Computer Architecture ELE 475 / COS 475 Slide Deck 2: Microcode and **Pipelining Review** David Wentzlaff **Department of Electrical Engineering Princeton University**





Agenda

- Microcoded Microarchitectures
- Pipeline Review
 - Pipelining Basics
 - Structural Hazards
 - Data Hazards
 - Control Hazards

Agenda

- Microcoded Microarchitectures
- Pipeline Review
 - Pipelining Basics
 - Structural Hazards
 - Data Hazards
 - Control Hazards

What Happens When the Processor is Too Large?

What Happens When the Processor is Too Large?

• Time Multiplex Resources!

Microcontrol Unit Maurice Wilkes, 1954

op conditional code flip-flop *First used in EDSAC-2, completed 1958*



Microcoded Microarchitecture



A Bus-based Datapath for RISC



Microinstruction: register to register transfer (17 control signals)

Agenda

- Microcoded Microarchitectures
- Pipeline Review
 - Pipelining Basics
 - Structural Hazards
 - Data Hazards
 - Control Hazards

An Ideal Pipeline



- All objects go through the same stages
- No sharing of resources between any two stages
- Propagation delay through all pipeline stages is equal
- Scheduling of a transaction entering the pipeline is not affected by the transactions in other stages

An Ideal Pipeline



- All objects go through the same stages
- No sharing of resources between any two stages
- Propagation delay through all pipeline stages is equal
- Scheduling of a transaction entering the pipeline is not affected by the transactions in other stages
- These conditions generally hold for industry assembly lines, but instructions depend on each other causing various hazards

Unpipelined Datapath for MIPS



Simplified Unpipelined Datapath



Pipelined Datapath



 $t_{C} > max \{t_{IM}, t_{RF}, t_{ALU}, t_{DM}, t_{RW}\} (= t_{DM} probably)$

Pipelined Control



 $t_{C} > max \{t_{IM}, t_{RF}, t_{ALU}, t_{DM}, t_{RW}\} (= t_{DM} probably)$

Pipelined Control



 $t_{C} > max \{t_{IM}, t_{RF}, t_{ALU}, t_{DM}, t_{RW}\} (= t_{DM} probably)$

Pipelined Control



t_C > max {t_{IM}, t_{RF}, t_{ALU}, t_{DM}, t_{RW}} (= t_{DM} probably) However, CPI will increase unless instructions are pipelined

17



 $t_{C} > max \{t_{IM}, t_{RF}, t_{ALU}, t_{DM}, t_{RW}\} (= t_{DM} probably)$

However, CPI will increase unless instructions are pipelined ¹⁸

"Iron Law" of Processor Performance

Time	=	<u>Instructions</u>		Cycles		<u>Time</u>
Program		Program	*	Instruction	*	Cycle

- Instructions per program depends on source code, compiler technology, and ISA
- -Cycles per instructions (CPI) depends upon the ISA and the microarchitecture
- Time per cycle depends upon the microarchitecture and the base technology

Microarchitecture	CPI	cycle time
Microcoded	>1	short
Single-cycle unpipelined	1	long
Pipelined	1	short

"Iron Law" of Processor Performance

Time	=	<u>Instructions</u>		Cycles		<u>Time</u>
Program		Program *	<	Instruction	*	Cycle

- Instructions per program depends on source code, compiler technology, and ISA
- -Cycles per instructions (CPI) depends upon the ISA and the microarchitecture
- -Time per cycle depends upon the microarchitecture and the base technology

Microarchitecture	CPI	cycle time
Microcoded	>1	short
Single-cycle unpipelined	1	long
Pipelined	1	short
Multi-cycle, unpipelined control		

"Iron Law" of Processor Performance

Time	=	Instructions		Cycles		<u>Time</u>
Program		Program [*]	*	Instruction	*	Cycle

- Instructions per program depends on source code, compiler technology, and ISA
- -Cycles per instructions (CPI) depends upon the ISA and the microarchitecture
- -Time per cycle depends upon the microarchitecture and the base technology

Microarchitecture	CPI	cycle time
Microcoded	>1	short
Single-cycle unpipelined	1	long
Pipelined	1	short
Multi-cycle, unpipelined control	>1	short

CPI Examples



3 instructions, 22 cycles, CPI=7.33

Unpipelined machine



3 instructions, 3 cycles, CPI=1

Pipelined machine



3 instructions, 3 cycles, CPI=1

Technology Assumptions

- A small amount of very fast memory (caches) backed up by a large, slower memory
- Fast ALU (at least for integers)
- Multiported Register files (slower!)

