


## Lecture 2: Review of Pipelines

DAP Spr. '98 ©UCB 1

## The Instruction Set: a Critical Interface



1/22/02

CS252/Culler  
Lec 1.2


## Instruction Set Architecture

... the attributes of a [computing] system as seen by the programmer, *i.e.* the conceptual structure and functional behavior, as distinct from the organization of the data flows and controls the logic design, and the physical implementation.

- Amdahl, Blaaw, and

Brooks, 1964

- Organization of Programmable Storage
- Data Types & Data Structures: Encodings & Representations
- Instruction Formats
- Instruction (or Operation Code) Set
- Modes of Addressing and Accessing Data Items and Instructions
- Exceptional Conditions



1/22/02

CS252/Culler  
Lec 1.3


## Organization

- Capabilities & Performance Characteristics of Principal Functional Units
  - (e.g., Registers, ALU, Shifters, Logic Units, ...)
- Ways in which these components are interconnected
- Information flows between components
- Logic and means by which such information flow is controlled.
- Choreography of FUs to realize the ISA
- Register Transfer Level (RTL) Description

Logic Designer's View

-----  
ISA Level


FUs & Interconnect



1/22/02

CS252/Culler  
Lec 1.4

## Review: MIPS R3000 (core)



2<sup>32</sup> x bytes

31 x 32-bit GPRs (R0=0)

32 x 32-bit FP regs (paired DP)

HI, LO, PC

### Arithmetic logical

Add, AddU, Sub, SubU, And, Or, Xor, Nor, SLT, SLTU,

AddI, AddIU, SLTI, SLTIU, AndI, OrI, XorI, LUI

SLL, SRL, SRA, SLLV, SRLV, SRAV

### Memory Access

LB, LBU, LH, LHU, LW, LWL, LWR

SB, SH, SW, SWL, SWR

### Control

J, JAL, JR, JALR

BEq, BNE, BLEZ, BGTZ, BLTZ, BGEZ, BLTZAL, BGEZAL

32-bit instructions on word boundary

1/22/02

CS252/Culler  
Lec 1.5

## Review: Basic ISA Classes

### Accumulator:

1 address    add A             $acc \leftarrow acc + mem[A]$   
 1+x address    addx A             $acc \leftarrow acc + mem[A + x]$

### Stack:

0 address    add             $tos \leftarrow tos + next$

### General Purpose Register:

2 address    add A B             $EA(A) \leftarrow EA(A) + EA(B)$   
 3 address    add A B C             $EA(A) \leftarrow EA(B) + EA(C)$

### Load/Store:

3 address    add Ra Rb Rc     $Ra \leftarrow Rb + Rc$   
 load Ra Rb     $Ra \leftarrow mem[Rb]$   
 store Ra Rb     $mem[Rb] \leftarrow Ra$

1/22/02

CS252/Culler  
Lec 1.6

## Instruction Formats

Variable:    ...

Fixed:

Hybrid:

### Addressing modes

-each operand requires address specifier => variable format

• code size => variable length instructions

• performance => fixed length instructions

-simple decoding, predictable operations

• With load/store instruction arch, only one memory address and few addressing modes

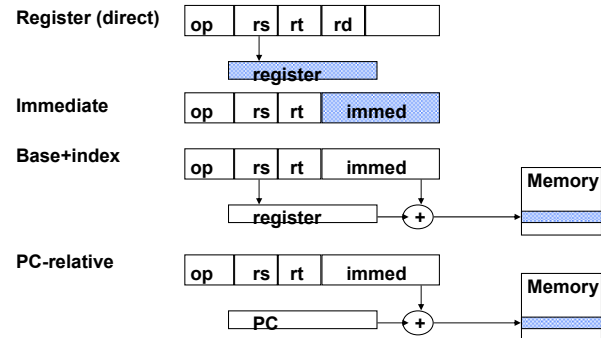
• => simple format, address mode given by opcode

1/22/02

CS252/Culler  
Lec 1.7

## MIPS Addressing Modes & Formats

- Simple addressing modes
- All instructions 32 bits wide




• Register Indirect?

1/22/02


CS252/Culler  
Lec 1.8

## Cray-1: the original RISC

Register-Register



Load, Store and Branch




1/22/02

CS252/Culler  
Lec 1.9

## VAX-11: the canonical CISC

Variable format, 2 and 3 address instruction





- Rich set of orthogonal address modes
  - immediate, offset, indexed, autoinc/dec, indirect, indirect+offset
  - applied to any operand
- Simple and complex instructions
  - synchronization instructions
  - data structure operations (queues)
  - polynomial evaluation

1/22/02

CS252/Culler  
Lec 1.10

## Review: Load/Store Architectures

- 3 address GPR
- Register to register arithmetic
- Load and store with simple addressing modes (reg + immediate)
- Simple conditionals
  - compare ops + branch z
  - compare&branch
  - condition code + branch on condition
- Simple fixed-format encoding



- Substantial increase in instructions
- Decrease in data BW (due to many registers)
- Even more significant decrease in CPI (pipelining)
- Cycle time, Real estate, Design time, Design complexity

1/22/02


CS252/Culler  
Lec 1.11

## MIPS R3000 ISA (Summary)

### Instruction Categories

- Load/Store
- Computational
- Jump and Branch
- Floating Point
  - » coprocessor
- Memory Management
- Special

