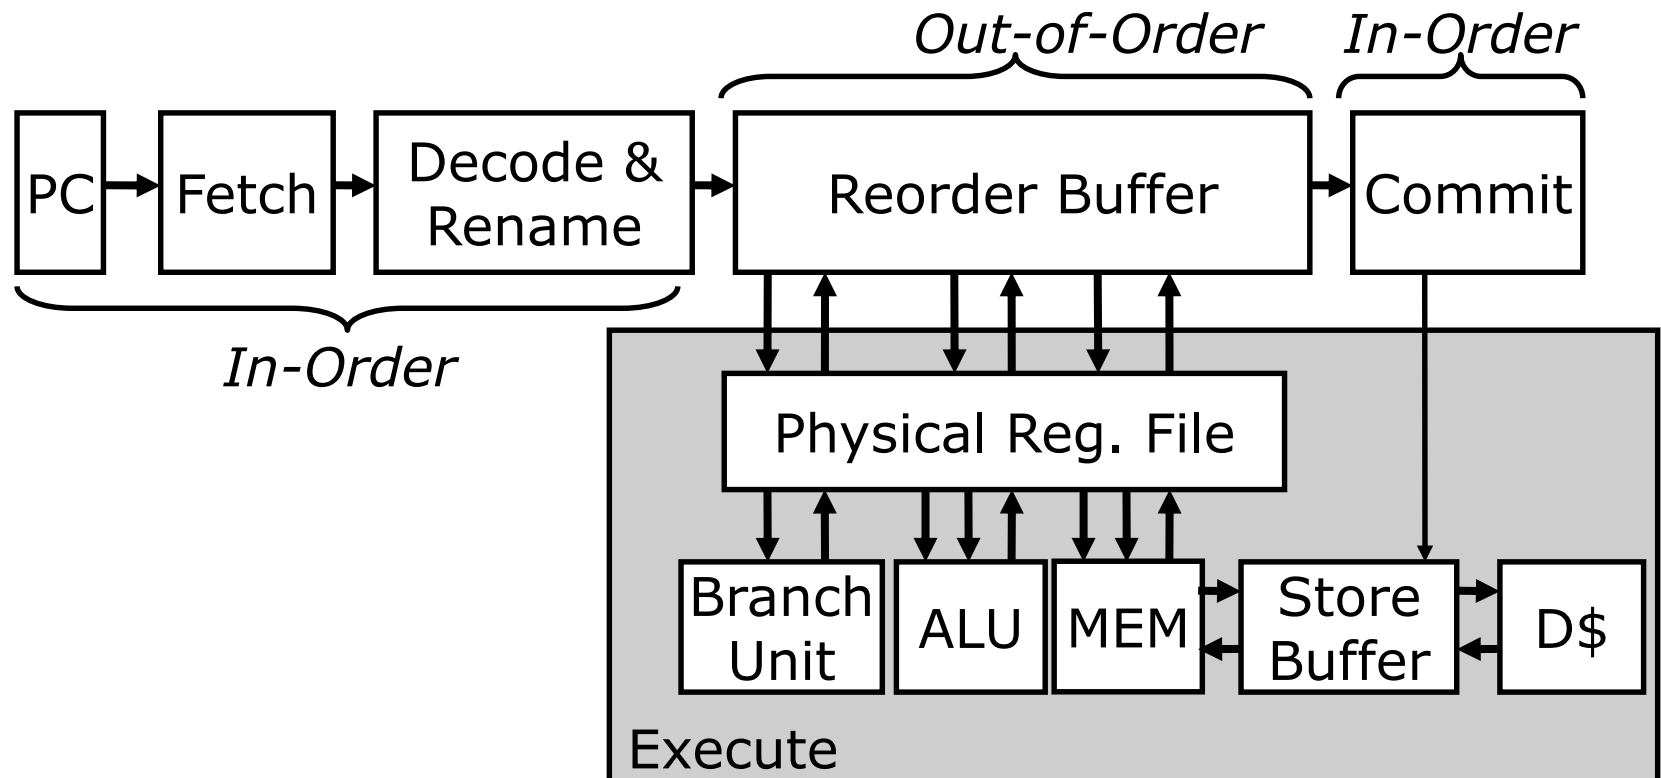
Last time in Lecture 11

- Register renaming removes WAR, WAW hazards
- In-order fetch/decode, out-of-order execute, in-order commit gives high performance and precise exceptions
- Need to rapidly recover on branch mispredictions
- Unified physical register file machines remove data values from ROB
 - All values only read and written during execution
 - Only register tags held in ROB

Question of the Day

- How many in-order cores do you think take up the same area as an out-of-order core?

Reminder



Separate Pending Instruction Window from ROB

Instructions that committed are not in the instruction window any more

The instruction window holds instructions that have been decoded and renamed but not issued into execution. Has register tags and presence bits, and pointer to ROB entry.

use	ex	op	p1	PR1	p2	PR2	PRd	ROB#

Reorder buffer used to hold exception information for commit.

