

18-447

Computer Architecture  
Lecture 11: Precise Exceptions,  
State Maintenance, State Recovery

Prof. Onur Mutlu

Carnegie Mellon University

Spring 2014, 2/10/2014

# Announcements

---

- Homework 2 due Wednesday (Feb 12)
- Lab 3 available online (due Feb 21)

# Readings for Next Few Lectures (I)

---

- P&H Chapter 4.9-4.11
- Smith and Sohi, “The Microarchitecture of Superscalar Processors,” Proceedings of the IEEE, 1995
  - More advanced pipelining
  - Interrupt and exception handling
  - Out-of-order and superscalar execution concepts
- McFarling, “Combining Branch Predictors,” DEC WRL Technical Report, 1993.
- Kessler, “The Alpha 21264 Microprocessor,” IEEE Micro 1999.

# Readings for Next Few Lectures (II)

---

- Smith and Plezskun, “[Implementing Precise Interrupts in Pipelined Processors](#),” IEEE Trans on Computers 1988 (earlier version in ISCA 1985).

# Readings Specifically for Today

---

- Smith and Plezskun, “[Implementing Precise Interrupts in Pipelined Processors](#),” IEEE Trans on Computers 1988 (earlier version in ISCA 1985).
- Smith and Sohi, “[The Microarchitecture of Superscalar Processors](#),” Proceedings of the IEEE, 1995
  - More advanced pipelining
  - Interrupt and exception handling
  - Out-of-order and superscalar execution concepts

# Review: How to Handle Control Dependences

---

- Critical to keep the pipeline full with correct sequence of dynamic instructions.
- Potential solutions if the instruction is a control-flow instruction:
  - **Stall** the pipeline until we know the next fetch address
  - Guess the next fetch address (**branch prediction**)
  - Employ delayed branching (**branch delay slot**)
  - Do something else (**fine-grained multithreading**)
  - Eliminate control-flow instructions (**predicated execution**)
  - Fetch from both possible paths (if you know the addresses of both possible paths) (**multipath execution**)

# Remember: Branch Types

Type	Direction at fetch time	Number of possible next fetch addresses?	When is next fetch address resolved?
Conditional	Unknown	2	Execution (register dependent)
Unconditional	Always taken	1	Decode (PC + offset)
Call	Always taken	1	Decode (PC + offset)
Return	Always taken	Many	Execution (register dependent)
Indirect	Always taken	Many	Execution (register dependent)

Different branch types can be handled differently

# Call and Return Prediction

---

- **Direct calls are easy to predict**

- Always taken, single target
- Call marked in BTB, target predicted by BTB

- **Returns are indirect branches**

- A function can be called from many points in code
- **A return instruction can have many target addresses**
  - Next instruction after each call point for the same function
- **Observation: Usually a return matches a call**
- **Idea: Use a stack to predict return addresses (Return Address Stack)**
  - A fetched call: pushes the return (next instruction) address on the stack
  - A fetched return: pops the stack and uses the address as its predicted target
  - Accurate most of the time: 8-entry stack → > 95% accuracy

Call X

...

Call X

...

Call X

...


Return

Return


Return

# Indirect Branch Prediction (I)

- Register-indirect branches have multiple targets



Conditional (Direct) Branch



Indirect Jump

- Used to implement
  - ❑ Switch-case statements
  - ❑ Virtual function calls
  - ❑ Jump tables (of function pointers)
  - ❑ Interface calls

# Indirect Branch Prediction (II)

---

- No direction prediction needed
- Idea 1: Predict the last resolved target as the next fetch address
  - + Simple: Use the BTB to store the target address
  - Inaccurate: 50% accuracy (empirical). Many indirect branches switch between different targets
- Idea 2: Use history based target prediction
  - E.g., Index the BTB with GHR XORed with Indirect Branch PC
  - Chang et al., “Target Prediction for Indirect Jumps,” ISCA 1997.
  - + More accurate
  - An indirect branch maps to (too) many entries in BTB
    - Conflict misses with other branches (direct or indirect)
    - Inefficient use of space if branch has few target addresses

# More Ideas on Indirect Branches?

---

- Virtual Program Counter prediction
  - Idea: Use conditional branch prediction structures *iteratively* to make an indirect branch prediction
  - i.e., devirtualize the indirect branch in hardware
- Curious?
  - Kim et al., “VPC Prediction: Reducing the Cost of Indirect Branches via Hardware-Based Dynamic Devirtualization,” ISCA 2007.

