# **Pipelining**\*

### Maurizio Palesi

\* Adapted from David A. Patterson's CS252 lecture slides,

http://www.cs.berkeley/~pattrsn/252S98/index.html

Copyright 1998 UCB

## References

John L. Hennessy and David A. Patterson, Computer Architecture a Quantitative Approach, second edition, Morgan Kaufmann

Chapter 3

## **Basic Steps of Execution in DLX**

- Instruction fetch step (IF)
- Instruction decode/register fetch step (ID)
- **3.** Execution/effective address step (EX)
- 4. Memory access/branch completion step (MEM)
- **5.** Register write-back step (WB)

# **DLX Datapath**



## **DLX Instruction Format**

#### **Register-Register – Register ALU operations**

0	5	6	10	11 15	16 2	0 21	27	26	31	
	Ор	Rs	1	Rs2	Rd			Орх		к-туре

#### Register-Immediate – Load/Store, ALU immed, Jump

0	5	6	10	11	15	16		31
Ор		Rs	1	R	d		Immediate	

#### Branch

0	5	6	10 11	15	16	31
Ор		Rs1	F	≀s2	Immediate	

#### Jump / Call

0	5	6 3	1
Ор		Target	J-Type

I-Type

1. Instruction fetch (IF)

 $IR \leftarrow Mem[PC], NPC \leftarrow PC + 4$ 



2. Instruction decode/register fetch (ID)

 $A \leftarrow \text{Regs}[\text{IR}(6:10)], B \leftarrow \text{Regs}[\text{IR}(11:15)]$ Imm  $\leftarrow (\text{IR}(16)^{16} \# \text{IR}(16:31))$ 





- 4. Memory access/Branch completion (MEM)
  - PC ← NPC, LMD ← Mem[ALU output] or Mem[ALU output] ← B
  - if (cond) PC  $\leftarrow$  ALU ouput



5. Register write-back (WB)

Regs[IR(16:20)] ← ALU output Regs[IR(11:15)] ← ALU output Regs[IR(11:15)] ← LMD



## Discussion

- Branch & Store instructions: 4 cycles
- All other instructions: 5 cycles
- Example
  - $\rightarrow$  Branch 12%, Store 5%
  - → CPI = (12% + 5%)\*4 + 83%\*5 = 4.83
- Improvement
  - Complete ALU instruction during MEM
  - $\rightarrow$  Ex., 47% ALU instructions
    - ✓ CPI = (12% + 5% + 47%)\*4 + 36%\*5 = 4.36
    - ✓ Speedup = 4.83/4.36 = 1.1

### Any other attempts to decrease CPI may increase the clock cycle time

## **Multicycle Implementation**

#### A simple FSM could be used to implement the control logic

 $\rightarrow$  Microcode control could be used for a much more complex machine

#### Hardware redoundancies

- → Two ALU
- $\rightarrow$  Separate instruction and data memories



# **Pipelining: Its Natural!**

- Laundry Example
- Ann, Brian, Cathy, Dave each have one load of clothes to wash, dry, and fold



"Folder" takes 20 minutes







## **Sequential Laundry**



### **Pipelined Laundry Start work ASAP**



Pipelined laundry takes 3.5 hours for 4 loads

# **Pipelining Lessons**



- Pipelining doesn't help latency of single task, it helps throughput of entire workload
- Pipeline rate limited by slowest pipeline stage
- Multiple tasks operating simultaneously
- Potential speedup = Number pipe stages
- Unbalanced lengths of pipe stages reduces speedup
- Time to "fill" pipeline and time to "drain" it reduces speedup

# Pipelining

- Technique for having multiple instructions execute in an overlapped manner
- Assembly Line
- Throughput
  - Determined by how often an instruction exits the pipeline
- Time to move one step is machine cycle time
  - Determined by the slowest step in the pipeline
  - Often it is clock cycle though clocks may have multiple phases
- Length of Pipe No of stages
  - Determine latency

# **Pipeline Performance**

Under ideal conditions

- Time per instruction = Time per instruction on non-pipelined machine/Number of pipe stages
- Speedup = Number of stages

But...