Registers



3 Instruction Formats: all 32 bits wide



1/22/02

CS252/Culler  
Lec 1.12

## Levels of Representation (61C Review)

High Level Language Program

Compiler

Assembly Language Program

Assembler

Machine Language Program

Machine Interpretation

Control Signal Specification

```
temp = v[k];
v[k] = v[k+1];
v[k+1] = temp;
```

```
lw $15,0($2)
lw $16,4($2)
sw  $16,0($2)
sw  $15,4($2)
```


```
0000 1001 1100 0110 1010 1111 0101 1000
1010 1111 0101 1000 0000 1001 1100 0110
1100 0110 1010 1111 0101 1000 0000 1001
0101 1000 0000 1001 1100 0110 1010 1111
```

```
ALUOP[0:3] <= InstReg[9:11] & MASK
```

1/22/02

CS252/Culler  
Lec 1.13

## Execution Cycle



Obtain instruction from program storage

Determine required actions and instruction size

Locate and obtain operand data

Compute result value or status


Deposit results in storage for later use

Determine successor instruction

1/22/02

CS252/Culler  
Lec 1.14

## What's a Clock Cycle?





- Old days: 10 levels of gates
- Today: determined by numerous time-of-flight issues + gate delays
  - clock propagation, wire lengths, drivers

1/22/02

CS252/Culler  
Lec 1.15

## Fast, Pipelined Instruction Interpretation




1/22/02

CS252/Culler  
Lec 1.16


## Pipelining: It's Natural!

- Laundry Example
- Ann, Brian, Cathy, Dave each have one load of clothes to wash, dry, and fold
- Washer takes 30 minutes
- Dryer takes 40 minutes
- "Folder" takes 20 minutes



DAP Spr: '98 eUCB 17


## Sequential Laundry



- Sequential laundry takes 6 hours for 4 loads
- If they learned pipelining, how long would laundry take?

DAP Spr: '98 eUCB 18


## Pipelined Laundry Start work ASAP



- Pipelined laundry takes 3.5 hours for 4 loads

DAP Spr: '98 eUCB 19

## 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

DAP Spr: '98 eUCB 20

## Computer Pipelines

- Execute billions of instructions, so throughput is what matters
- DLX desirable features: all instructions same length, registers located in same place in instruction format, memory operands only in loads or stores
- + N'est pas visible au programmeur

DAP Spr.'98 eUCB 21

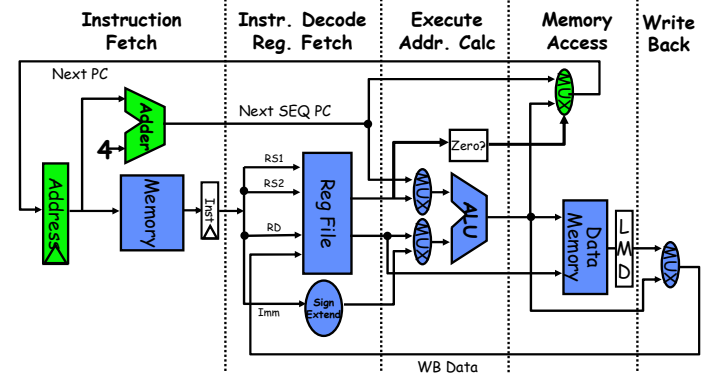
## Steps 1 & 2

- IF - instruction fetch step
  - IR  $\leftarrow$  Mem[ PC ]: fetch the next instruction from memory
  - NPC  $\leftarrow$  PC + 4 : compute the new PC
 done in parallel with opcode decode
- ID - instruction decode and register fetch step
  - A  $\leftarrow$  Regs[ IR 6.. 10 ]
  - B  $\leftarrow$  Regs[ IR 11.. 16 ]
- Possible since register specifiers are encoded in *fixed fields*
- We may fetch register contents that we don't use but OK since the operands will be ready if the opcode is of the type that does use them
- Also calculate the sign extended immediate in case that's the value that the opcode needs