Thus, the following timing assumption is reasonable

$$t_{IM} \approx t_{RF} \approx t_{ALU} \approx t_{DM} \approx t_{RW}$$

A 5-stage pipeline will be the focus of our detailed design

- some commercial designs have over 30 pipeline stages to do an integer add!



We need some way to show multiple simultaneous transactions in both space and time

Pipeline Diagrams: Transactions vs. Time



Pipeline Diagrams: Space vs. Time



Instructions Interact With Each Other in Pipeline

- **Structural Hazard:** An instruction in the pipeline needs a resource being used by another instruction in the pipeline
- **Data Hazard:** An instruction depends on a data value produced by an earlier instruction
- Control Hazard: Whether or not an instruction should be executed depends on a control decision made by an earlier instruction

Agenda

- Microcoded Microarchitectures
- Pipeline Review
 - Pipelining Basics
 - Structural Hazards
 - Data Hazards
 - Control Hazards

Overview of Structural Hazards

- Structural hazards occur when two instructions need the same hardware resource at the same time
- Approaches to resolving structural hazards
 - Schedule: Programmer explicitly avoids scheduling instructions that would create structural hazards
 - Stall: Hardware includes control logic that stalls until earlier instruction is no longer using contended resource
 - Duplicate: Add more hardware to design so that each instruction can access independent resources at the same time
- Simple 5-stage MIPS pipeline has no structural hazards specifically because ISA was designed that way







Agenda

- Microcoded Microarchitectures
- Pipeline Review
 - Pipelining Basics
 - Structural Hazards
 - Data Hazards
 - Control Hazards

Overview of Data Hazards

- Data hazards occur when one instruction depends on a data value produced by a preceding instruction still in the pipeline
- Approaches to resolving data hazards
 - Schedule: Programmer explicitly avoids scheduling instructions that would create data hazards
 - Stall: Hardware includes control logic that freezes earlier stages until preceding instruction has finished producing data value
 - Bypass: Hardware datapath allows values to be sent to an earlier stage before preceding instruction has left the pipeline
 - Speculate: Guess that there is not a problem, if incorrect kill speculative instruction and restart

Example Data Hazard



. . .

 $r1 \leftarrow r0 + 10$ (ADDI R1, R0, #10) $r4 \leftarrow r1 + 17$ (ADDI R4, R1, #17) *r1 is stale. Oops!*

. . .

Feedback to Resolve Hazards



- Later stages provide dependence information to earlier stages which can *stall (or kill) instructions*
- Controlling a pipeline in this manner works provided the instruction at stage i+1 can complete without any interaction from instructions in stages 1 to i (otherwise deadlock)

Resolving Data Hazards with Stalls (Interlocks)



Stalled Stages and Pipeline Bubbles



Resource Usage

time t5 t6 t7 t0 t2 t3 t4 t1 I_3 I₃ I₂ I_3 I₃ IF I_1 I_5 I_4 I₂ ID I_1 $I_2 I_2 I_2$ I_4 I₃ 15 I₃ EX I_1 nop nop I_2 I_4 I_5 nop I₃ MA I_1 nop nop nop I_2 I_4 Is WB I_1 I I_{4} I_{5} nop nop nop I_2

 $nop \Rightarrow pipeline bubble$

Stall Control Logic



Compare the source registers of the instruction in the decode stage with the destination register of the uncommitted instructions.

Stall Control Logic (ignoring jumps & branches)



Should we always stall if the rs field matches some rd? not every instruction writes a register \Rightarrow we not every instruction reads a register \Rightarrow re

So	ource 8	L Dest	inati	on Regis	sters
F	R-type:	op r	s rt	rd f	unc
Ι	-type:	op r	s rt	immediate1	6
J	-type:	ор	imm	nediate26	
ALU ALUI LW SW BZ	rd \leftarrow (rs) further (rs) op rt \leftarrow (rs) op rt \leftarrow M [(rs) M [(rs) + in cond (rs) true: PC \leftarrow false: PC \leftarrow	nc (rt) immedia + immed nmediate] - (PC) + i - (PC) + 4	te iate] \leftarrow (rt) mmediat	source(s) rs, rt rs rs rs, rt rs, rt	<i>destination</i> rd rt rt
JAL	$rC \leftarrow (PC)$ $r31 \leftarrow (PC)$	$PC \leftarrow (P)$	C) + imr	nediate	31
JR JALR	$PC \leftarrow (rs)$ $r31 \leftarrow (PC)$	PC ← (rs	5)	rs rs	31

Deriving the Stall Signal



re	re1	= <i>Case</i> opcode ALU, ALUi, LW, SW, BZ, JR, JALR	⇒on
		J, JAL	\Rightarrow off
	re2	= <i>Case</i> opcode	
		ALU, SW	⇒on
			\Rightarrow OTT

$$C_{stall}$$

$$stall = ((rs_D = ws_E).we_E + (rs_D = ws_M).we_M + (rs_D = ws_W).we_W) \cdot re1_D + ((rt_D = ws_E).we_E + (rt_D = ws_M).we_M + (rt_D = ws_W).we_M) \cdot re2_D$$

This is not on ! This full story !