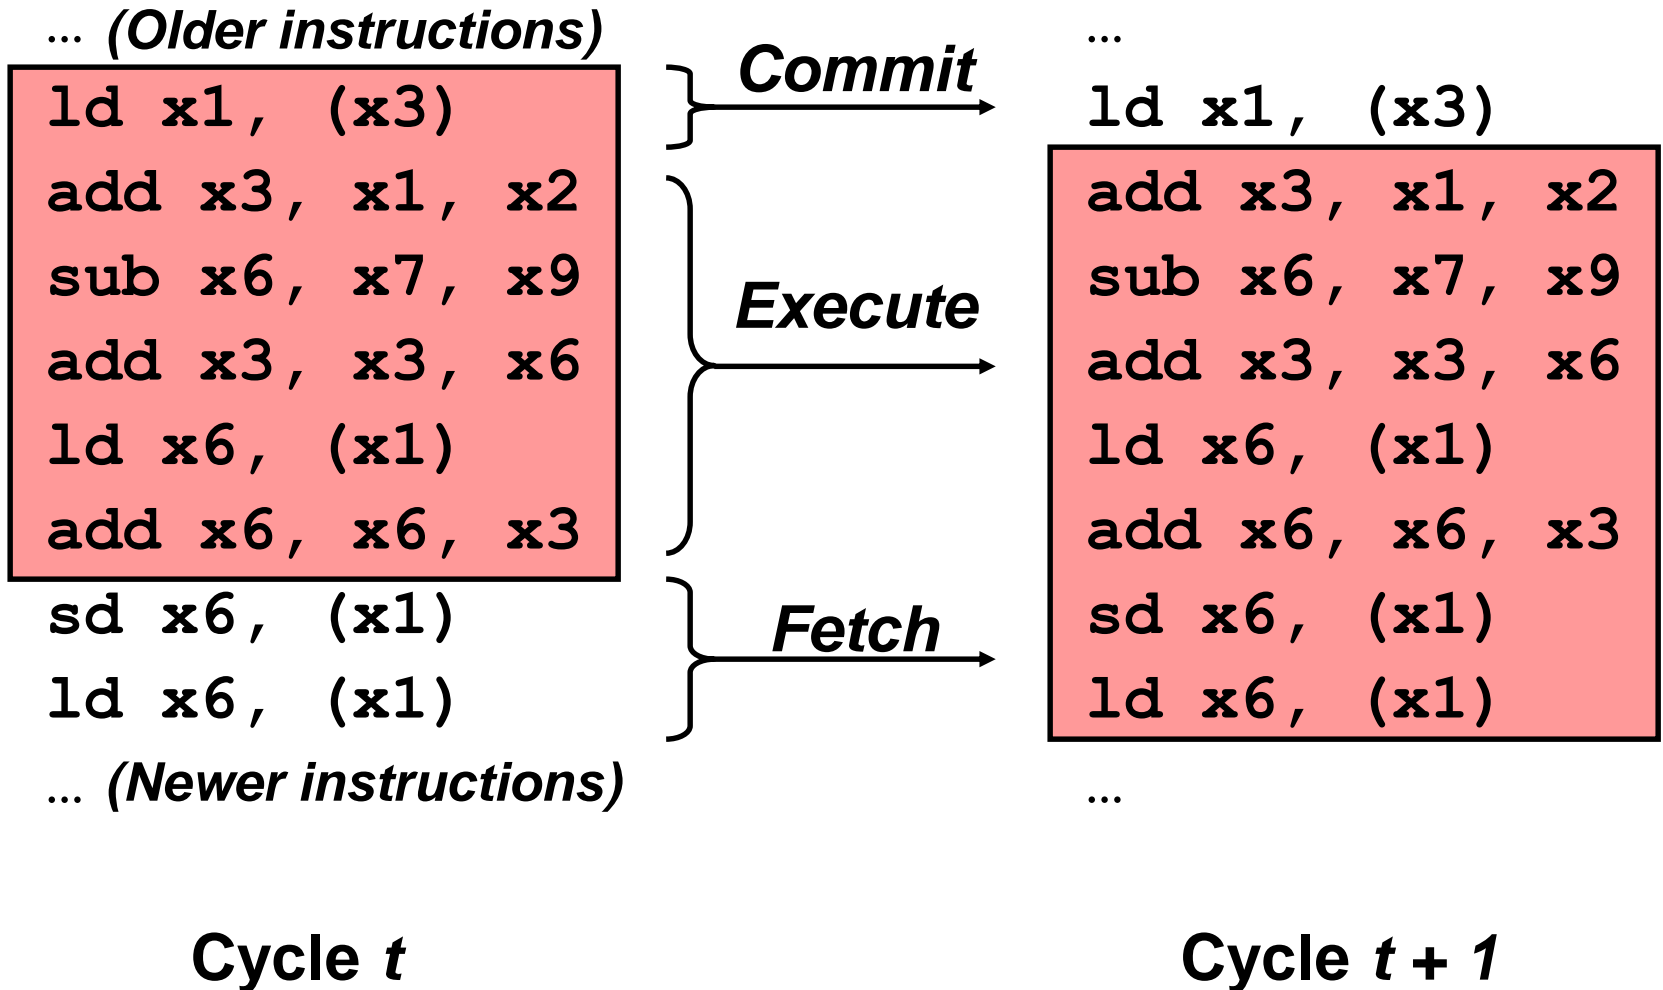
Ptr₂
next to commit

Ptr₁
next available

Done?	Rd	LPRd	PC	Except?

ROB is usually several times larger than instruction window – why?

Reorder Buffer Holds Active Instructions (Decoded but not Committed)



Issue Timing

i1	Add R1,R1,#1	Issue ₁	Execute ₁		
i2	Sub R1,R1,#1			Issue ₂	Execute ₂

How can we issue earlier?


Using knowledge of execution latency (bypass)

i1	Add R1,R1,#1	Issue ₁	Execute ₁		
i2	Sub R1,R1,#1		Issue ₂	Execute ₂	

What makes this schedule fail?

If execution latency wasn't as expected

Issue Queue with latency prediction



Issue Queue (Reorder buffer)

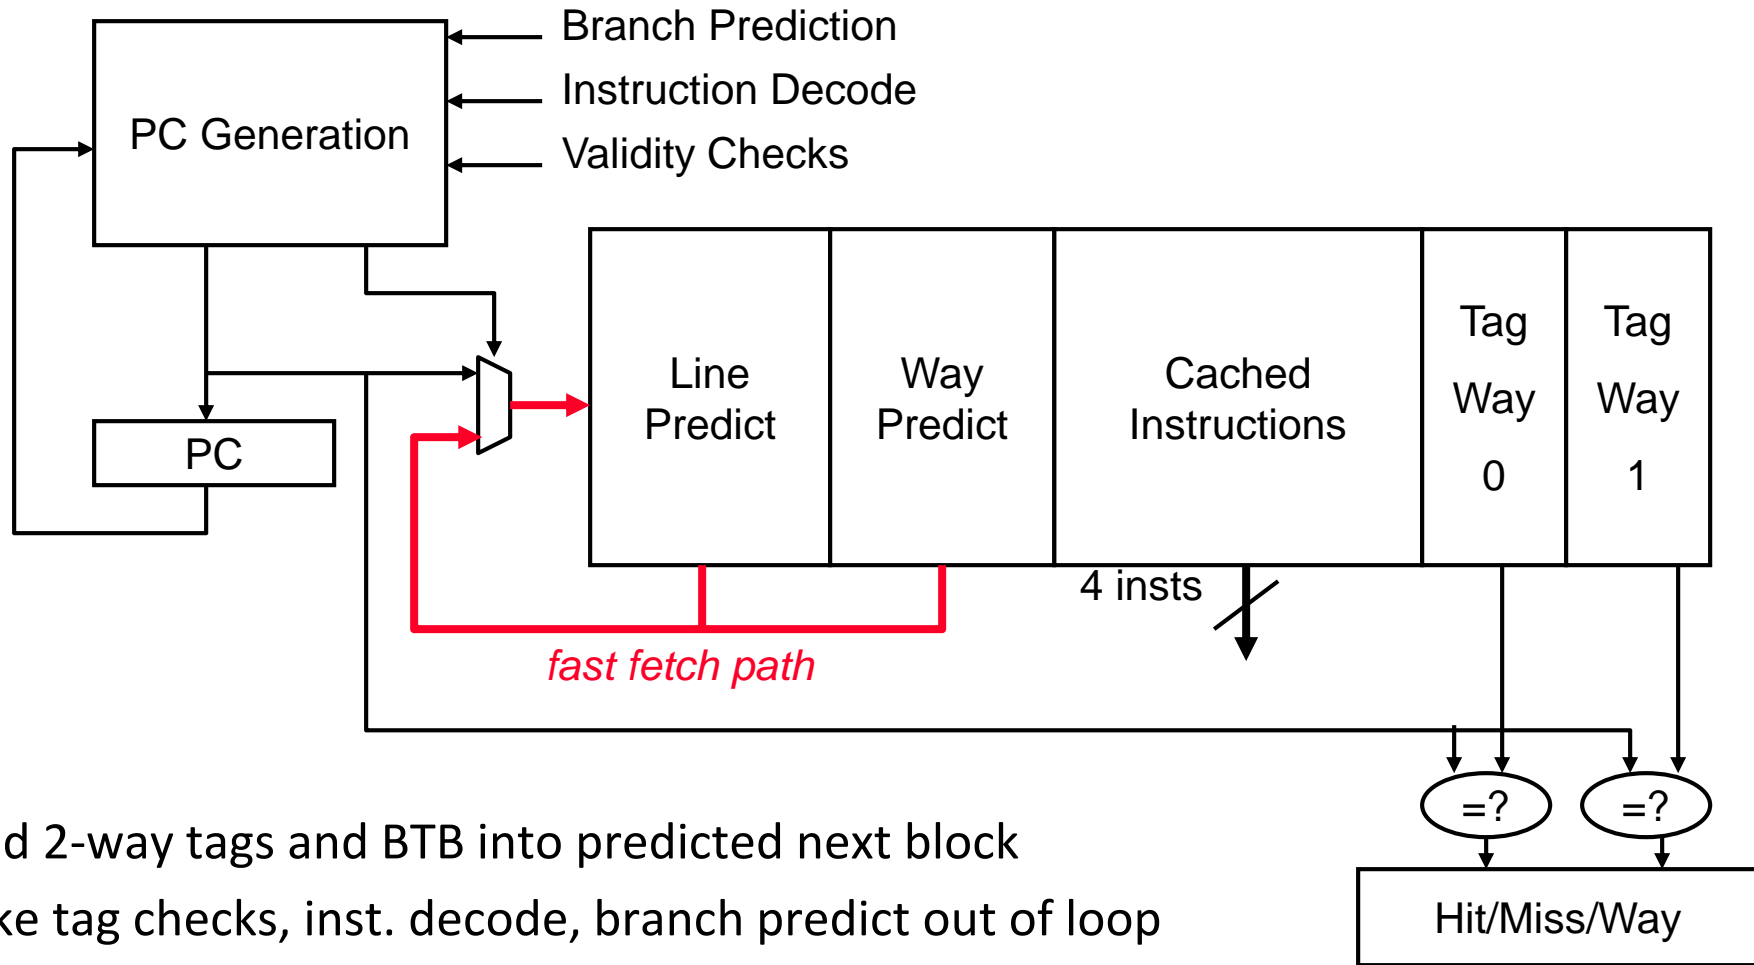
- Fixed latency: latency included in queue entry (‘bypassed’)
- Predicted latency: latency included in queue entry (speculated)
- Variable latency: wait for completion signal (stall)