# Issues in Branch Prediction (I)


---

- Need to identify a branch before it is fetched
- How do we do this?
  - BTB hit → indicates that the fetched instruction is a branch
  - BTB entry contains the “type” of the branch
- What if no BTB?
  - Bubble in the pipeline until target address is computed
  - E.g., IBM POWER4

# Issues in Branch Prediction (II)

---

- **Latency:** Prediction is latency critical
  - ❑ Need to generate next fetch address for the next cycle
  - ❑ Bigger, more complex predictors are more accurate but slower




# Complications in Superscalar Processors

---

- “Superscalar” processors
  - attempt to execute more than 1 instruction-per-cycle
  - must fetch multiple instructions per cycle
- Consider a 2-way superscalar fetch scenario
  - (case 1) Both insts are not taken control flow inst
    - $nPC = PC + 8$
  - (case 2) One of the insts is a taken control flow inst
    - $nPC = \text{predicted target addr}$
    - \*NOTE\* both instructions could be control-flow; prediction based on the first one predicted taken
    - If the 1<sup>st</sup> instruction is the predicted taken branch
      - nullify 2<sup>nd</sup> instruction fetched

# Multiple Instruction Fetch: Concepts




← Fetch 1 inst/cycle

- Downside:

Flynn's bottleneck

If you fetch 1 inst/cycle  
you cannot finish  $> 1$  inst/cycle



← Fetch 4 inst/cycle

Two major approaches

1) VLIW

Compiler decides what insts.  
can be executed in parallel  
→ Simple hardware

2) Superscalar

Hardware detects dependencies  
between instructions that  
are fetched in the same  
cycle.

# Review of Last Few Lectures

---

- Control dependence handling in pipelined machines
  - Delayed branching
  - Fine-grained multithreading
  - Branch prediction
    - Compile time (static)
      - Always NT, Always T, Backward T Forward NT, Profile based
    - Run time (dynamic)
      - Last time predictor
      - Hysteresis: 2BC predictor
      - Global branch correlation → Two-level global predictor
      - Local branch correlation → Two-level local predictor
  - Predicated execution
  - Multipath execution

# Pipelining and Precise Exceptions: Preserving Sequential Semantics


# Multi-Cycle Execution

---

- Not all instructions take the same amount of time for “execution”
- Idea: Have multiple different functional units that take different number of cycles
  - Can be pipelined or not pipelined
  - Can let independent instructions to start execution on a different functional unit before a previous long-latency instruction finishes execution

# Issues in Pipelining: Multi-Cycle Execute

- Instructions can take different number of cycles in EXECUTE stage
  - Integer ADD versus FP MULtiply



- What is wrong with this picture?
  - What if FMUL incurs an exception?
  - Sequential semantics of the ISA NOT preserved!

# Exceptions vs. Interrupts

---

## ■ Cause

- ❑ Exceptions: internal to the running thread
- ❑ Interrupts: external to the running thread

## ■ When to Handle

- ❑ Exceptions: when detected (and known to be non-speculative)
- ❑ Interrupts: when convenient
  - Except for very high priority ones
    - ❑ Power failure
    - ❑ Machine check

## ■ Priority: process (exception), depends (interrupt)

## ■ Handling Context: process (exception), system (interrupt)

---

# Precise Exceptions/Interrupts

---

- The architectural state should be consistent when the exception/interrupt is ready to be handled

1. All previous instructions should be completely retired.

2. No later instruction should be retired.

Retire = commit = finish execution and update arch. state

# Why Do We Want Precise Exceptions?


---

- Semantics of the von Neumann model ISA specifies it
  - Remember von Neumann vs. dataflow
- Aids software debugging
- Enables (easy) recovery from exceptions, e.g. page faults
- Enables (easily) restartable processes
- Enables traps into software (e.g., software implemented opcodes)

# Ensuring Precise Exceptions in Pipelining

- Idea: Make each operation take the same amount of time

FMUL R3 ← R1, R2  
ADD R4 ← R1, R2



- Downside
  - ❑ What about memory operations?
  - ❑ Each functional unit takes 500 cycles?