Hazards due to Loads & Stores



 $M[(r1)+7] \leftarrow (r2)$ r4 \leftarrow M[(r3)+5] *Is there any possible data hazard in this instruction sequence?*

Data Hazards Due to Loads and Store

- Example instruction sequence
 - Mem[Regs[r1] + 7] <- Regs[r2]</pre>
 - Regs[r4] <- Mem[Regs[r3] + 5]</p>
- What if Regs[r1]+7 == Regs[r3]+5 ?
 - Writing and reading to/from the same address
 - Hazard is avoided because our memory system completes writes in a single cycle
 - More realistic memory system will require more careful handling of data hazards due to loads and stores

Overview of Data Hazards

- Data hazards occur when one instruction depends on a data value produced by a preceding instruction still in the pipeline
- Approaches to resolving data hazards
 - Schedule: Programmer explicitly avoids scheduling instructions that would create data hazards
 - Stall: Hardware includes control logic that freezes earlier stages until preceding instruction has finished producing data value
 - Bypass: Hardware datapath allows values to be sent to an earlier stage before preceding instruction has left the pipeline
 - Speculate: Guess that there is not a problem, if incorrect kill speculative instruction and restart

Adding Bypassing to the Datapath



Deriving the Bypass Signal



47

A new datapath, i.e., a bypass, can get the data from the output of the ALU to its input

The Bypass Signal

Deriving it from the Stall Signal

stall = $(\frac{((rs_D = ws_E).we_E}{(rt_D = ws_B).we_M} + (rs_D = ws_W).we_W).re1_D + ((rt_D = ws_E).we_E + (rt_D = ws_M).we_M + (rt_D = ws_W).we_W).re2_D)$

we = *Case* opcode ALU, ALUi, LW \Rightarrow (ws \neq 0) JAL, JALR \Rightarrow on ... \Rightarrow off

 $ASrc = (rs_D = ws_E).we_E.re1_D$

Is this correct?

No because only ALU and ALUi instructions can benefit from this bypass Split we_E into two components: we-bypass, we-stall

Bypass and Stall Signals

Split we_E into two components: we-bypass, we-stall

we-bypass _F = <i>Case</i> opcode _F					
ALU, ALUi	\Rightarrow (ws \neq 0)				
	\Rightarrow off				

we-stall _F = <i>Case</i> opcode _F				
LW	\Rightarrow (ws $\neq 0$)			
JAL, JALR	\Rightarrow on			
	\Rightarrow off			

ASrc =
$$(rs_D = ws_E)$$
.we-bypass_E . re1_D

stall =
$$((rs_D = ws_E).we-stall_E + (rs_D = ws_M).we_M + (rs_D = ws_W).we_W).re1_D + ((rt_D = ws_E).we_E + (rt_D = ws_M).we_M + (rt_D = ws_W).we_W).re2_D$$

Fully Bypassed Datapath



Overview of Data Hazards

- Data hazards occur when one instruction depends on a data value produced by a preceding instruction still in the pipeline
- Approaches to resolving data hazards
 - Schedule: Programmer explicitly avoids scheduling instructions that would create data hazards
 - Stall: Hardware includes control logic that freezes earlier stages until preceding instruction has finished producing data value
 - **Bypass:** Hardware datapath allows values to be sent to an earlier stage before preceding instruction has left the pipeline
 - Speculate: Guess that there is not a problem, if incorrecturse kill speculative instruction and restart
 Jater in Contract Jater Ja

Agenda

- Microcoded Microarchitectures
- Pipeline Review
 - Pipelining Basics
 - Structural Hazards
 - Data Hazards
 - Control Hazards

Control Hazards

- What do we need to calculate next PC?
 - For Jumps
 - Opcode, offset and PC
 - For Jump Register
 - Opcode and Register value
 - For Conditional Branches
 - Opcode, PC, Register (for condition), and offset
 - For all other instructions
 - Opcode and PC
 - have to know it's not one of above!