Improving Instruction Fetch

Performance of speculative out-of-order machines often limited by instruction fetch bandwidth

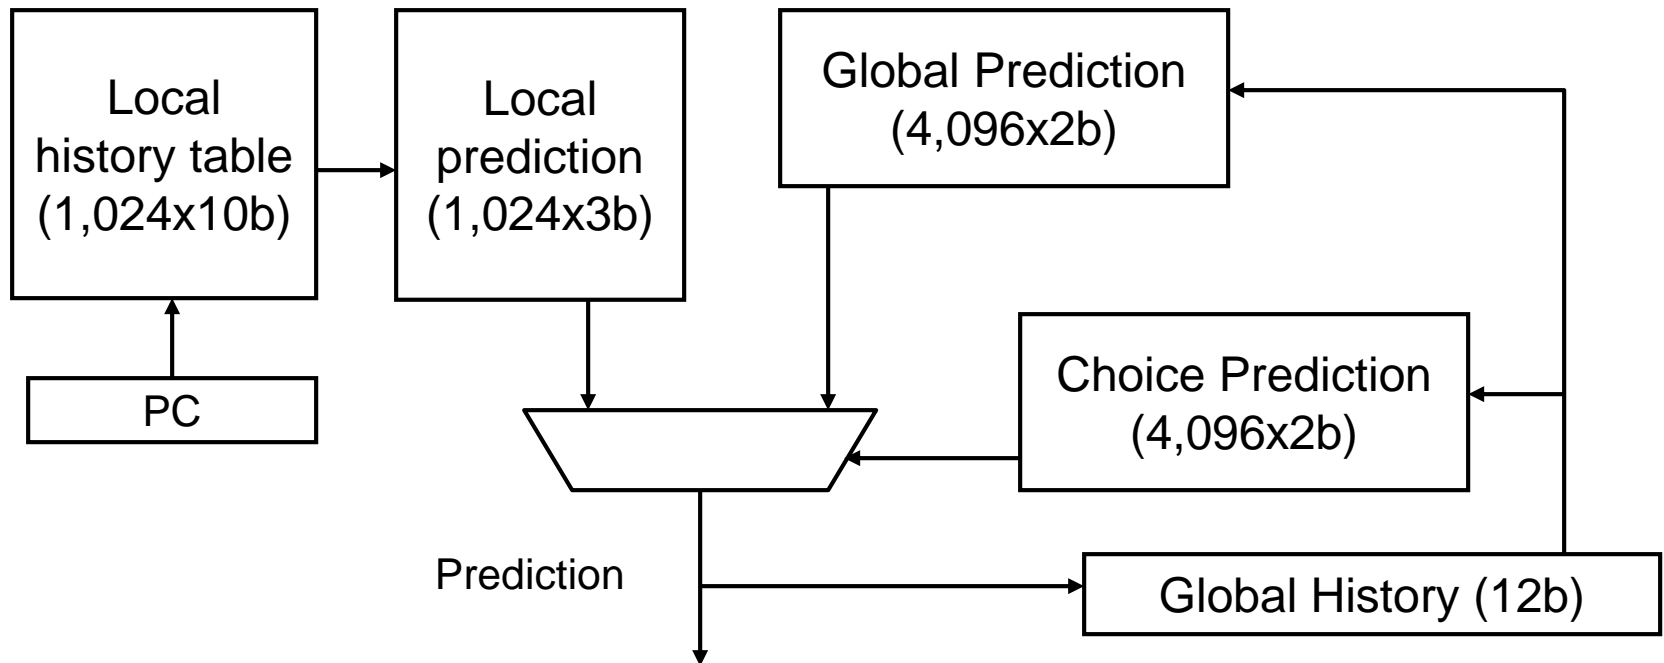
- speculative execution can fetch 2-3x more instructions than are committed
- mispredict penalties dominated by time to refill instruction window
- *taken branches* are particularly troublesome

Increasing Taken Branch Bandwidth (Alpha 21264 I-Cache)



- Fold 2-way tags and BTB into predicted next block
- Take tag checks, inst. decode, branch predict out of loop
- Raw RAM speed on critical loop (1 cycle at ~1 GHz)
- 2-bit hysteresis counter per block prevents overtraining

Tournament Branch Predictor (Alpha 21264)