# Solutions


---

- Reorder buffer
- History buffer
- Future register file
- Checkpointing
- Recommended Reading
  - Smith and Plezskun, “[Implementing Precise Interrupts in Pipelined Processors](#),” IEEE Trans on Computers 1988 and ISCA 1985.
  - Hwu and Patt, “[Checkpoint Repair for Out-of-order Execution Machines](#),” ISCA 1987.

# Solution I: Reorder Buffer (ROB)

---

- Idea: Complete instructions out-of-order, but reorder them before making results visible to architectural state
- When instruction is decoded it reserves an entry in the ROB
- When instruction completes, it writes result into ROB entry
- When instruction oldest in ROB and it has completed without exceptions, its result moved to reg. file or memory



# What's in a ROB Entry?


---

V	DestRegID	DestRegVal	StoreAddr	StoreData	PC	Valid bits for reg/data + control bits	Exc?
---	-----------	------------	-----------	-----------	----	---	------

- Need valid bits to keep track of readiness of the result(s)

# Reorder Buffer: Independent Operations


- Results first written to ROB, then to register file at commit time



- What if a later operation needs a value in the reorder buffer?
  - Read reorder buffer in parallel with the register file. **How?**

# Reorder Buffer: How to Access?

- A register value can be in the register file, reorder buffer, (or bypass/forwarding paths)



# Simplifying Reorder Buffer Access

---

- Idea: Use indirection
- Access register file first
  - If register not valid, register file stores the ID of the reorder buffer entry that contains (or will contain) the value of the register
  - Mapping of the register to a ROB entry
- Access reorder buffer next
- What is in a reorder buffer entry?

V	DestRegID	DestRegVal	StoreAddr	StoreData	PC/IP	Control/val id bits	Exc?
---	-----------	------------	-----------	-----------	-------	------------------------	------

- Can it be simplified further?


# Aside: Register Renaming with a Reorder Buffer

---

- Output and anti dependencies are not true dependencies
  - WHY? The same register refers to values that have nothing to do with each other
  - **They exist due to lack of register ID' s (i.e. names) in the ISA**
- The register ID is **renamed** to the reorder buffer entry that will hold the register' s value
  - Register ID → ROB entry ID
  - Architectural register ID → Physical register ID
  - After renaming, ROB entry ID used to refer to the register
- This eliminates anti- and output- dependencies
  - Gives the illusion that there are a large number of registers

# In-Order Pipeline with Reorder Buffer

- **Decode (D)**: Access regfile/ROB, allocate entry in ROB, check if instruction can execute, if so **dispatch** instruction
- **Execute (E)**: Instructions can complete out-of-order
- **Completion (R)**: Write result to **reorder buffer**
- **Retirement/Commit (W)**: Check for exceptions; if none, write result to architectural register file or memory; else, flush pipeline and start from exception handler
- **In-order dispatch/execution, out-of-order completion, in-order retirement**



# Reorder Buffer Tradeoffs

---


- Advantages
  - Conceptually simple for supporting precise exceptions
  - Can eliminate false dependencies
- Disadvantages
  - Reorder buffer needs to be accessed to get the results that are yet to be written to the register file
    - CAM or indirection → increased latency and complexity
- Other solutions aim to eliminate the disadvantages
  - History buffer
  - Future file
  - Checkpointing

# Solution II: History Buffer (HB)

---

- Idea: Update the register file when instruction completes, but UNDO UPDATES when an exception occurs
- When instruction is decoded, it reserves an HB entry
- When the instruction completes, it stores the old value of its destination in the HB
- When instruction is oldest and no exceptions/interrupts, the HB entry discarded
- When instruction is oldest and an exception needs to be handled, old values in the HB are written back into the architectural state from tail to head

# History Buffer



Used only on exceptions

- Advantage:

- Register file contains up-to-date values. History buffer access not on critical path

- Disadvantage:

- Need to read the old value of the destination register
- Need to unwind the history buffer upon an exception → increased exception/interrupt handling latency


