

EEL 5764: Graduate Computer Architecture

Appendix A - Pipelining Review

Ann Gordon-Ross
Electrical and Computer Engineering
University of Florida

<http://www.ann.ece.ufl.edu/>

These slides are provided by:
David Patterson
Electrical Engineering and Computer Sciences, University of California, Berkeley
Modifications/additions have been made from the originals

What is Pipelining?

- **Overlapping execution to produce faster results**
 - Washing and drying dishes
 - Washing and drying laundry
 - Automobile assembly line
 - Chipotle, Quiznos, etc
- **Pipelining in computer architecture**
 - Multiple instructions are overlapped in execution
 - Exploits parallelism
 - Not visible to programmer
- **Each stage is a pipeline “cycle”**
 - Each stage happens simultaneously so results are produced only as fast as the *longest* pipeline cycle
 - Determines clock cycle time

9/5/07

2

Outline

- MIPS – An ISA for Pipelining
- 5 stage pipelining
- Structural and Data Hazards
- Forwarding
- Branch Schemes
- Exceptions and Interrupts

9/5/07

3

A "Typical" RISC ISA (Load/Store)


- 32-bit fixed format instruction (3 formats)
- 32 32-bit GPR (R0 contains zero)
- ALU instructions
 - 3-address, reg-reg arithmetic instruction
 - 2-address, reg-im arithmetic instruction
- Single address mode for load/store:
base + displacement
 - no indirection
- Simple branch conditions
- Delayed branch

9/5/07


4

Example: MIPS (- MIPS)


Register-Register




Register-Immediate



Branch




Jump / Call



9/5/07

5

Datapath vs Control (FSM+D)



- **Datapath:** Storage, FU, interconnect sufficient to perform the desired functions
 - Inputs are Control Points
 - Outputs are signals
- **Controller:** State machine to orchestrate operation on the data path

9/5/07 - Based on desired function and signals

6

Approaching an ISA

- **Instruction Set Architecture**
 - Defines set of operations, instruction format, hardware supported data types, named storage, addressing modes, sequencing
- **Meaning of each instruction is described by RTL on *architected registers* and memory**
- **Given technology constraints assemble adequate datapath**
 - Architected storage mapped to actual storage
 - Function units to do all the required operations
 - Possible additional storage (eg. MAR, MBR, ...)
 - Interconnect to move information among regs and FUs
- **Implement controller (Finite State Machine (FSM))**

9/5/07

7

Outline

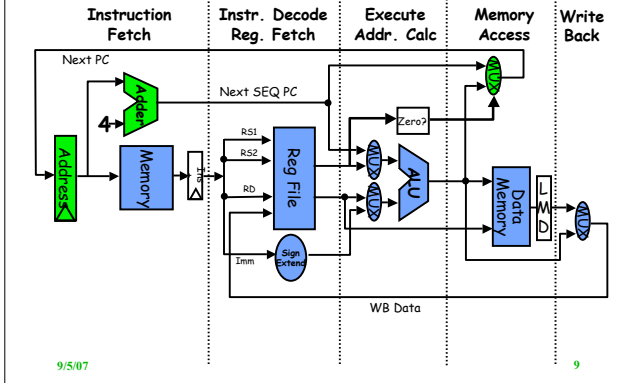
- MIPS – An ISA for Pipelining
- **5 stage pipelining**
- **Structural and Data Hazards**
- **Forwarding**
- **Branch Schemes**
- **Exceptions and Interrupts**