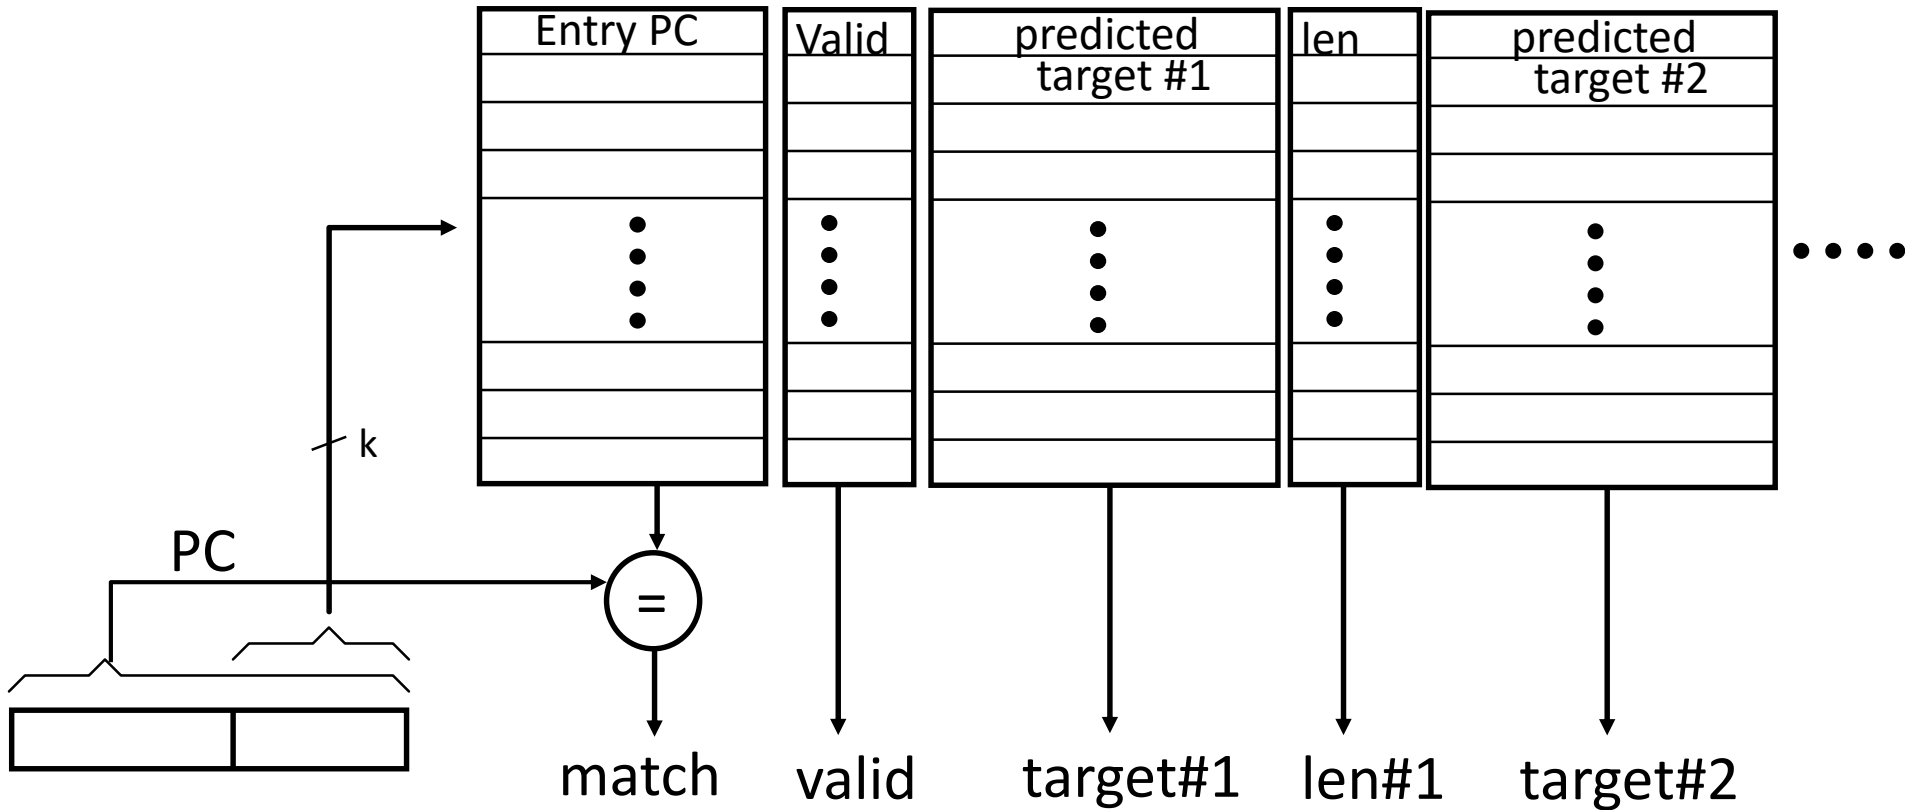


- Choice predictor learns whether best to use local or global branch history in predicting next branch (best in each case)
- Global history is speculatively updated but restored on mispredict
- Claim 90-100% success on range of applications

Taken Branch Limit

- Integer codes have a taken branch every 6-9 instructions
- To avoid fetch bottleneck, must execute multiple taken branches per cycle when increasing performance
- This implies:
 - predicting multiple branches per cycle
 - fetching multiple non-contiguous blocks per cycle

Branch Address Cache (Yeh, Marr, Patt)



Extend BTB to return multiple branch predictions per cycle

Fetching Multiple Basic Blocks


Requires either

- multiported cache: expensive
- interleaving: bank conflicts will occur

Merging multiple blocks to feed to decoders adds latency
increasing mispredict penalty and reducing branch
throughput