- Stages are not balanced
- Overhead (10%)

# **Basic Performance Issues**

- Pipeline increases instruction throughput
  - It does not decrease the time of execution of any single instruction
    - ✓It may increase it!
- Stage imbalance may yield further inefficiencies
- Overheads due to
  - Register and latches adding to delays and clock skew

## **Pipeline Example: 3 stages**

Assume 2 nsec latch delay

#### **Unpipelined (mono cycle)**



Latency = 42 nsec, Throughput = 1/42

#### **Pipelined**



## **CPU Time**

#### CPU Time = # Instructions $\times$ CPI $\times$ T<sub>CK</sub>

Strategy	# instructions	CPI	Tck
Single long clock cycle	=	1	big
Multiple cycles	=	many	small
Pipelining	=	1 (ideal)	small

# **Basic Pipeline DLX**

	Clock cycle							
Instruction number	1	2	3	4	5	6	7	8
Instruction i	IF	ID	EX	MEM	WB			
Instruction i+1		IF	ID	EX	MEM	WB		
Instruction i+2			IF	ID	EX	MEM	WB	
Instruction i+3				IF	ID	EX	MEM	WB
Instruction i+4					IF	ID	EX	MEM

# Making Pipeline Work

- Determine what happens at each clock tick
- Assure no resource conflicts
  - Can not use one ALU for address calculation and ALU functions
- All operations in a pipe stage must complete in one clock cycle
  - May have to elongate the clock cycle to accommodate this

# **DLX Pipeline: Datapath**



# **Pipeline Structure**

Major functional units are used in different cycles

### Three observations

- Basic datapath uses separate instruction and data memory
  - $\checkmark$  Have separate instruction and data caches
  - $\checkmark$  Memory bandwidth increases in a pipelined system
- $\rightarrow$ Register file is used in two stages ID and WB
- To start a new instruction every clock we must increment and store the PC every clock cycle during IF stage

 $\checkmark$  What happens when branches occur

# **Pipeline Structure**

- Every pipe stage is active in every cycle
  - $\rightarrow$  Values passed from one stage to next must be placed in registers
  - $\rightarrow$  Use pipeline registers or pipeline latches between stages
- Registers used to transfer information from one stage to the next
- Pipeline registers carry both control and data from stage to stage
- Values may have to be copied from one to the next stage if it is needed at a later stage
- An instruction is active in exactly one stage of the pipe at any moment

## **DLX Datapath w/o Pipelining**



## **DLX Pipeline Stages**



#### IF Stage

- → IF/ID.IR  $\leftarrow$  Mem[PC]
- → IF/ID.NPC, PC ← if (EX/MEM.Opcode == Branch && EX/MEM.cond) ← EX/MEM.ALU output else ← PC+4
- ID Stage
  - → ID/EX.A  $\leftarrow$  Regs[IF/ID.IR(6:10)]; ID/EX.B  $\leftarrow$  Regs[IF/ID.IR(11:15)]; ID/EX.NPC  $\leftarrow$  IF/ID.NPC;
  - → ID/EX.IR  $\leftarrow$  IF/ID.IR; ID/EX.Imm  $\leftarrow$  (IF/ID.IR(16))<sup>16</sup>#IF/ID.IR(16:31)



	<b>ALU</b> instruction	Load or store instruction	<b>Branch instruction</b>
EX	EX/MEM.IR ← ID/EX.IR; EX/MEM.ALUOutput← ID/EX.A op ID/EX.B; or EX/MEM.ALUOutput ←	EX/MEM.IR← ID/EX.IR EX/MEM.ALUOutput ← ID/EX.A + ID/EX.Imm;	EX/MEM.ALUOutput ← ID/EX.NPC+ID.EX.Imm;
	ID/EX.A op ID/EX.Imm; EX/MEM.cond $\leftarrow$ 0;	EX/MEM.cond $\leftarrow$ 0; EX/MEM.B $\leftarrow$ ID/EX.B;	EX/MEM.cond $\leftarrow$ (ID/EX.A op 0);
PC	4 - ADD - IF/ID Instruction IR MEM/WB.IR MEM/WB.IR 16 0	ID/EX EX Branch taken Registers	/MEM MEM/WB Data memory