DAP Spr.'98 eUCB 23

## 5 Steps of MIPS Datapath

Figure 3.1, Page 130, CA:AQA 2e

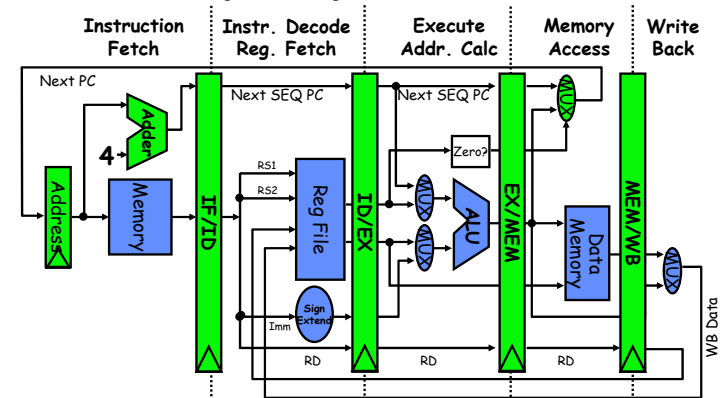


1/22/02

CS252/Culler  
Lec. 1.22

## 5 Steps of MIPS Datapath

Figure 3.4, Page 134, CA:AQA 2e



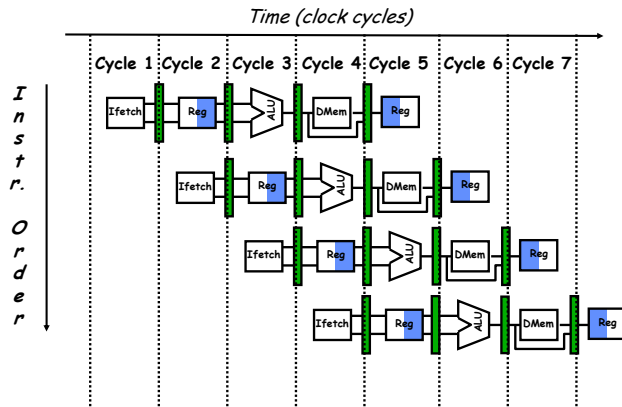
- Data stationary control
  - local decode for each instruction phase / pipeline stage

1/22/02

CS252/Culler  
Lec. 1.24

## Visualizing Pipelining

Figure 3.3, Page 133, CA:AQA Ze



## Its Not That Easy for Computers

- Limits to pipelining: **Hazards** prevent next instruction from executing during its designated clock cycle
  - **Structural hazards:** HW cannot support this combination of instructions (single person to fold and put clothes away)
  - **Data hazards:** Instruction depends on result of prior instruction still in the pipeline (missing sock)
  - **Control hazards:** Pipelining of branches & other instructions that change the PC
  - Common solution is to **stall** the pipeline until the hazard is resolved, inserting one or more **"bubbles"** in the pipeline

DAP Spr.'98 ©UCB 26

## One Memory Port/Structural Hazards

Figure 3.6, Page 142

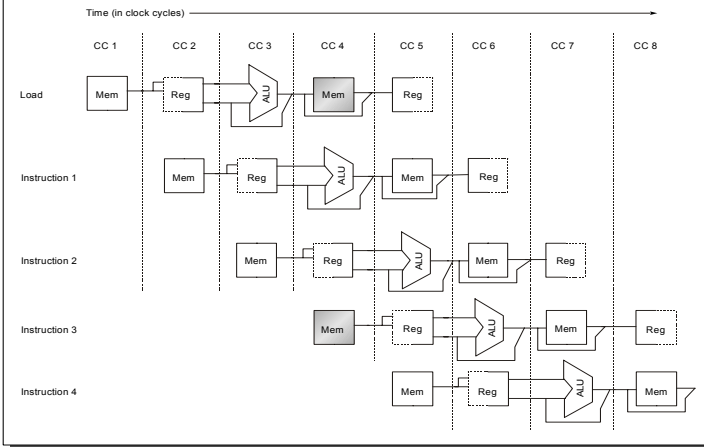
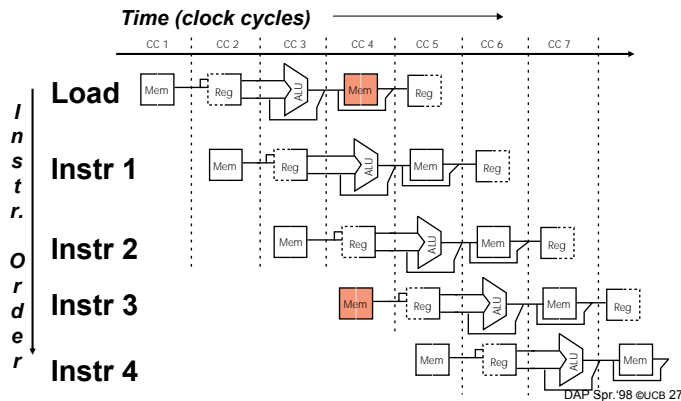
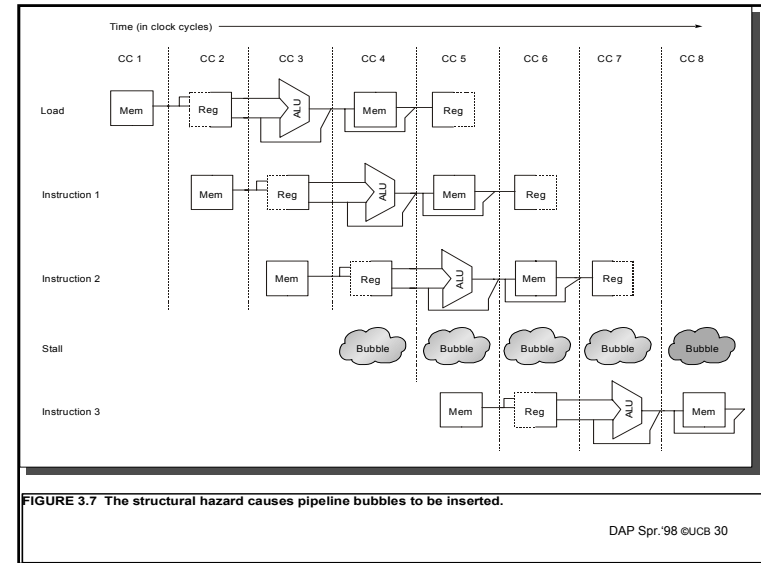
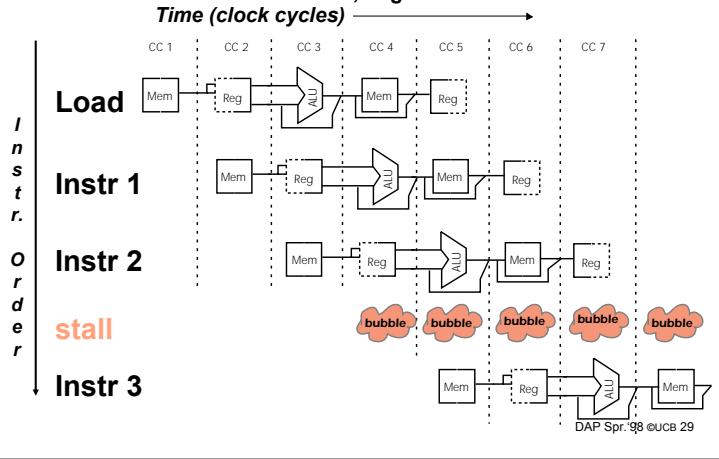