Trace Cache

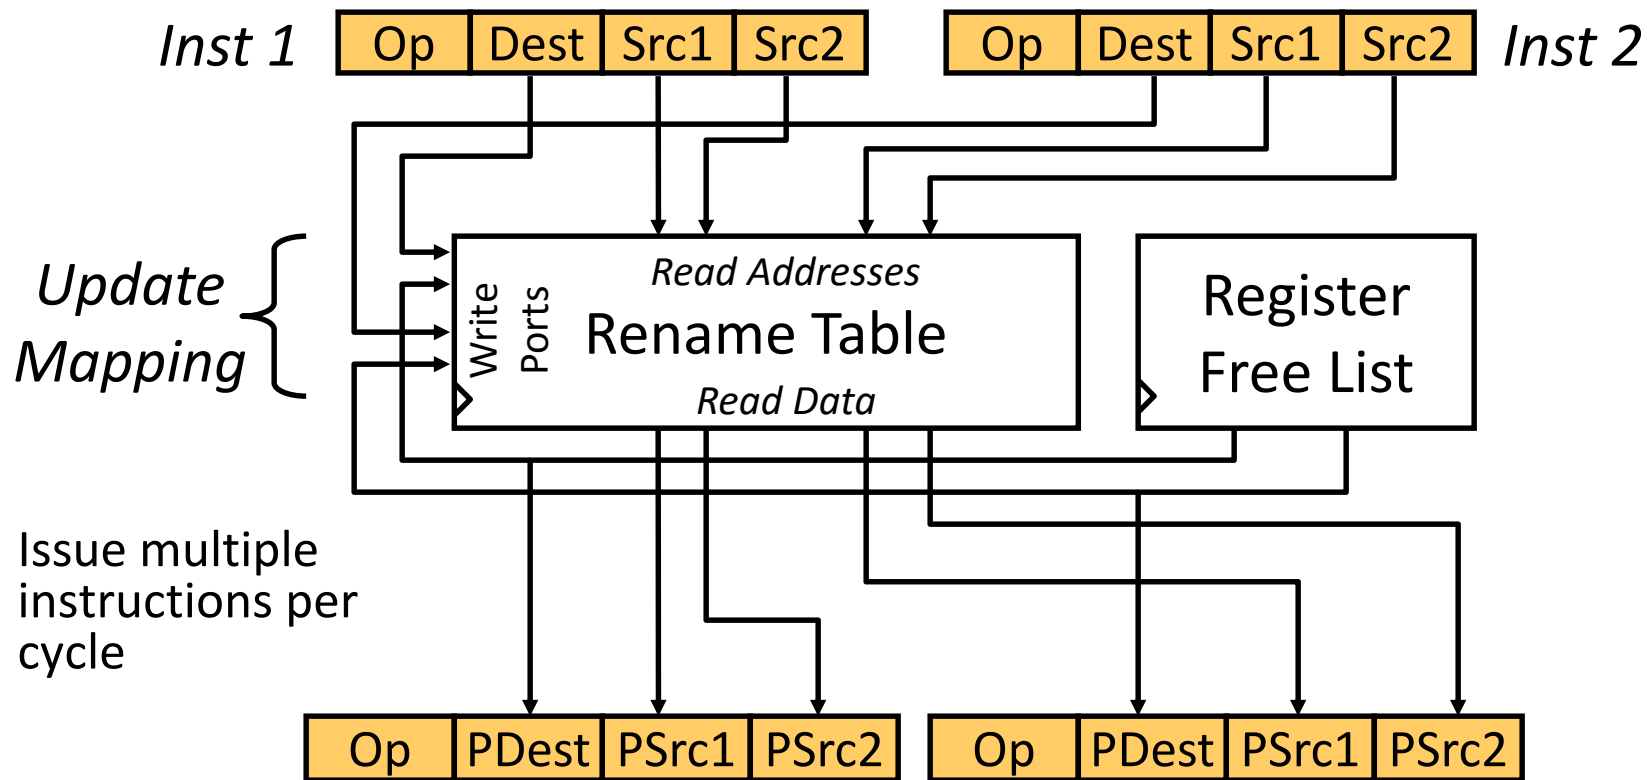
Key Idea: Pack multiple non-contiguous basic blocks into one contiguous trace cache line



- Single fetch brings in multiple basic blocks
- Trace cache indexed by start address *and* next n branch predictions
- Used in Intel Pentium-4 processor to hold decoded uops

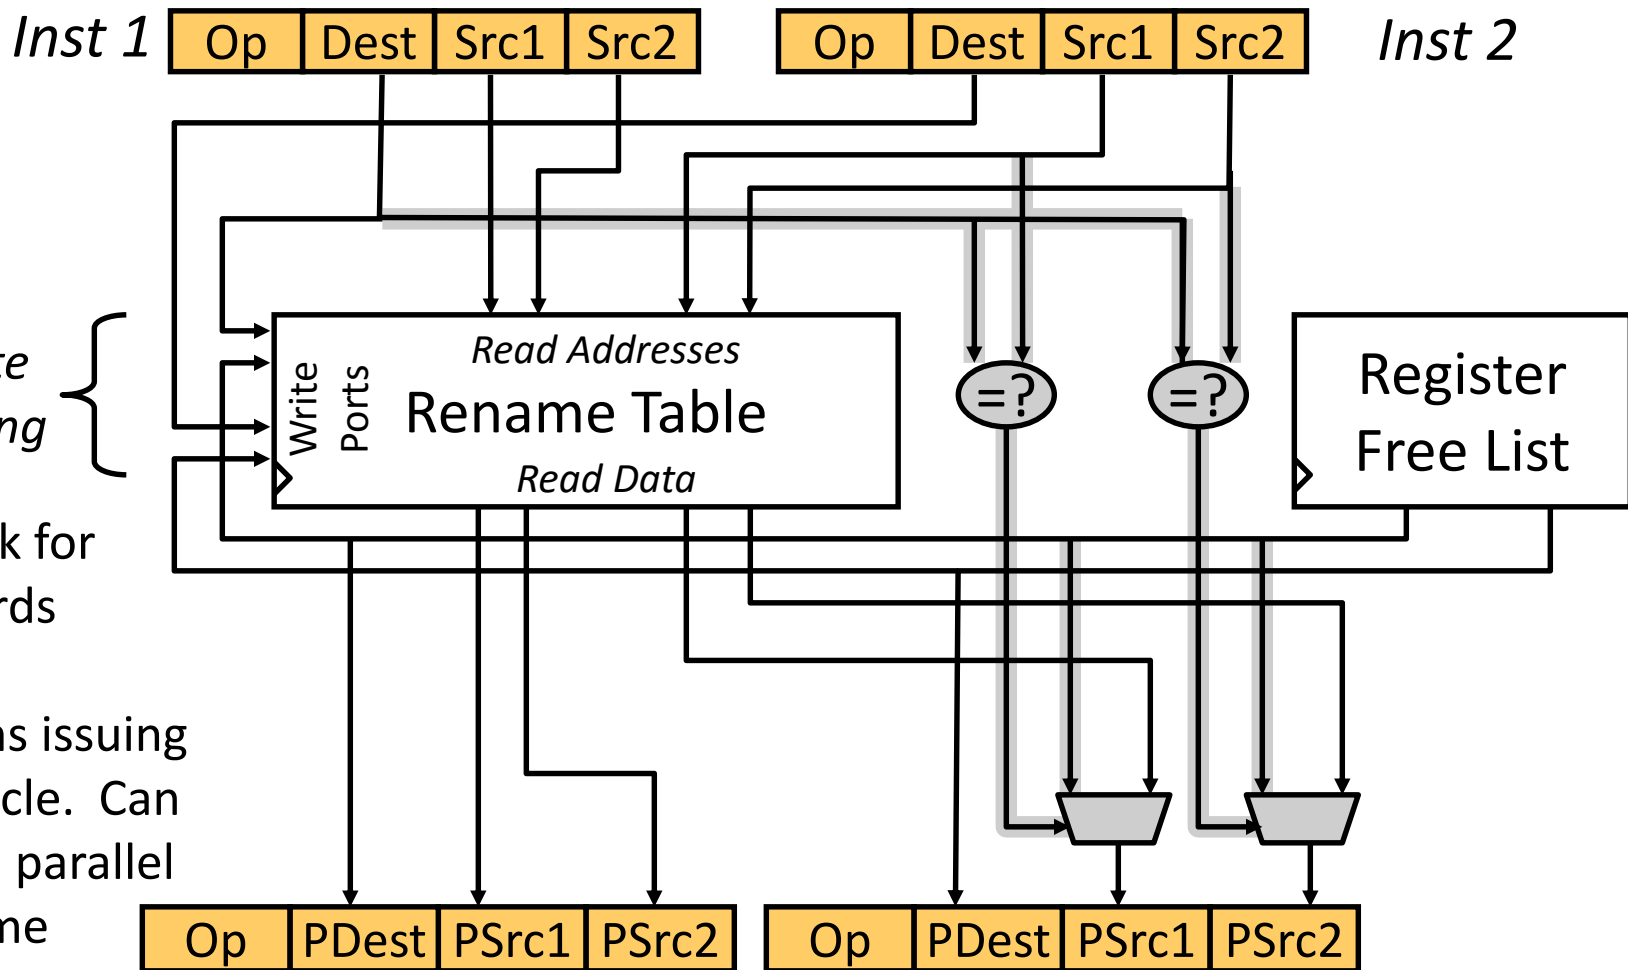
Superscalar Register Renaming

- During decode, instructions allocated new physical destination register
- Source operands renamed to physical register with newest value
- Execution unit only sees physical register numbers



Does this work?