MEM MEM/WB.IR ← EX/MEM.IR; MEM/WB.ALUOutput ← EX/MEM.ALUOutput; MEM/WB.IR ← EX/MEM.IR; MEM/WB.LMD ← Mem[EX/MEM.ALUOutput]; or Mem[EX/MEM.ALUOutput] ← EX/MEM.B;



WB Regs[MEM/WB.IR<sub>16..20</sub>] ← Regs[MEM/WB.IR<sub>11..15</sub>] ← MEM/WB.ALUOutput; MEM/WB.LMD; or Regs[MEM/WB.IR<sub>11..15</sub>] ← MEM/WB.ALUOutput;



## **Pipeline Hazards: Major Hurdles**

- Structural Hazards
  - Resource conflicts
- Data Hazards
  - $\rightarrow$ Data dependencies
- Control Hazards
  - $\rightarrow$ Branches and other instructions that change PC

## **Pipeline Hazards: Major Hurdles**

Structural Hazards
Resource conflicts
Data Hazards
Data dependencies
Control Hazards

 $\rightarrow$ Branches and other instructions that change PC

## **Structural Hazard**

- Some combination of instructions result in resource conflicts
  - $\rightarrow$ Some functional units are not fully pipelined
  - Some resource is not duplicated enough
    - Register file write port
    - $\checkmark$  Single memory pipeline for instruction and data
- Stall pipeline for one cycle
  - Also called a Bubble

## **Structural Hazard**



# **Pipeline Stall**

### Single memory for data and instructions

	Clock cycle								
Instruction number	1	2	3	4	5	6	7	8	9
Load instruction	IF	ID	EX	MEM	WB				
Instruction i+1		IF	ID	EX	MEM	WB			
Instruction i+2			IF	ID	EX	MEM	WB		
Instruction i+3				Stall	IF	ID	EX	MEM	WB
Instruction i+4						IF	ID	EX	MEM

## **Structural Hazard and Bubbles**



## **Solutions to Structural Hazard**

### Resource Duplication

→Example

Separate I and D caches for memory access conflict
Time-multiplexed or multiport register file for register file access conflict

## **Pipeline Hazards: Major Hurdles**

Structural Hazards
Resource conflicts
Data Hazards
Data dependencies
Control Hazards
Branches and other instructions that change PC

## Data Hazards

 Order of access to operands is changed by the pipeline vs the normal order
ADD R1, R2, R3
SUB R4, R1, R5



# **Types of Data Hazard**

- **RAW** (read after write)
  - j tries to read a source before i writes it, so j may get wrong value
- WAR (write after read)
  - $\rightarrow j$  tries to write a destination before it is read by i
  - $\rightarrow$ As reads are early this can not happen in our examples
  - $\rightarrow$ Autoincrement addressing can create it
- WAW (write after write)
  - →j tries to write an operand before it is written by i. Wrong value may remain in the register
  - $\rightarrow$ Occurs in pipes which write in more than one stage



WAR						
SW 0(R1),R2	IF	ID	EX	MEM1	MEM2	WB
ADD R2,R3,R4		IF	ID	EX	WB	
WAW						
LW R1,0(R2)	IF	ID	EX	MEM1	MEM2	WB
ADD R1,R2,R3		IF	ID	EX	WB	

# **Solutions to Data Hazard**

- (Internal) Forwarding
  - Extra hardware
- Freezing the pipeline
  - Stalls/bubbles; pipeline interlock
- Compiler (instruction) scheduling
  - Delay slot