# Solution III: Future File (FF) + ROB

---

- Idea: **Keep two register files (speculative and architectural)**
  - Arch reg file: Updated in program order for precise exceptions
    - Use a reorder buffer to ensure in-order updates
  - Future reg file: Updated as soon as an instruction completes (if the instruction is the youngest one to write to a register)
- **Future file is used for fast access to latest register values (speculative state)**
  - Frontend register file
- **Architectural file is used for state recovery on exceptions (architectural state)**
  - Backend register file

# Future File

---



Used only on exceptions

## ■ Advantage

- ❑ No need to read the values from the ROB (no CAM or indirection)

## ■ Disadvantage

- ❑ Multiple register files
- ❑ Need to copy arch. reg. file to future file on an exception



# Checking for and Handling Exceptions in Pipelining

---

- When the **oldest instruction ready-to-be-retired is detected to have caused an exception**, the control logic
  - Recovers architectural state (register file, IP, and memory)
  - Flushes all younger instructions in the pipeline
  - Saves IP and registers (as specified by the ISA)
  - Redirects the fetch engine to the exception handling routine
    - Vectored exceptions

# Pipelining Issues: Branch Mispredictions

---

- A branch misprediction resembles an “exception”
  - Except it is not visible to software
- What about branch misprediction recovery?
  - Similar to exception handling except can be initiated before the branch is the oldest instruction
  - All three state recovery methods can be used
- Difference between exceptions and branch mispredictions?
  - Branch mispredictions are much more common
    - need fast state recovery to minimize performance impact of mispredictions

# How Fast Is State Recovery?

---

- Latency of state recovery affects
  - Exception service latency
  - Interrupt service latency
  - Latency to supply the correct data to instructions fetched after a branch misprediction
  
- Which ones above need to be fast?
  
- How do the three state maintenance methods fare in terms of recovery latency?
  - Reorder buffer
  - History buffer
  - Future file

# Branch State Recovery Actions and Latency

---

- Reorder Buffer
  - Wait until branch is the oldest instruction in the machine
  - Flush entire pipeline
- History buffer
  - Undo all instructions after the branch by rewinding from the tail of the history buffer until the branch & restoring old values one by one into the register file
  - Flush instructions in pipeline younger than the branch
- Future file
  - Wait until branch is the oldest instruction in the machine
  - Copy arch. reg. file to future file
  - Flush entire pipeline

# Can We Do Better?

---

- Goal: Restore the frontend state (future file) such that the correct next instruction after the branch can execute right away after the branch misprediction is resolved
- Idea: Checkpoint the frontend register state at the time a branch is fetched and keep the checkpointed state updated with results of instructions older than the branch
- Hwu and Patt, “Checkpoint Repair for Out-of-order Execution Machines,” ISCA 1987.

# Checkpointing

---

- **When a branch is decoded**
  - Make a copy of the future file and associate it with the branch
- **When an instruction produces a register value**
  - All future file checkpoints that are younger than the instruction are updated with the value
- **When a branch misprediction is detected**
  - Restore the checkpointed future file for the mispredicted branch when the branch misprediction is resolved
  - Flush instructions in pipeline younger than the branch
  - Deallocate checkpoints younger than the branch

# Checkpointing

---

- Advantages?

- Disadvantages?

# Registers versus Memory

---

- So far, we considered mainly registers as part of state
- What about memory?
- What are the fundamental differences between registers and memory?
  - Register dependences known statically – memory dependences determined dynamically
  - Register state is small – memory state is large
  - Register state is not visible to other threads/processors – memory state is shared between threads/processors (in a shared memory multiprocessor)

# Maintaining Speculative Memory State: Stores


---

- Handling out-of-order completion of memory operations
  - UNDOing a memory write more difficult than UNDOing a register write. **Why?**
  - **One idea:** Keep store address/data in reorder buffer
    - How does a load instruction find its data?
  - **Store/write buffer:** Similar to reorder buffer, but used only for store instructions
    - Program-order list of un-committed store operations
    - When store is decoded: Allocate a store buffer entry
    - When store address and data become available: Record in store buffer entry
    - When the store is the oldest instruction in the pipeline: Update the memory address (i.e. cache) with store data

# Out-of-Order Execution (Dynamic Instruction Scheduling)

# An In-order Pipeline

---



- Problem: A true data dependency stalls dispatch of younger instructions into functional (execution) units
- Dispatch: Act of sending an instruction to a functional unit