FIGURE 3.6 A machine with only one memory port will generate a conflict whenever a memory reference occurs.

DAP Spr.'98 ©UCB 28

## One Memory Port/Structural Hazards

Figure 3.7, Page 143



## Speed Up Equation for Pipelining

$$CPI_{\text{pipelined}} = \frac{\text{Ideal CPI}}{\text{Ideal CPI} + \text{Pipeline stall clock cycles per instr}}$$

$$\text{Speedup} = \frac{\text{Ideal CPI} \times \text{Pipeline depth}}{\text{Ideal CPI} + \text{Pipeline stall CPI}} \times \frac{\text{Clock Cycle}_{\text{unpipelined}}}{\text{Clock Cycle}_{\text{pipelined}}}$$

$$\text{Speedup} = \frac{\text{Pipeline depth}}{1 + \text{Pipeline stall CPI}} \times \frac{\text{Clock Cycle}_{\text{unpipelined}}}{\text{Clock Cycle}_{\text{pipelined}}}$$

DAP Spr: '98 ©UCB 31

## Example: Dual-port vs. Single-port

- Machine A: Dual ported memory
- Machine B: Single ported memory, but its pipelined implementation has a 1.05 times faster clock rate
- Ideal CPI = 1 for both
- Loads are 40% of instructions executed

$$\text{SpeedUp}_A = \frac{\text{Pipeline Depth}}{1 + 0} \times \frac{\text{clock}_{\text{unpipe}}}{\text{clock}_{\text{pipe}}} = \text{Pipeline Depth}$$

$$\begin{aligned} \text{SpeedUp}_B &= \frac{\text{Pipeline Depth}}{1 + 0.4 \times 1} \times \frac{\text{clock}_{\text{unpipe}}}{(\text{clock}_{\text{unpipe}} / 1.05)} \\ &= (\text{Pipeline Depth} / 1.4) \times 1.05 \\ &= 0.75 \times \text{Pipeline Depth} \end{aligned}$$

$$\text{SpeedUp}_A / \text{SpeedUp}_B = \text{Pipeline Depth} / (0.75 \times \text{Pipeline Depth}) = 1.33$$

- Machine A is 1.33 times faster

DAP Spr: '98 ©UCB 32



## Data Hazard on R1

Figure 3.9, page 147

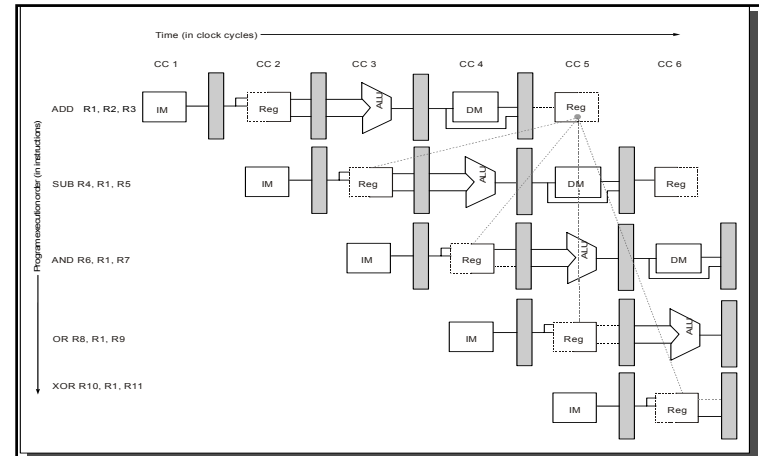
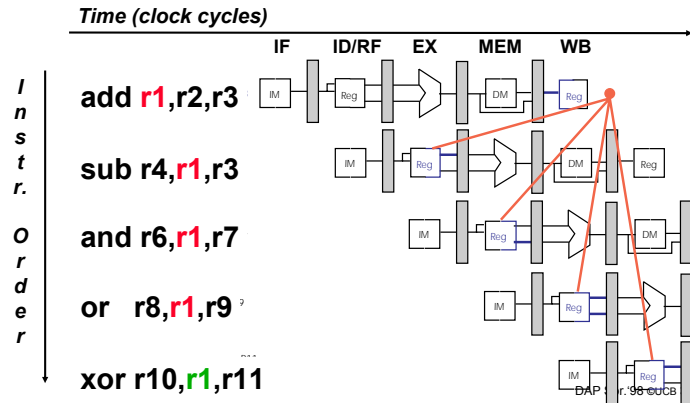


FIGURE 3.9 The use of the result of the ADD instruction in the next three instructions causes a hazard, since the register is not written until after those instructions read it.