Opcode Decoding Bubble

(assuming no branch delay slots for now)





Speculate next address is PC+4



096 ADD 100 J 304 104 ADD *kill* 304 ADD

 I_2

 I_3

 I_4

A jump instruction kills (not stalls) the following instruction

How?

Pipelining Jumps



Jump Pipeline Diagrams





 $nop \Rightarrow pipeline bubble_{57}$

Pipelining Conditional Branches



 I_4

096 ADD [100 BEQZ r1 +200 e 104 ADD 108 ... 304 ADD

Branch condition is not known until the execute stage what action should be taken in the decode stage ? 58

Pipelining Conditional Branches



If the branch is taken

096

100

104

108

304

ADD

ADD

ADD

. . .

BEQZ r1 +200

 I_1

 I_2

 I_3

 I_4

- kill the two following instructions
- the instruction at the decode stage is not valid

 \Rightarrow stall signal is not valid

Pipelining Conditional Branches

PCSrc (pc+4 / jabs / rind / br) stall



096 ADD 100 BEQZ r1 +200 104 ADD 108 ... 304 ADD

 I_1

 I_2

I₃

 I_4

If the branch is taken

- kill the two following instructions
- the instruction at the decode stage is not valid

\Rightarrow stall signal is not valid

New Stall Signal

stall = ($((rs_D = ws_E).we_E + (rs_D = ws_M).we_M + (rs_D = ws_W).we_W).re1_D$ + $((rt_D = ws_E).we_E + (rt_D = ws_M).we_M + (rt_D = ws_W).we_W).re2_D)$. $!((opcode_E = BEQZ).z + (opcode_E = BNEZ).!z)$

Don't stall if the branch is taken. Why?

Instruction at the decode stage is invalid

Control Equations for PC and IR Muxes



Give priority to the older instruction, i.e., execute-stage instruction over decode-stage instruction

Branch Pipeline Diagrams

(resolved in execute stage)



Resource Usage

time t1 t2 t3 t4 t5 t6 t7 t0 $I_2 \quad I_3 \quad I_4 \quad I_5$ IF \mathbf{I}_{1} ID $I_1 I_2 I_3 \text{ nop } I_5$ I_1 I_2 nop nop I_5 EX MA I_1 I_2 nop nop I_5 WB I_2 nop nop I_5 I_1

 $nop \Rightarrow pipeline bubble$

Reducing Branch Penalty

(resolve in decode stage)

- One pipeline bubble can be removed if an extra comparator is used in the Decode stage
 - But might elongate cycle time

PCSrc (pc+4 / jabs / rind/ br)



Pipeline diagram now same as for $jumps_{64}$

Branch Delay Slots

(expose control hazard to software)

- Change the ISA semantics so that the instruction that follows a jump or branch is always executed
 - gives compiler the flexibility to put in a useful instruction where normally a pipeline bubble would have resulted.

I ₁	096	ADD	
$\overline{I_2}$	100	BEQZ r1 +200	Delay slot instruction executed
$\overline{I_3}$	104	ADD -	Delay slot instruction executed
I ₄	304	ADD	regaraless of branch outcome

• Other techniques include more advanced branch prediction, which can dramatically reduce the branch penalty... *to come later*

Branch Pipeline Diagrams

(branch delay slot)





Why an Instruction may not be dispatched every cycle (CPI>1)

- Full bypassing may be too expensive to implement
 - typically all frequently used paths are provided
 - some infrequently used bypass paths may increase cycle time and counteract the benefit of reducing CPI
- Loads have two-cycle latency
 - Instruction after load cannot use load result
 - MIPS-I ISA defined *load delay slots*, a software-visible pipeline hazard (compiler schedules independent instruction or inserts NOP to avoid hazard). Removed in MIPS-II (pipeline interlocks added in hardware)
 - MIPS: "Microprocessor without Interlocked Pipeline Stages"
- Conditional branches may cause bubbles
 - kill following instruction(s) if no delay slots

Machines with software-visible delay slots may execute significant number of NOP instructions inserted by the compiler. NOPs not counted in useful CPI (alternatively, increase instructions/program)

Other Control Hazards

- Exceptions
- Interrupts

More on this later in the course

Agenda

- Microcoded Microarchitectures
- Pipeline Review
 - Pipelining Basics
 - Structural Hazards
 - Data Hazards
 - Control Hazards

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)
 - Christopher Batten (Cornell)
- MIT material derived from course 6.823
- UCB material derived from course CS252 & CS152
- Cornell material derived from course ECE 4750

Copyright © 2013 David Wentzlaff