# Can We Do Better?

---

- What do the following two pieces of code have in common (with respect to execution in the previous design)?

```
IMUL R3 ← R1, R2
ADD  R3 ← R3, R1
ADD  R1 ← R6, R7
IMUL R5 ← R6, R8
ADD  R7 ← R3, R5
```

```
LD   R3 ← R1 (0)
ADD  R3 ← R3, R1
ADD  R1 ← R6, R7
IMUL R5 ← R6, R8
ADD  R7 ← R3, R5
```

- **Answer: First ADD stalls the whole pipeline!**
  - ADD cannot dispatch because its source registers unavailable
  - **Later independent instructions cannot get executed**
- How are the above code portions different?
  - **Answer: Load latency is variable (unknown until runtime)**
  - What does this affect? Think compiler vs. microarchitecture

# Preventing Dispatch Stalls

---

- Multiple ways of doing it
- You have already seen THREE:
  - 1.
  - 2.
  - 3.
- What are the disadvantages of the above three?
- Any other way to prevent dispatch stalls?
  - Actually, you have briefly seen the basic idea before
    - Dataflow: fetch and “fire” an instruction when its inputs are ready
  - Problem: in-order dispatch (scheduling, or execution)
  - Solution: out-of-order dispatch (scheduling, or execution)


# Out-of-order Execution (Dynamic Scheduling)

---

- Idea: Move the dependent instructions out of the way of independent ones
  - Rest areas for dependent instructions: Reservation stations
- Monitor the source “values” of each instruction in the resting area
- When all source “values” of an instruction are available, “fire” (i.e. dispatch) the instruction
  - Instructions dispatched in **dataflow (not control-flow) order**
- Benefit:
  - **Latency tolerance**: Allows independent instructions to execute and complete in the presence of a long latency operation

# In-order vs. Out-of-order Dispatch


- In order dispatch + precise exceptions:



```

IMUL R3 ← R1, R2
ADD R3 ← R3, R1
ADD R1 ← R6, R7
IMUL R5 ← R6, R8
ADD R7 ← R3, R5
    
```

- Out-of-order dispatch + precise exceptions:



- 16 vs. 12 cycles

# Enabling OoO Execution

---


1. Need to link the consumer of a value to the producer
  - ❑ Register renaming: Associate a “tag” with each data value
2. Need to buffer instructions until they are ready to execute
  - ❑ Insert instruction into reservation stations after renaming
3. Instructions need to keep track of readiness of source values
  - ❑ Broadcast the “tag” when the value is produced
  - ❑ Instructions compare their “source tags” to the broadcast tag  
→ if match, source value becomes ready
4. When all source values of an instruction are ready, need to dispatch the instruction to its functional unit (FU)
  - ❑ Instruction wakes up if all sources are ready
  - ❑ If multiple instructions are awake, need to select one per FU

# Tomasulo's Algorithm

---


- OoO with register renaming invented by Robert Tomasulo
  - Used in IBM 360/91 Floating Point Units
  - **Read:** Tomasulo, “**An Efficient Algorithm for Exploiting Multiple Arithmetic Units,**” IBM Journal of R&D, Jan. 1967.
  
- What is the major difference today?
  - **Precise exceptions:** IBM 360/91 did NOT have this
  - Patt, Hwu, Shebanow, “**HPS, a new microarchitecture: rationale and introduction,**” MICRO 1985.
  - Patt et al., “**Critical issues regarding HPS, a high performance microarchitecture,**” MICRO 1985.
  
- Variants used in most high-performance processors
  - Initially in Intel Pentium Pro, AMD K5
  - Alpha 21264, MIPS R10000, IBM POWER5, IBM z196, Oracle UltraSPARC T4, ARM Cortex A15

# Two Humps in a Modern Pipeline




- Hump 1: Reservation stations (scheduling window)
- Hump 2: Reordering (reorder buffer, aka instruction window or active window)

# General Organization of an OOO Processor



- Smith and Sohi, “**The Microarchitecture of Superscalar Processors,**” Proc. IEEE, Dec. 1995.