DAP Spr.'98 ©UCB 34

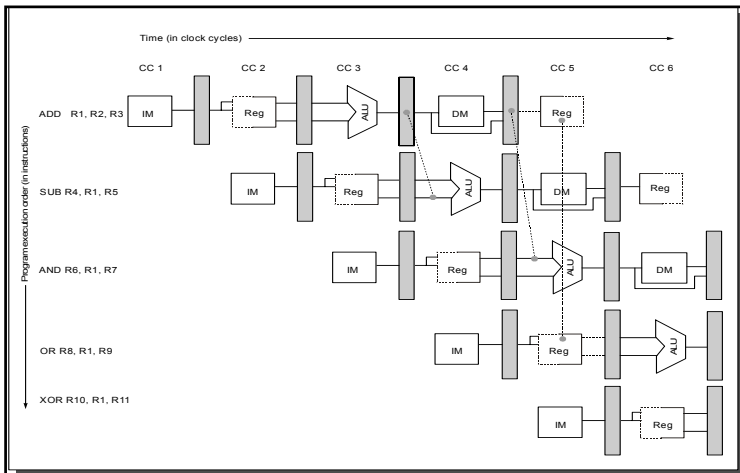


FIGURE 3.10 A set of instructions that depend on the ADD result use forwarding paths to avoid the data hazard.

DAP Spr.'98 ©UCB 35

## Three Generic Data Hazards

Instr<sub>i</sub> followed by Instr<sub>j</sub>

- **Read After Write (RAW)**  
Instr<sub>j</sub> tries to read operand before Instr<sub>i</sub> writes it

DAP Spr.'98 ©UCB 36

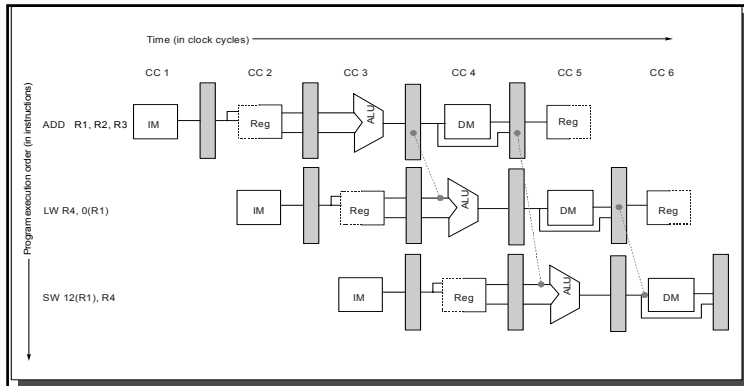


FIGURE 3.11 Stores require an operand during MEM, and forwarding of that operand is shown here.

DAP Spr: '98 eUCB 37

## Three Generic Data Hazards

Instr<sub>i</sub> followed by Instr<sub>j</sub>

- **Write After Read (WAR)**  
Instr<sub>j</sub> tries to write operand before Instr<sub>i</sub> reads it
  - Gets wrong operand
- **Can't happen in DLX 5 stage pipeline because:**
  - All instructions take 5 stages, and
  - Reads are always in stage 2, and
  - Writes are always in stage 5

DAP Spr: '98 eUCB 38

## Three Generic Data Hazards

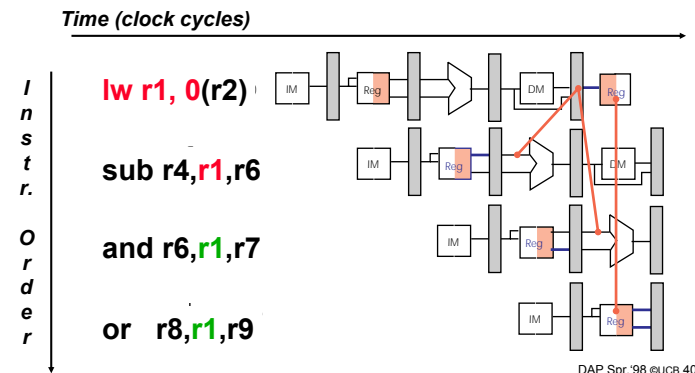
Instr<sub>i</sub> followed by Instr<sub>j</sub>

- **Write After Write (WAW)**  
Instr<sub>j</sub> tries to write operand before Instr<sub>i</sub> writes it
  - Leaves wrong result ( Instr<sub>i</sub> not Instr<sub>j</sub> )
- **Can't happen in DLX 5 stage pipeline because:**
  - All instructions take 5 stages, and
  - Writes are always in stage 5
- Will see WAR and WAW in later more complicated pipes