9/5/07

8

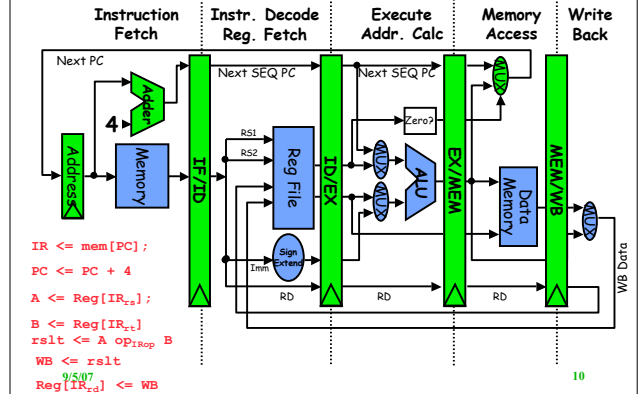
5 Steps of MIPS Datapath

Figure A.2, Page A-8

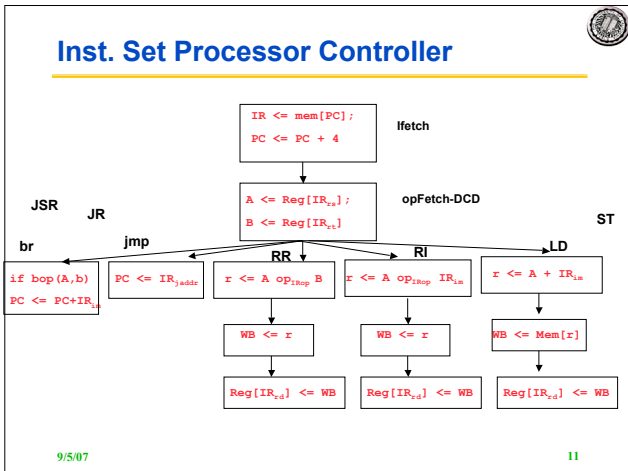


5 Steps of MIPS Datapath

Figure A.3, Page A-9

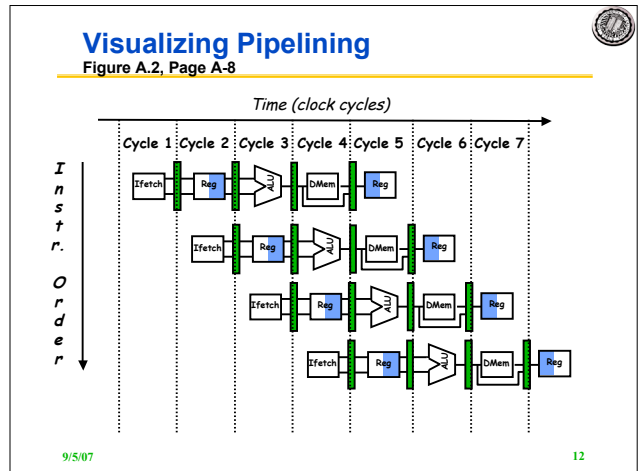


Inst. Set Processor Controller



Visualizing Pipelining

Figure A.2, Page A-8



Pipelining is not quite that easy!

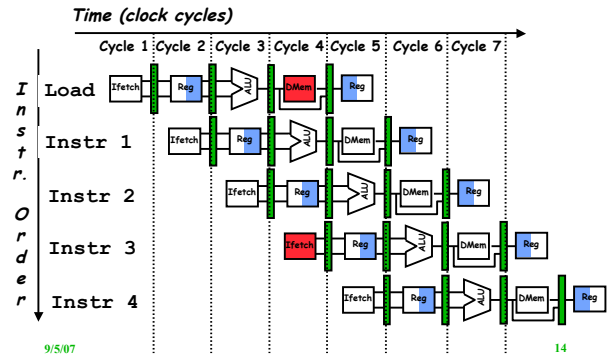
- Limits to pipelining: **Hazards** prevent next instruction from executing during its designated clock cycle
 - **Structural hazards:** HW cannot support this combination of instructions (single person to fold and put clothes away)
 - **Data hazards:** Instruction depends on result of prior instruction still in the pipeline (missing sock)
 - **Control hazards:** Caused by delay between the fetching of instructions and decisions about changes in control flow (branches and jumps).

9/5/07

13

One Memory Port/Structural Hazards

Figure A.4, Page A-14

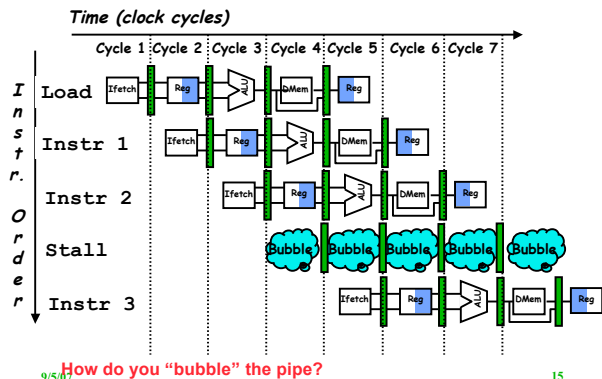


9/5/07

14

One Memory Port/Structural Hazards

(Similar to Figure A.5, Page A-15)



How do you "bubble" the pipe?

9/5/07

15

Speed Up Equation for Pipelining

$$CPI_{\text{pipelined}} = \text{Ideal CPI} + \text{Average Stall cycles per Inst}$$

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

For simple RISC pipeline, $CPI = 1$:

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

9/5/07

16

Example: Dual-port vs. Single-port

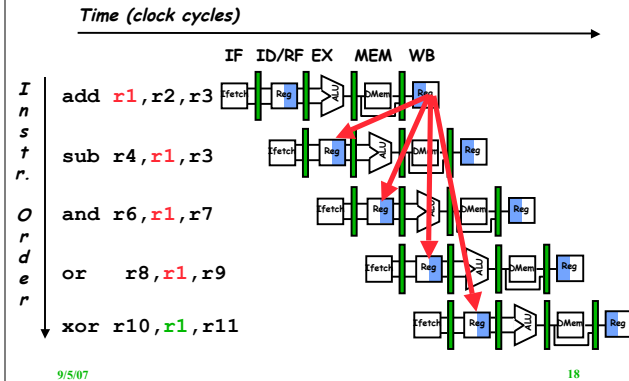
- Machine A: Dual ported memory ("Harvard Architecture")
- Machine B: Single ported memory, but its pipelined implementation has a 1.05 times faster clock rate
- Ideal CPI = 1 for both
- Loads are 40% of instructions executed
 - $SpeedUp_A = \text{Pipeline Depth} / (1 + 0) \times (\text{clock}_{unpipe} / \text{clock}_{pipe})$
= Pipeline Depth
 - $SpeedUp_B = \text{Pipeline Depth} / (1 + 0.4 \times 1) \times (\text{clock}_{unpipe} / (\text{clock}_{unpipe} / 1.05))$
= (Pipeline Depth / 1.4) \times 1.05
= 0.75 \times Pipeline Depth
 - $SpeedUp_A / SpeedUp_B = \text{Pipeline Depth} / (0.75 \times \text{Pipeline Depth}) = 1.33$
- Machine A is 1.33 times faster

9/5/07

17

Data Hazard on R1

Figure A.6, Page A-17



18

Three Generic Data Hazards

- Read After Write (RAW)**
Instr_j tries to read operand before Instr_i writes it

```

I: add r1, r2, r3
J: sub r4, r1, r3
    
```

- Caused by a "Dependence" (in compiler nomenclature). This hazard results from an actual need for communication.

9/5/07

19

Three Generic Data Hazards

- Write After Read (WAR)**
Instr_j writes operand *before* Instr_i reads it

```

I: sub r4, r1, r3
J: add r1, r2, r3
K: mul r6, r1, r7
    
```

- Called an "anti-dependence" by compiler writers. This results from reuse of the name "r1".
- Can't happen in MIPS 5 stage pipeline because:
 - All instructions take 5 stages, and
 - Reads are always in stage 2, and
 - Writes are always in stage 5

9/5/07

20

Three Generic Data Hazards

- **Write After Write (WAW)**
Instr_j writes operand *before* Instr_i writes it.
- ```

 I: sub r1, r4, r3
 J: add r1, r2, r3
 K: mul r6, r1, r7

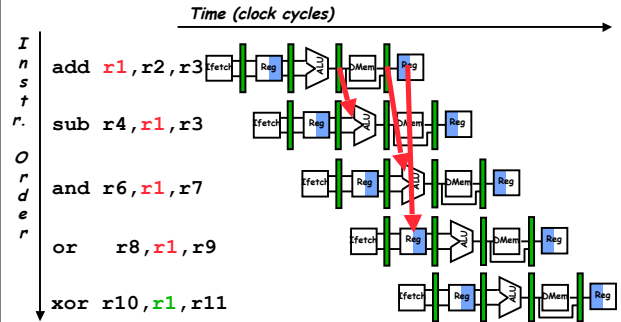
```
- Called an “**output dependence**” by compiler writers  
This also results from the reuse of name “**r1**”.
  - Can’t happen in MIPS 5 stage pipeline because:
    - All instructions take 5 stages, and
    - Writes are always in stage 5
  - Will see WAR and WAW in more complicated pipes

9/5/07

21

### Forwarding to Avoid Data Hazard

Figure A.7, Page A-19




9/5/07

22

### HW Change for Forwarding

Figure A.23, Page A-37



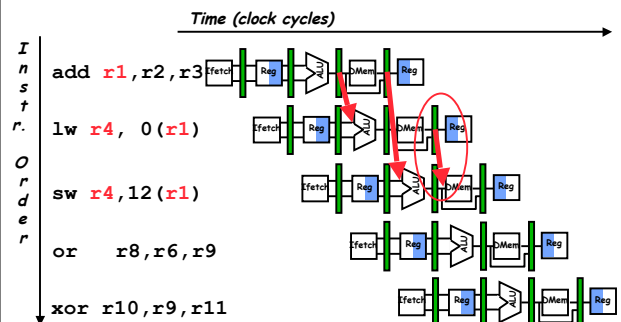
What circuit detects and resolves this hazard?