Superscalar Register Renaming



Must check for RAW hazards between instructions issuing in same cycle. Can be done in parallel with rename lookup.


MIPS R10K renames 4 serially-RAW-dependent insts/cycle

Speculative Loads / Stores

Just like register updates, stores should not modify the memory until after the instruction is committed

- A speculative store buffer is a structure introduced to hold speculative store data.

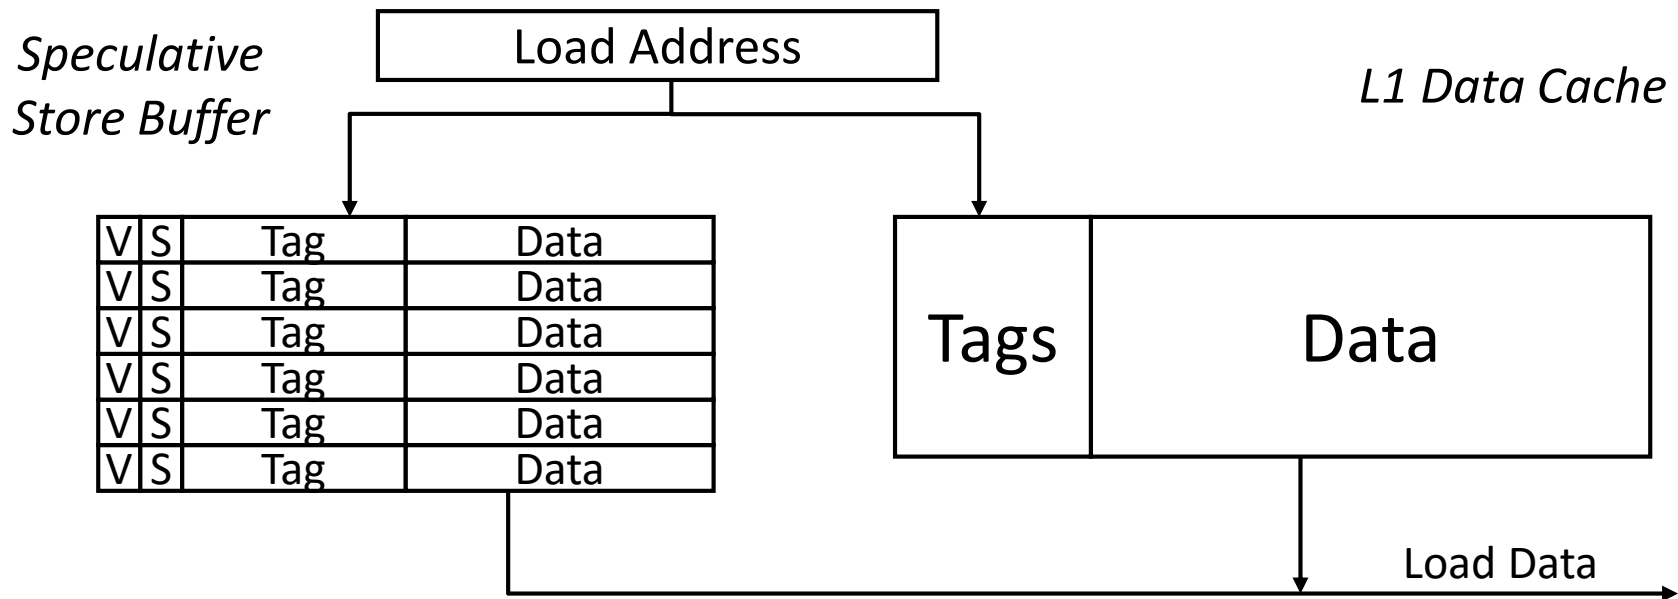
Speculative Store Buffer



Just like register updates, stores should not modify the memory until after the instruction is committed. A speculative store buffer is a structure introduced to hold speculative store data.

- During decode, store buffer slot allocated in program order
- Stores split into “store address” and “store data” micro-operations
- “Store address” execute writes tag
- “Store data” execute writes data
- Store commits when oldest instruction and both address and data available:
 - clear speculative bit and eventually move data to cache
- On store abort:
 - clear valid bit

Load bypass from speculative store buffer



- If data in both store buffer and cache, which should we use?
Speculative store buffer
- If same address in store buffer twice, which should we use?
Youngest store older than load

Memory Dependencies

```
sd x1, (x2)
```

```
ld x3, (x4)
```

When can we execute the load?

In-Order Memory Queue

- Execute all loads and stores in program order
- => Load and store cannot leave ROB for execution until all previous loads and stores have completed execution
- Can still execute loads and stores speculatively, and out-of-order with respect to other instructions
 - Need a structure to handle memory ordering...

Conservative O-o-O Load Execution

```
sd x1, (x2)
ld x3, (x4)
```

- Can execute load before store, if addresses known and $x4 \neq x2$
- Each load address compared with addresses of all previous uncommitted stores
 - *can use partial conservative check i.e., bottom 12 bits of address, to save hardware*
- Don't execute load if any previous store address not known

(MIPS R10K, 16-entry address queue)

Address Speculation

```
sd x1, (x2)
```

```
ld x3, (x4)
```

- Guess that $x4 \neq x2$
- Execute load before store address known
- Need to hold all completed but uncommitted load/store addresses in program order
- If subsequently find $x4 == x2$, squash load and *all* following instructions