# Tomasulo's Machine: IBM 360/91



# Register Renaming

---

- Output and anti dependencies are not true dependencies
  - WHY? The same register refers to values that have nothing to do with each other
  - **They exist because not enough register ID's (i.e. names) in the ISA**
- The register ID is **renamed** to the reservation station entry that will hold the register's value
  - Register ID → RS entry ID
  - Architectural register ID → Physical register ID
  - After renaming, RS entry ID used to refer to the register
- This eliminates anti- and output- dependencies
  - Approximates the performance effect of a large number of registers even though ISA has a small number

# Tomasulo's Algorithm: Renaming

---

- Register rename table (register alias table)

	tag	value	valid?
R0			1
R1			1
R2			1
R3			1
R4			1
R5			1
R6			1
R7			1
R8			1
R9			1

# Tomasulo's Algorithm

---

- If reservation station available before renaming
  - Instruction + renamed operands (source value/tag) inserted into the reservation station
  - Only rename if reservation station is available
- Else stall
- While in reservation station, each instruction:
  - Watches common data bus (CDB) for tag of its sources
  - When tag seen, grab value for the source and keep it in the reservation station
  - When both operands available, instruction ready to be dispatched
- Dispatch instruction to the Functional Unit when instruction is ready
- After instruction finishes in the Functional Unit
  - Arbitrate for CDB
  - Put tagged value onto CDB (tag broadcast)
  - Register file is connected to the CDB
    - Register contains a tag indicating the latest writer to the register
    - If the tag in the register file matches the broadcast tag, write broadcast value into register (and set valid bit)
  - Reclaim rename tag
    - no valid copy of tag in system!

# An Exercise

---

MUL R3  $\leftarrow$  R1, R2  
ADD R5  $\leftarrow$  R3, R4  
ADD R7  $\leftarrow$  R2, R6  
ADD R10  $\leftarrow$  R8, R9  
MUL R11  $\leftarrow$  R7, R10  
ADD R5  $\leftarrow$  R5, R11



- Assume ADD (4 cycle execute), MUL (6 cycle execute)
- Assume one adder and one multiplier
- How many cycles
  - in a non-pipelined machine
  - in an in-order-dispatch pipelined machine with imprecise exceptions (no forwarding and full forwarding)
  - in an out-of-order dispatch pipelined machine imprecise exceptions (full forwarding)

# Exercise Continued

MUL R1, R2, → R3  
ADD R3, R4 → R5  
ADD R2, R6 → R7  
ADD R8, R9 → R10  
MUL R7, R10 → R11  
ADD R5, R11, → R5

MUL takes 6 cycles  
ADD takes 4 cycles

How many cycles total w/o data forwarding?  
" " " " w/ " " ?

Pipeline structure

F D E W

↓  
can take multiple cycles

# Exercise Continued

```

FD123456W
FD-----D1234W
F-----D1234W
          FD1234W
          FD-----D123456W
          F-----D          D1234W
    
```



Execution timeline w/ scoreboard

31 cycles


```

FD123456W
FD      ↘ E, D1234W
F       D 1 2 3 4 W
        FD1234W
        FD      ↘ 1 2 3 4.. 6 W
        F       D      ↘ 1 2 3 4 W
    
```

25 cycles

# Exercise Continued





MUL R3 ← R1, R2  
ADD R5 ← R3, R4  
ADD R7 ← R2, R6  
ADD R10 ← R8, R9  
MUL R11 ← R7, R10  
ADD R5 ← R5, R11



Tomasulo's algorithm + full forwarding

20 cycles

# How It Works



Assume  
adder &  
multiplier have  
separate  
buses



# Cycle 2

cycle 2:

MUL R1, R2 → R3 reads its sources from the RAT

- writes to its destination in the RAT  
(renames its destination)


→ allocates a reservation station entry

→ allocates a tag for its destination register

- places its sources in the reservation station entry that is allocated.

End of cycle 2:

	V	tag	value
R1	1	~	1
R2	1	~	2
R3	0	X	~
R4	1	~	4
R11	1	~	11



- MUL at X becomes ready to execute  
(What if multiple instructions become ready at the same time)  
→ both of its sources are valid in the reservation station X

cycle 3:

→ MUL at X starts execution

→ ADD R3, R4, → RS gets renamed and placed into the ADDPER reservation stations





end of cycle 3:

R1	1	~	1
R2	1	~	2
R3	0	X	~
R4	1	~	4
R5	0	9	~
R6	1	~	6
R11	1	~	11

	V	tag	value	V	tag	value
a	0	X	~	1	~	4
b	1					
c						
d						

X	1	~	1	1	~	2
Y						



- ADD at a cannot be ready to execute because one of its sources is not ready

→ It is waiting for the value with the tag X to be broadcast (by the MUL in X)

## Cycle 3

Aside: Does the tag need to be associated with the RS entry of the producer?

