

CSE 490/590 Computer Architecture

Complex Pipelining II

Steve Ko
Computer Sciences and Engineering
University at Buffalo

CSE 490/590, Spring 2011

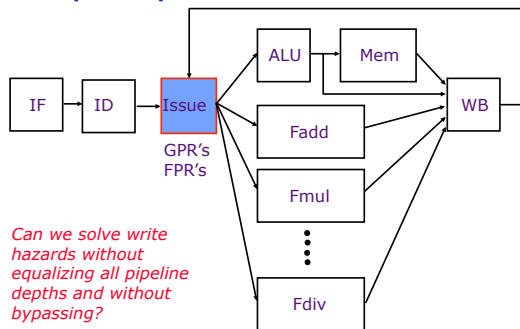
Last time...

- Complex pipelining
 - Multiple functional units with variable access time
- Types of data hazards
 - RAW, WAR, WAW
- Dependency graph
 - How instructions are dependent on each other
 - Basis for out-of-order

CSE 490/590, Spring 2011

2

Complex Pipeline



CSE 490/590, Spring 2011

3

When is it Safe to Issue an Instruction?

Suppose a data structure keeps track of all the instructions in all the functional units

The following checks need to be made before the Issue stage can dispatch an instruction

- Is the required function unit available?
- Is the input data available? \Rightarrow RAW?
- Is it safe to write the destination? \Rightarrow WAR? WAW?
- Is there a structural conflict at the WB stage?

CSE 490/590, Spring 2011

4

Types of Data Hazards

Consider executing a sequence of

type of instructions
 $r_k \leftarrow r_i \text{ op } r_j$

Data-dependence

$r_3 \leftarrow r_1 \text{ op } r_2$
 $r_5 \leftarrow r_3 \text{ op } r_4$ Read-after-Write (RAW) hazard

Anti-dependence

$r_3 \leftarrow r_1 \text{ op } r_2$
 $r_1 \leftarrow r_4 \text{ op } r_5$ Write-after-Read (WAR) hazard

Output-dependence

$r_3 \leftarrow r_1 \text{ op } r_2$
 $r_3 \leftarrow r_6 \text{ op } r_7$ Write-after-Write (WAW) hazard

CSE 490/590, Spring 2011

5

A Data Structure for Correct Issues

Keeps track of the status of Functional Units

Name	Busy	Op	Dest	Src1	Src2
Int					
Mem					
Add1					
Add2					
Add3					
Mult1					
Mult2					
Div					

The instruction i at the Issue stage consults this table

FU available? check the busy column
RAW? search the dest column for i 's sources
WAR? search the source columns for i 's destination
WAW? search the dest column for i 's destination

An entry is added to the table if no hazard is detected;
An entry is removed from the table after Write-Back

CSE 490/590, Spring 2011

6

Simplifying the Data Structure Assuming In-order Issue

Suppose the instruction is not dispatched by the Issue stage if a RAW hazard exists or the required FU is busy, and that operands are latched by functional unit on issue:

- Can the dispatched instruction cause a WAR hazard?
 - NO: Operands read at issue
- WAW hazard?
 - YES: Out-of-order completion

Simplifying the Data Structure ...

No WAR hazard
 => no need to keep src1 and src2

The Issue stage does not dispatch an instruction in case of a WAW hazard
 => a register name can occur at most once in the dest column

WP[reg#] : a bit-vector to record the registers for which writes are pending

These bits are set to true by the Issue stage and set to false by the WB stage

=> Each pipeline stage in the FU's must carry the dest field and a flag to indicate if it is valid "the (we, ws) pair"

Scoreboard for In-order Issues

Busy[FU#] : a bit-vector to indicate FU's availability. (FU = Int, Add, Mult, Div)
 These bits are hardwired to FU's.

WP[reg#] : a bit-vector to record the registers for which writes are pending.
 These bits are set to true by the Issue stage and set to false by the WB stage

Issue checks the instruction (opcode dest src1 src2) against the scoreboard (Busy & WP) to dispatch

FU available?	Busy[FU#]
RAW?	WP[src1] or WP[src2]
WAR?	cannot arise
WAW?	WP[dest]

Scoreboard Dynamics

		Functional Unit Status				Registers Reserved for Writes	
		Int(1)	Add(1)	Mult(3)	Div(4)	WB	
t0	I ₁				f6		f6
t1	I ₂	f2			f6		f6, f2
t2					f6	f2	f6, f2
t3	I ₃		f0		f6		f6, f0
t4			f0		f6		f6, f0
t5	I ₄		f0	f8		f0	f0, f8
t6				f8		f0	f0, f8
t7	I ₅	f10			f8	f10	f8, f10
t8					f8		f8, f10
t9					f8		f8
t10	I ₆	f6					f6
t11					f6		f6


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	

CSE 490/590 Administrivia

- Midterm on Friday, 3/4
- Review on Wednesday, 3/2
- Project 1 deadline: Friday, 3/11
- Office hours this week: Wed after class until 2pm

In-Order Issue Limitations: an example


1	LD	F2, 34(R2)	latency 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

In-order restriction prevents instruction 4 from being dispatched

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.

CSE 490/590, Spring 2011

13

Issue Limitations: In-Order and Out-of-Order

1	LD	F2,	34(R2)	latency 1	1	2
2	LD	F4,	45(R3)	long	4	3
3	MULTD	F6,	F4, F2	3	5	6
4	SUBD	F8,	F2, F2	1	3	6
5	DIVD	F4,	F2, F8	4	4	6
6	ADDD	F10,	F6, F4	1	6	6

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

Out-of-order execution did not allow any significant improvement!

CSE 490/590, Spring 2011

14

How many instructions can be in the pipeline?

Which features of an ISA limit the number of instructions in the pipeline?

Number of Registers

Out-of-order dispatch by itself does not provide any significant performance improvement!

CSE 490/590, Spring 2011

15

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

CSE 490/590, Spring 2011

16

Instruction-level Parallelism via Renaming

1	LD	F2,	34(R2)	latency 1	1	2
2	LD	F4,	45(R3)	long	4	3
3	MULTD	F6,	F4, F2	3	5	6
4	SUBD	F8,	F2, F2	1	3	6
5	DIVD	F4',	F2, F8	4	4	6
6	ADDD	F10,	F6, F4'	1	6	6

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 => additional storage)
 Can it be done in hardware? yes!*

CSE 490/590, Spring 2011

17

Acknowledgements

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

CSE 490/590, Spring 2011

18