DAP Spr: '98 eUCB 39

## Data Hazard Even with Forwarding

Figure 3.12, Page 153



DAP Spr: '98 eUCB 40

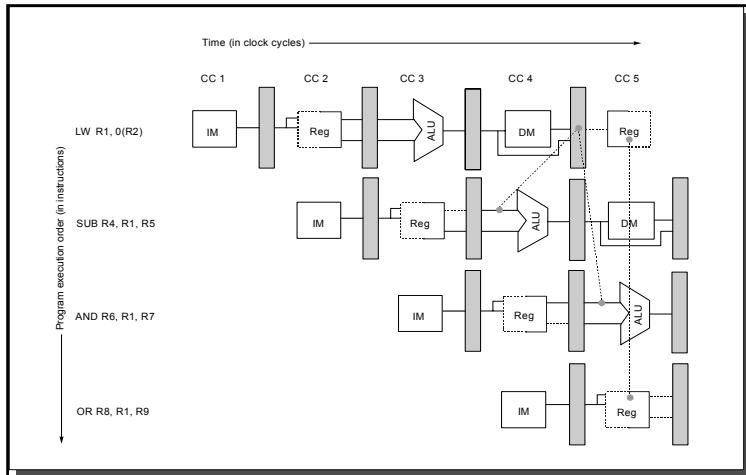


FIGURE 3.12 The load instruction can bypass its results to the AND and OR instructions, but not to the SUB, since that would mean forwarding the result in "negative time."  
DAP Spr: '98 ©UCB 41

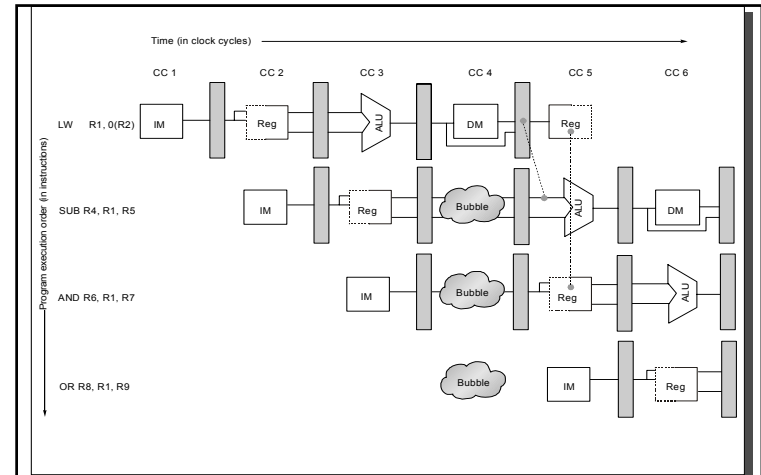


FIGURE 3.13 The load interlock causes a stall to be inserted at clock cycle 4, delaying the SUB instruction and those that follow by one cycle.  
DAP Spr: '98 ©UCB 42

$$A = B + C$$

lw rb, b	IF	ID	EX	MEM	WB				
lw rc, c		IF	ID	EX	MEM	WB			
add ra, rb, rc			IF	ID	Cale	EX	MEM	WB	
sw a, ra				IF	Cale	ID	EX	MEM	WB

DAP Spr: '98 ©UCB 43

## Software Scheduling to Avoid Load Hazards

Try producing fast code for

$a = b + c;$

$d = e - f;$

assuming a, b, c, d, e, and f in memory.

Slow code:


```

LW   Rb,b
LW   Rc,c
ADD  Ra,Rb,Rc
SW   a,Ra
LW   Re,e
LW   Rf,f
SUB  Rd,Re,Rf
SW   d,Rd


```

DAP Spr: '98 ©UCB 44

## Data Flow Graph



## Scheduling using DFG



## Software Scheduling to Avoid Load Hazards

Try producing fast code for

$a = b + c;$

$d = e - f;$

assuming  $a, b, c, d, e,$  and  $f$  in memory.

Slow code:

LW Rb,b  
 LW Rc,c  
 ADD Ra,Rb,Rc  
 SW a,Ra  
 LW Re,e  
 LW Rf,f  
 SUB Rd,Re,Rf  
 SW d,Rd

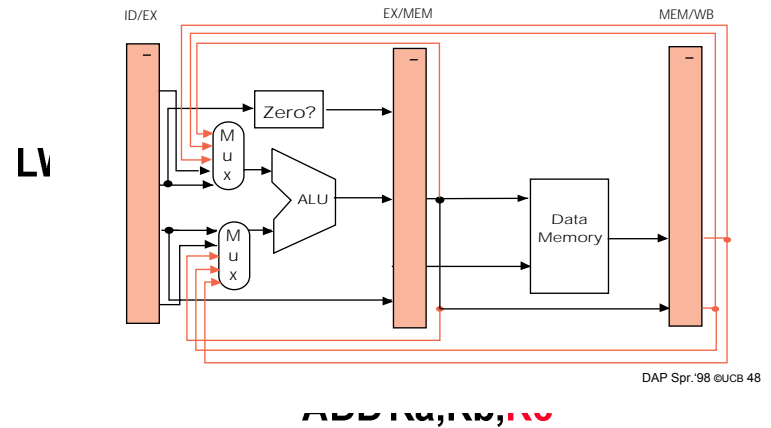
Fast code:

LW Rb,b  
 LW Rc,c  
 LW Re,e  
 ADD Ra,Rb,Rc  
 LW Rf,f  
 SW a,Ra  
 SUB Rd,Re,Rf  
 SW d,Rd