Answer: No: Tag is a tag for the value that is communicated.

← enables data-flow like value communication

RS is a place to hold the instructions while they become ready.

These two are completely orthogonal.

# Cycle 4



cycle 4: — ADD R2, R6 → R7 gets renamed and placed into RS (5)

end of cycle 4:

R1	1	~	1
R2	1	~	2
R3	0	X	~
R4	1	~	4
RS	0	a	~
R6	1	~	6
R7	0	b	~
R11	1	~	11

a	0	X	~	1	~	4
b	1	~	2	1	~	6
c						
d						

Same as cycle 3



— ADD at b becomes ready to execute  
(both sources are ready!)

— At cycle 5, it is sent to the adder out-of-program order!

→ It is executed before the add in a


# Cycle 7

end of cycle 7:

	V	tag	value
R1	1	~	1
R2	1	~	2
R3	0	X	~
R4	1	~	4
R5	0	d	~
R6	1	~	6
R7	0	b	~
R8	1	~	8
R9	1	~	9
R10	0	c	~
R11	0	Y	~

a	0	X	~	1	~	4
b	1	~	2	1	~	6
c	1	~	8	1	~	9
d	0	a	~	0	Y	~

X	1	~	1	1	~	2
Y	0	b	~	0	c	~



- \* All 6 instructions renamed.
- Note what happened to R5

# Cycle 8

---

cycle 8:

- MUL at X and ADD at b  
broadcast their tags and values

- RS entries waiting for these tags capture the values  
and set the Valid bit accordingly

→ (What is needed in HW to accomplish this?)

CAM on tags that are broadcast for all RS  
entries & sources

- RAT entries waiting for these tags also capture the  
values and set the Valid bits accordingly

# An Exercise, with Precise Exceptions

---

MUL R3  $\leftarrow$  R1, R2  
ADD R5  $\leftarrow$  R3, R4  
ADD R7  $\leftarrow$  R2, R6  
ADD R10  $\leftarrow$  R8, R9  
MUL R11  $\leftarrow$  R7, R10  
ADD R5  $\leftarrow$  R5, R11




- Assume ADD (4 cycle execute), MUL (6 cycle execute)
- Assume one adder and one multiplier
- How many cycles
  - in a non-pipelined machine
  - in an in-order-dispatch pipelined machine with reorder buffer (no forwarding and full forwarding)
  - in an out-of-order dispatch pipelined machine with reorder buffer (full forwarding)

# Out-of-Order Execution with Precise Exceptions

---

- **Idea:** Use a reorder buffer to reorder instructions before committing them to architectural state
- An instruction updates the register alias table (essentially a future file) when it completes execution
- An instruction updates the **architectural register file** when it is the oldest in the machine and has completed execution

# Out-of-Order Execution with Precise Exceptions



in order

out of order

in order

- Hump 1: Reservation stations (scheduling window)
- Hump 2: Reordering (reorder buffer, aka instruction window or active window)

# Enabling OoO Execution, Revisited

---

1. Link the consumer of a value to the producer
  - ❑ **Register renaming:** Associate a “tag” with each data value
2. Buffer instructions until they are ready
  - ❑ Insert instruction into **reservation stations** after renaming
3. Keep track of readiness of source values of an instruction
  - ❑ **Broadcast the “tag”** when the value is produced
  - ❑ Instructions **compare their “source tags”** to the broadcast tag
    - if match, source value becomes ready
4. When all source values of an instruction are ready, dispatch the instruction to functional unit (FU)
  - ❑ **Wakeup and select/schedule** the instruction

# Summary of OOO Execution Concepts

---

- Register renaming eliminates false dependencies, enables linking of producer to consumers
- Buffering enables the pipeline to move for independent ops
- Tag broadcast enables communication (of readiness of produced value) between instructions
- Wakeup and select enables out-of-order dispatch

# OOO Execution: Restricted Dataflow

---

- An out-of-order engine dynamically builds the dataflow graph of a piece of the program
  - which piece?
- The dataflow graph is limited to the instruction window
  - Instruction window: all decoded but not yet retired instructions
- Can we do it for the whole program?
- Why would we like to?
- In other words, how can we have a large instruction window?
- Can we do it efficiently with Tomasulo's algorithm?