## Data Hazard

- ADD R1, R2, R3
- **SUB** R4, R5, R1
- **AND R6**, **R1**, **R7**
- **OR R8, R1, R9**
- XOR R10, R1, R11

## Data Hazard



## **Forwarding Paths**





## **Forwarding**



## **Data Hazards Requiring Stalls**

### Situations were forwarding is not possible

### Load interlock

LW	R1, 0(R2)
SUB	R4, R1, R5
AND	<b>R6, R1, R7</b>
OR	<b>R8, R1, R9</b>

### Pipeline interlock hardware

- Detects Hazard
- $\rightarrow$ Stalls the pipe until hazard is cleared

### Load cannot Bypass Results to SUB



## **Pipeline Interlock Solution**



## Load Interlocks

Opcode field of ID/EX	Opcode field of IF/ID	Matching operand fields
ID/EX.IR(0.5)	IF/ID.IR(0.5)	
Load	Reg-reg ALU	ID/EX.IR(11:15)==IF/ID.IR(6:10)
Load	Reg-reg ALU	ID/EX.IR(11:15)==IF/ID.IR(11:15)
	Load, Store, ALU	
Load	imm, Branch	ID/EX.IR(11:15)==IF/ID.IR(6:10)

## **Pipeline Hazards: Major Hurdles**

Structural Hazards
Resource conflicts
Data Hazards
Data dependencies

### Control Hazards

 $\rightarrow$ Branches and other instructions that change PC

## **Control Hazards**

- Can cause a greater performance loss than do data hazards
- Recall that if instruction *i* is a *taken branch* 
  - $\rightarrow$  PC is not changed until the end of MEM



# **Stalling the Pipeline**

- The simplest method to deal with branches is stall the pipeline
- We do not want to stall the pipeline until we know that the instruction is a branch
  - $\rightarrow$  The stall does not occur until after the ID stage

Branch instruction	IF	ID	EX	MEM	WB					
Branch successor		IF	stall	stall	IF	ID	EX	MEM	WB	
Branch successor + 1		3 ()	cles ne	nalty		IF	ID	EX	MEM	WB
Branch successor + 2				Jucity			IF	ID	EX	MEM
Branch successor + 3								IF	ID	EX
Branch successor + 4									IF	ID
Branch successor + 5										IF

- The number of clock cycles in a branch stall can be reduced in two steps
  - Find out whether the branch is taken or not taken earlier in the pipeline
  - Compute the taken PC (i.e., the address of the branch target) earlier





With a separate adder and a branch decision made during ID, there is only 1 clock cycle stall on branches

Branch instruction	IF	ID	EX	MEM	WB			
Branch successor		IE	IF	ID	EX	MEM	WB	
Branch successor + 1	1 cy	cle pen	alty	IF	ID	EX	MEM	WB
Branch successor + 2					IF	ID	EX	MEM
Branch successor + 3						IF	ID	EX
Branch successor + 4							IF	ID
Branch successor + 5								IF

#### *Predict-not-taken* scheme

### Treat every branch as not taken

- Allowing the hw to continue as if the branch were not executed
  - $\checkmark$  If the branch is taken, it needs to turn the fetched instruction into a nop

<u>Dianch nul laken</u>	-				<b>.</b>				
Untaken branch instr.	IF	ID	EX	MEM	WB				
Instruction I + 1		IF	ID	EX	MEM	WB			
Instruction I + 2			IF	ID	EX	MEM	WB		
Instruction I + 3				IF	ID	EX	MEM	WB	
Instruction I + 4					IF	ID	EX	MEM	WB
	•				•				

#### 

#### **Branch taken**

Taken branch instr.	IF	ID	EX	MEM	WB				
Instruction I + 1		IF	idle	idle	idle	idle			
Branch target			IF	ID	EX	MEM	WB		
Branch target + 1				IF	ID	EX	MEM	WB	
Branch target + 2					IF	ID	EX	MEM	WB