DAP Spr: '98 eucb 47

## HW Change for Forwarding

Figure 3.20, Page 161



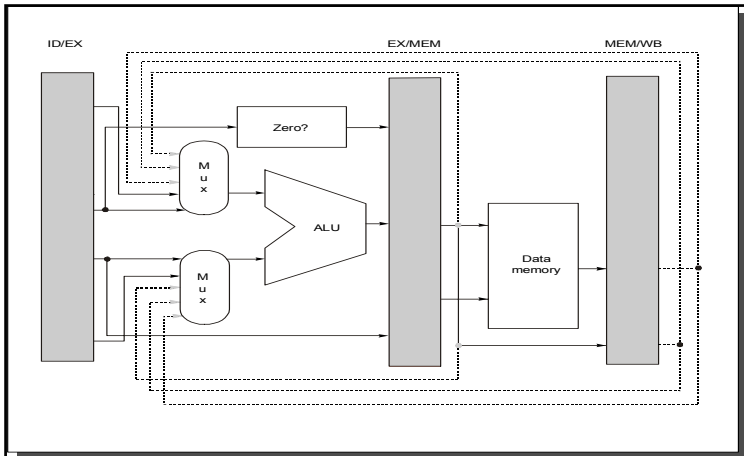


FIGURE 3.20 Forwarding of results to the ALU requires the addition of three extra inputs on each ALU multiplexer and the addition of three paths to the new inputs. DAP Spr: '98 ©UCB 49

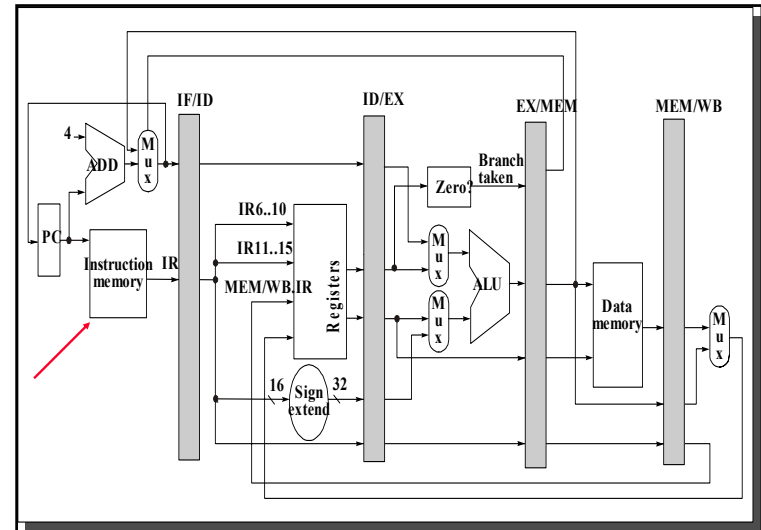



FIGURE 3.4 The datapath is pipelined by adding a set of registers, one between each pair of pipe stages. DAP Spr: '98 ©UCB 49

## Branch Hazards

- When we decide to branch, other instructions are in the pipeline!



- We are predicting “branch not taken”
  - need to add hardware for flushing instructions if we are wrong

Control Hazard on Branches Three Stage Stall 51

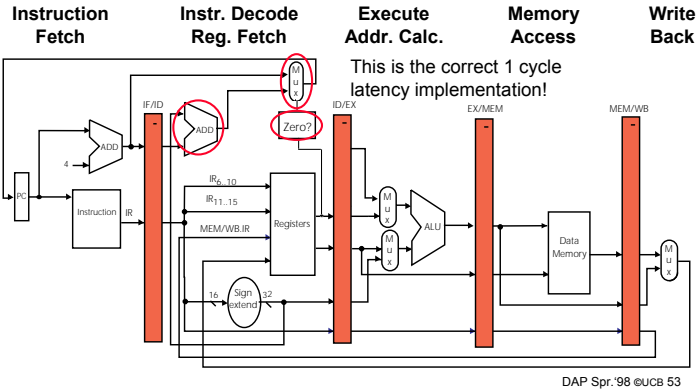
## Branch Stall Impact

- If CPI = 1, 30% branch, Stall 3 cycles => new CPI = 1.9!
- Two part solution:
  - Determine branch taken or not sooner, AND
  - Compute taken branch address earlier
- DLX branch tests if register = 0 or not 0
- DLX Solution:
  - Move Zero test to ID/RF stage
  - Adder to calculate new PC in ID/RF stage
  - 1 clock cycle penalty for branch versus 3

DAP Spr: '98 ©UCB 52

# Pipelined DLX Datapath

Figure 3.22, page 163



DAP Spr: '98 eUCB 53

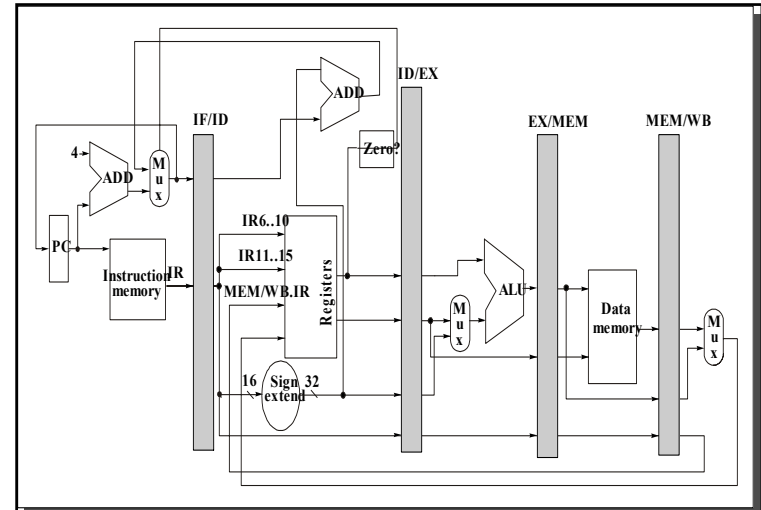




FIGURE 3.22 The stall from branch hazards can be reduced by moving the zero test and branch target calculation into the ID phase of the pipeline.