=> Large penalty for inaccurate address speculation

Memory Dependence Prediction

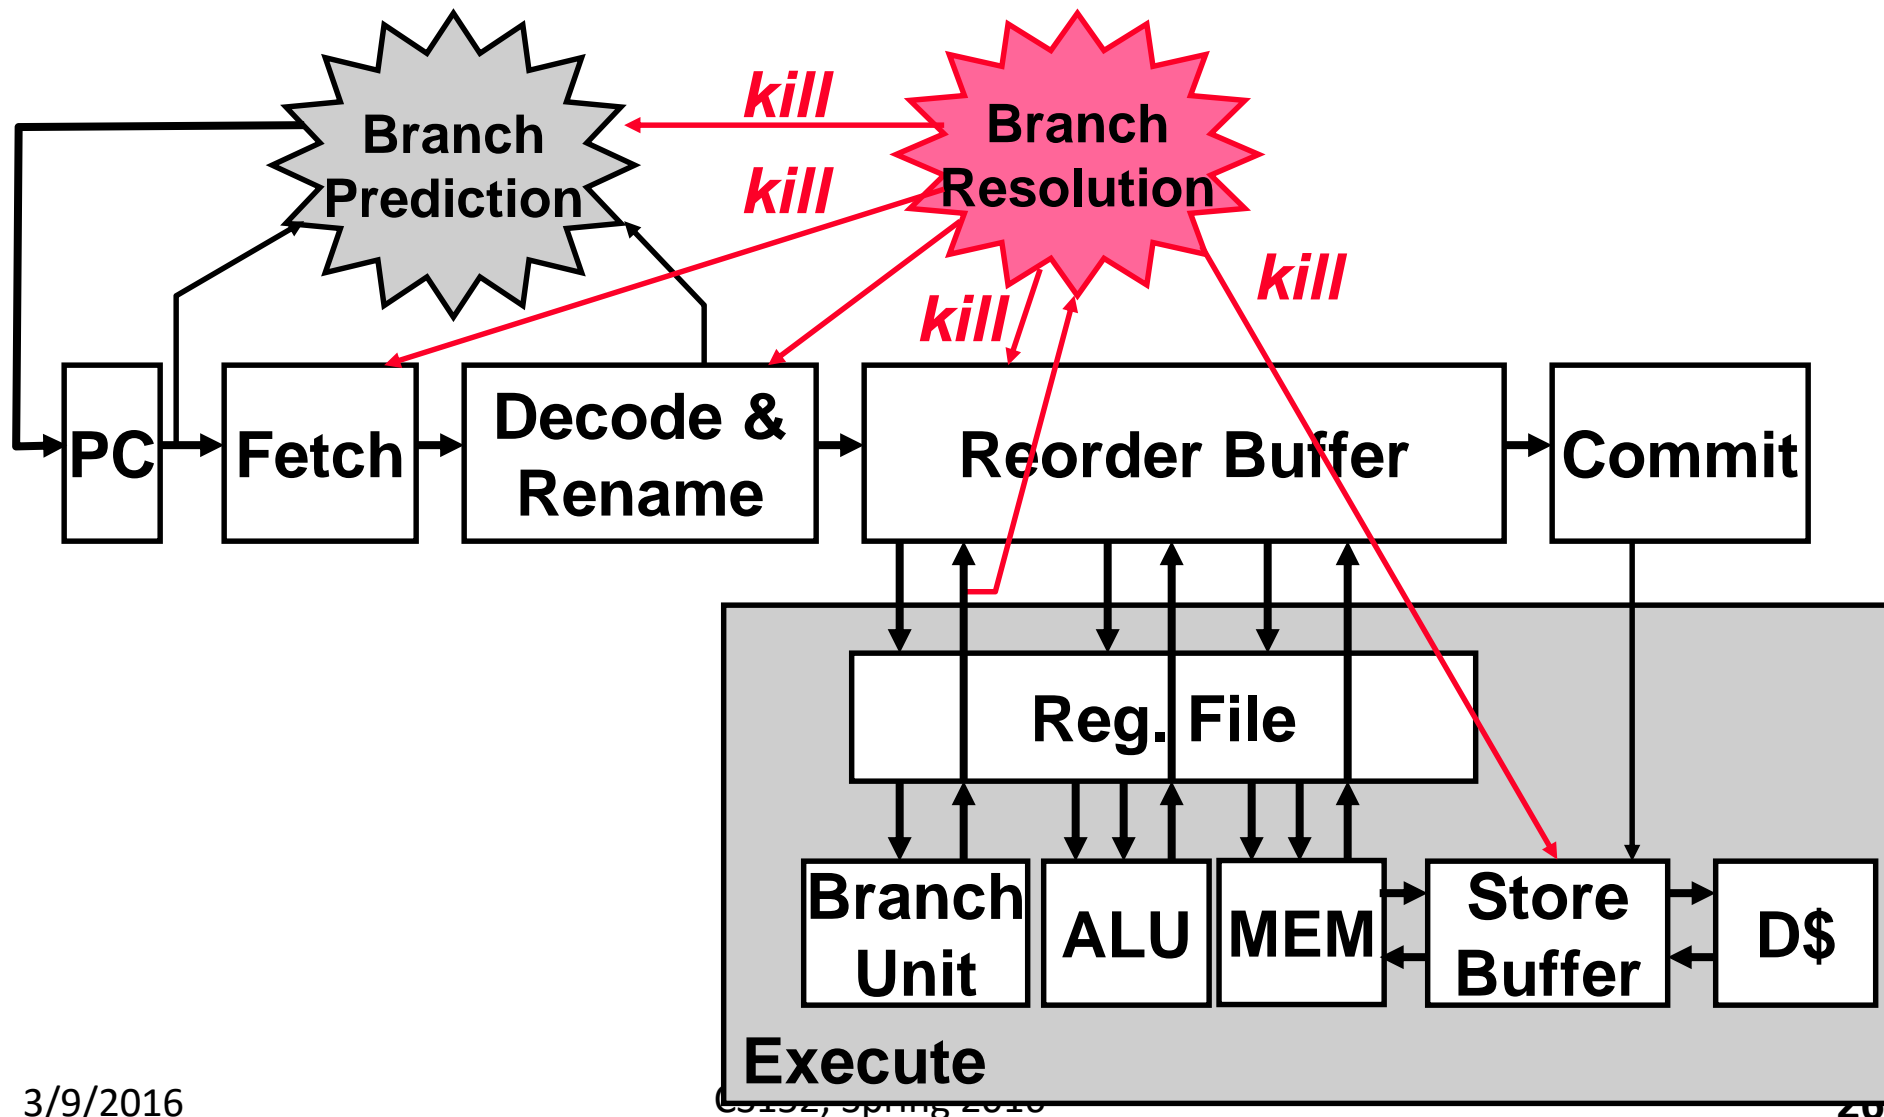
(Alpha 21264)

```
sd x1, (x2)
```

```
ld x3, (x4)
```

- Guess that $x4 \neq x2$ and execute load before store
- If later find $x4 == x2$, squash load and all following instructions, but mark load instruction as *store-wait*
- Subsequent executions of the same load instruction will wait for all previous stores to complete
- Periodically clear *store-wait* bits

Datapath: Branch Prediction and Speculative Execution



Instruction Flow in Unified Physical Register File Pipeline

■ Fetch

- Get instruction bits from current guess at PC, place in fetch buffer
- Update PC using sequential address or branch predictor (BTB)

■ Decode/Rename

- Take instruction from fetch buffer
- Allocate resources to execute instruction:
 - Destination physical register, if instruction writes a register
 - Entry in reorder buffer to provide in-order commit
 - Entry in issue window to wait for execution
 - Entry in memory buffer, if load or store
- Decode will stall if resources not available
- Rename source and destination registers
- Check source registers for readiness
- Insert instruction into issue window+reorder buffer+memory buffer

Memory Instructions

- Split store instruction into two pieces during decode:
 - Address calculation, store-address
 - Data movement, store-data
- Allocate space in program order in memory buffers during decode
- Store instructions:
 - Store-address calculates address and places in store buffer
 - Store-data copies store value into store buffer
 - Store-address and store-data execute independently out of issue window
 - Stores only commit to data cache at commit point
- Load instructions:
 - Load address calculation executes from window
 - Load with completed effective address searches memory buffer
 - Load instruction may have to wait in memory buffer for earlier store ops to resolve