9/5/07

23

### Forwarding to Avoid LW-SW Data Hazard

Figure A.8, Page A-20

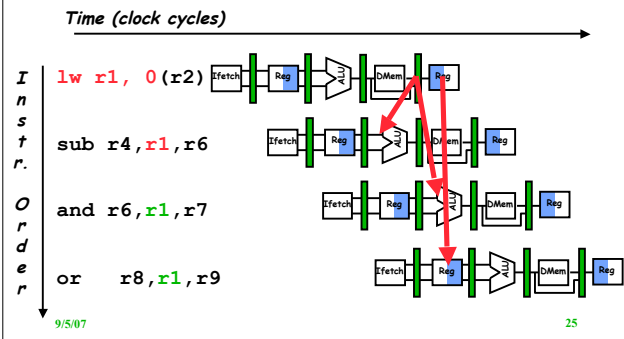


9/5/07

24

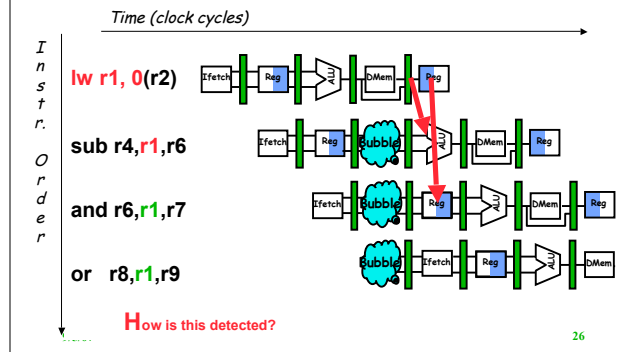
## Data Hazard Even with Forwarding

Figure A.9, Page A-21



## Data Hazard Even with Forwarding

(Similar to Figure A.10, Page A-21)



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

Fast code:

Compiler optimizes for performance. Hardware checks for safety.

9/5/07

27

## Outline

- MIPS – An ISA for Pipelining
- 5 stage pipelining
- Structural and Data Hazards
- Forwarding
- Branch Schemes
- Exceptions and Interrupts
- Conclusion

9/5/07


28

## Control Hazard on Branches Three Stage Stall

```

10: beq r1,r3,36
14: and r2,r3,r5
18: or r6,r1,r7
22: add r8,r1,r9
36: xor r10,r1,r11

```



What do you do with the 3 instructions in between?  
How do you do it?  
Where is the "commit"?

9/5/07

29

## Branch Stall Impact

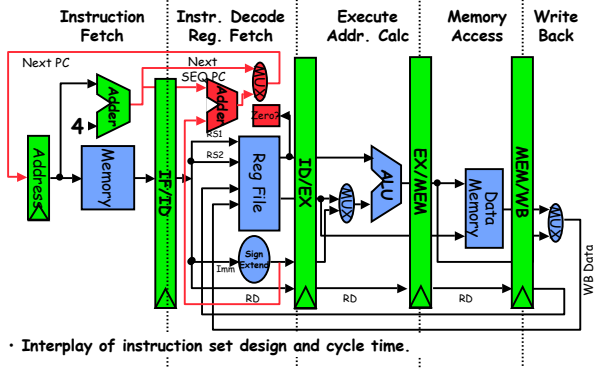
- If CPI = 1, 30% branch, Stall 3 cycles => new CPI = 1.9!
- Two part solution:
  - Determine branch taken or not sooner, AND
  - Compute taken branch address earlier
- MIPS branch tests if register = 0 or  $\neq 0$
- MIPS Solution:
  - Move Zero test to ID/RF stage
  - Adder to calculate new PC in ID/RF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

9/5/07

30

## Pipelined MIPS Datapath

Figure A.24, page A-38



9/5/07

31

## 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% MIPS branches not taken on average
  - PC+4 already calculated, so use it to get next instruction
- #3: Predict Branch Taken
  - 53% MIPS branches taken on average
  - But haven't calculated branch target address in MIPS
    - » MIPS still incurs 1 cycle branch penalty
    - » Other machines: branch target known before outcome

9/5/07

32



## 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
- MIPS uses this

9/5/07

33

## Scheduling Branch Delay Slots (Fig A.14)

A. From before branch      B. From branch target      C. From fall through



- A is the best choice, fills delay slot & reduces instruction count (IC)
- In B, the `sub` instruction may need to be copied, increasing IC
- In B and C, must be okay to execute `sub` when branch fails

9/5/07

34

## Delayed Branch

- 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: As processor go to deeper pipelines and multiple issue, the branch delay grows and need more than one delay slot**
  - Delayed branching has lost popularity compared to more expensive but more flexible dynamic approaches
  - Growth in available transistors has made dynamic approaches relatively cheaper

9/5/07

35

## Evaluating Branch Alternatives

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

Assume 4% unconditional branch, 6% conditional branch-untaken, 10% conditional branch-taken

| Scheduling scheme | Branch penalty | CPI  | speedup v. unpipelined | speedup v. stall |
|-------------------|----------------|------|------------------------|------------------|
| Stall pipeline    | 3              | 1.60 | 3.1                    | 1.0              |
| Predict taken     | 1              | 1.20 | 4.2                    | 1.33             |
| Predict not taken | 1              | 1.14 | 4.4                    | 1.40             |
| Delayed branch    | 0.5            | 1.10 | 4.5                    | 1.45             |

9/5/07

36