# Dataflow Graph for Our Example

---

MUL R3  $\leftarrow$  R1, R2

ADD R5  $\leftarrow$  R3, R4

ADD R7  $\leftarrow$  R2, R6

ADD R10  $\leftarrow$  R8, R9

MUL R11  $\leftarrow$  R7, R10

ADD R5  $\leftarrow$  R5, R11



# State of RAT and RS in Cycle 7

end of cycle 7:

	V	tag	value
R1	1	~	1
R2	1	~	2
R3	0	X	~
R4	1	~	4
R5	0	d	~
R6	1	~	6
R7	0	b	~
R8	1	~	8
R9	1	~	9
R10	0	c	~
R11	0	Y	~

a	0	X	~	1	~	4
b	1	~	2	1	~	6
c	1	~	8	1	~	9
d	0	a	~	0	Y	~

X	1	~	1	1	~	2
Y	0	b	~	0	c	~



- \* All 6 instructions renamed.
- Note what happened to R5


# Dataflow Graph

MUL R1, R2 → R3 (x)  
ADD R3, R4 → R5 (a)  
ADD R2, R6 → R7 (b)  
ADD R8, R9 → R10 (c)  
MUL R7, R10 → R11 (y)  
ADD R5, R11 → R5 (d)

## Dataflow graph

Nodes: operations performed by the instruction

Arcs: tags in Tomasulo's algorithm



# Restricted Data Flow

---

- An out-of-order machine is a “restricted data flow” machine
  - Dataflow-based execution is restricted to the microarchitecture level
  - ISA is still based on von Neumann model (sequential execution)
  
- Remember the data flow model (at the ISA level):
  - Dataflow model: An instruction is fetched and executed in data flow order
  - i.e., when its operands are ready
  - i.e., there is no instruction pointer
  - Instruction ordering specified by data flow dependence
    - Each instruction specifies “who” should receive the result
    - An instruction can “fire” whenever all operands are received

# Questions to Ponder

---

- Why is OoO execution beneficial?
  - What if all operations take single cycle?
  - **Latency tolerance**: OoO execution tolerates the latency of multi-cycle operations by executing independent operations concurrently
  
- What if an instruction takes 500 cycles?
  - How large of an instruction window do we need to continue decoding?
  - How many cycles of latency can OoO tolerate?
  - **What limits the latency tolerance scalability of Tomasulo's algorithm?**
    - **Active/instruction window size**: determined by register file, scheduling window, reorder buffer

# Registers versus Memory, Revisited

---

- So far, we considered register based value communication between instructions
- What about memory?
- What are the fundamental differences between registers and memory?
  - Register dependences known statically – memory dependences determined dynamically
  - Register state is small – memory state is large
  - Register state is not visible to other threads/processors – memory state is shared between threads/processors (in a shared memory multiprocessor)

# Memory Dependence Handling (I)

---

- Need to obey memory dependences in an out-of-order machine
  - and need to do so while providing high performance
- Observation and Problem: Memory address is not known until a load/store executes
- Corollary 1: Renaming memory addresses is difficult
- Corollary 2: Determining dependence or independence of loads/stores need to be handled after their execution
- Corollary 3: When a load/store has its address ready, there may be younger/older loads/stores with undetermined addresses in the machine

# Memory Dependence Handling (II)

---

- When do you schedule a load instruction in an OOO engine?
  - Problem: A younger load can have its address ready before an older store's address is known
  - Known as the **memory disambiguation** problem or the **unknown address** problem
  
- Approaches
  - **Conservative**: Stall the load until all previous stores have computed their addresses (or even retired from the machine)
  - **Aggressive**: Assume load is independent of unknown-address stores and schedule the load right away
  - **Intelligent**: Predict (with a more sophisticated predictor) if the load is dependent on the/any unknown address store