Issue Stage

- Writebacks from completion phase “wakeup” some instructions by causing their source operands to become ready in issue window
 - In more speculative machines, might wake up waiting loads in memory buffer
- Need to “select” some instructions for issue
 - Arbiter picks a subset of ready instructions for execution
 - Example policies: random, lower-first, oldest-first, critical-first
- Instructions read out from issue window and sent to execution

Execute Stage

- Read operands from physical register file and/or bypass network from other functional units
- Execute on functional unit
- Write result value to physical register file (or store buffer if store)
- Produce exception status, write to reorder buffer
- Free slot in instruction window

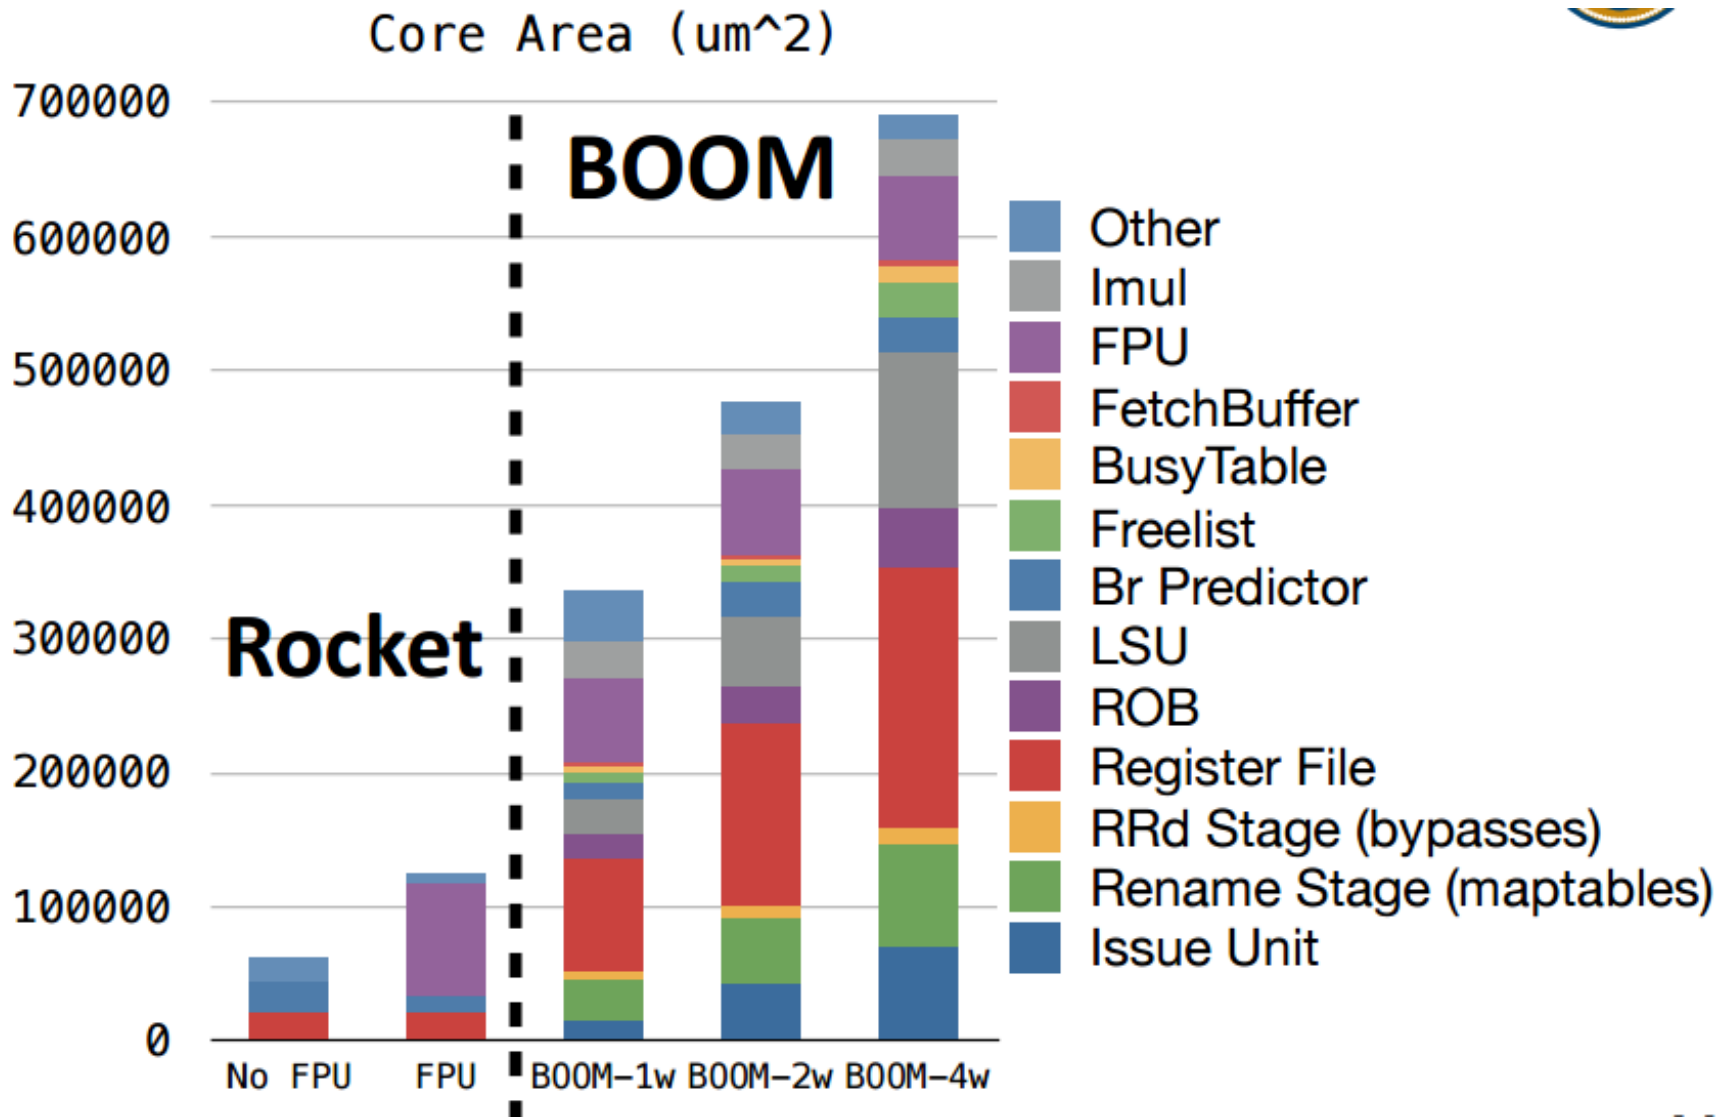
Commit Stage

- Read completed instructions in-order from reorder buffer
 - (may need to wait for next oldest instruction to complete)
- If exception raised
 - flush pipeline, jump to exception handler
- Otherwise, release resources:
 - Free physical register used by last writer to same architectural register
 - Free reorder buffer slot
 - Free memory reorder buffer slot


Question of the Day

- How many in-order cores do you think take up the same area as an out-of-order core?

How Many In-Order In The Same Area?



Performance?



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