DAP Spr: '98 eUCB 55

FIGURE 3.24 The frequency of instructions (branches, jumps, calls, and returns) that may change the PC.



DAP Spr: '98 eUCB 56

FIGURE 3.25 Together the forward and backward taken branches account for an average of 67% of all conditional branches.

## Four Branch Hazard Alternatives

**#1: Stall until branch direction is clear**

**#2: Predict Branch Not Taken**

- Execute successor instructions in sequence
- "Squash" instructions in pipeline if branch actually taken
- Advantage of late pipeline state update
- 47% DLX branches not taken on average
- PC+4 already calculated, so use it to get next instruction

**#3: Predict Branch Taken**

- 53% DLX branches taken on average
- **But haven't calculated branch target address in DLX**
  - » DLX still incurs 1 cycle branch penalty
  - » Other machines: branch target known before outcome

DAP Spr.'98 eucb 57

## Four Branch Hazard Alternatives

**#4: Delayed Branch**

- Define branch to take place **AFTER** a following instruction

```

branch instruction
sequential successor1
sequential successor2
.....
sequential successorn
branch target if taken
    
```

Branch delay of length *n*

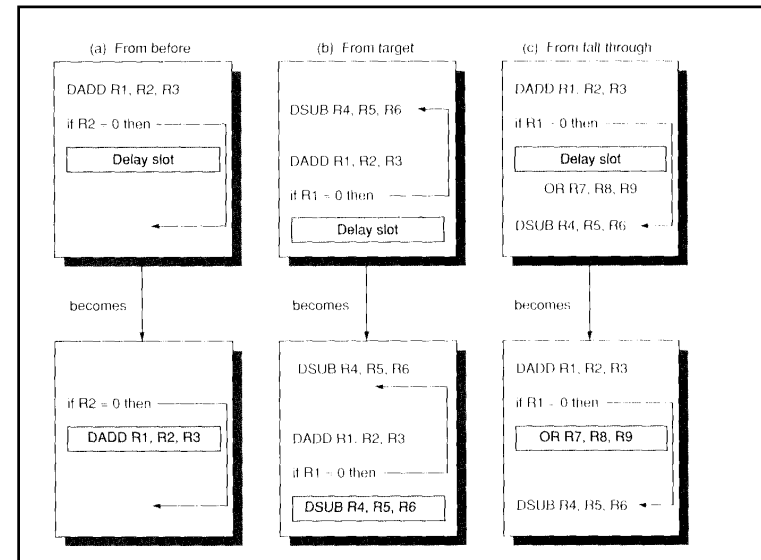
- 1 slot delay allows proper decision and branch target address in 5 stage pipeline
- DLX uses this

DAP Spr.'98 eucb 58

## Delayed Branch

- **Where to get instructions to fill branch delay slot?**
  - Before branch instruction
  - From the target address: only valuable when branch taken
  - From fall through: only valuable when branch not taken
  - Cancelling branches allow more slots to be filled
- **Compiler effectiveness for single branch delay slot:**
  - Fills about 60% of branch delay slots
  - About 80% of instructions executed in branch delay slots useful in computation
  - About 50% (60% x 80%) of slots usefully filled
- **Delayed Branch downside: 7-8 stage pipelines, multiple instructions issued per clock (superscalar)**

DAP Spr.'98 eucb 59



## Evaluating Branch Alternatives

$$\text{Pipeline speedup} = \frac{\text{Pipeline depth}}{1 + \text{Branch frequency} \times \text{Branch penalty}}$$

<i>Scheduling scheme</i>	<i>Branch penalty</i>	<i>CPI</i>	<i>speedup v. unpipelined</i>	<i>speedup v. stall</i>
Stall pipeline	3	1.42	3.5	1.0
Predict taken	1	1.14	4.4	1.26
Predict not taken	1	1.09	4.5	1.29
Delayed branch	0.5	1.07	4.6	1.31

Conditional & Unconditional = 14%, 65% change PC

DAP Spr.'98 ©UCB 61

## Pipelining Introduction Summary

- Just overlap tasks, and easy if tasks are independent
- Speed Up → Pipeline Depth; if ideal CPI is 1, then:

$$\text{Speedup} = \frac{\text{Pipeline Depth}}{1 + \text{Pipeline stall CPI}} \times \frac{\text{Clock Cycle Unpipelined}}{\text{Clock Cycle Pipelined}}$$

- Hazards limit performance on computers:
  - Structural: need more HW resources
  - Data (RAW,WAR,WAW): need forwarding, compiler scheduling
  - Control: delayed branch, prediction

DAP Spr.'98 ©UCB 62