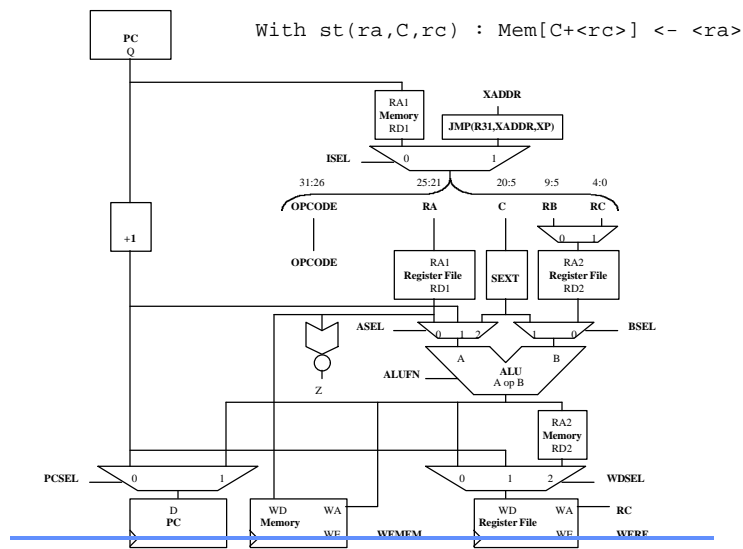


How Computers Work

Lecture 13

Details of the Pipelined Beta

Review: 1-Stage Beta (Top-Down View)



Review: Pipeline Stages

GOAL: Maintain (nearly) 1.0 CPI, but increase clock speed.

APPROACH: structure processor as 4-stage pipeline:


Instruction Fetch stage: Maintains PC, fetches one instruction per cycle and passes it to

Register File stage: Reads source operands from register file, passes them to

ALU stage: Performs indicated operation, passes result to

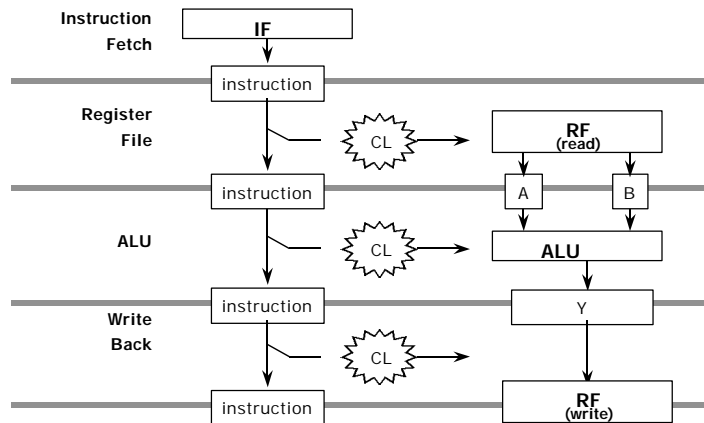
Write-Back stage: writes result back into register file.

WHAT OTHER information do we have to pass down the pipeline?



How Computers Work Lecture 13 Page 3

Sketch of 4-Stage Pipeline

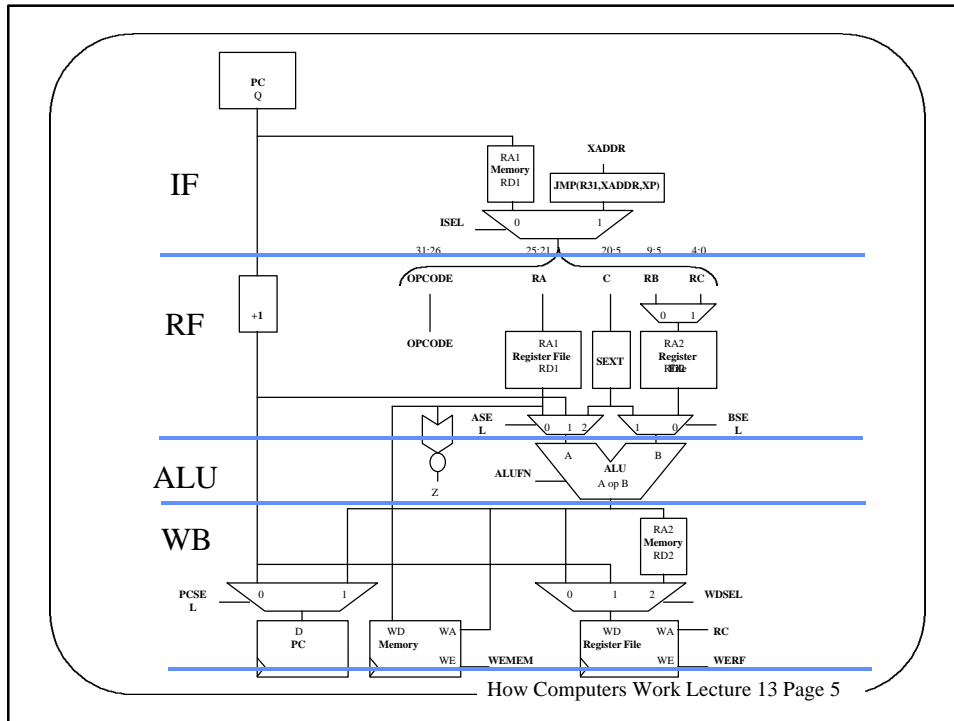


NEED ALSO:

<PC> - for Branch/JMP instrs!

Lets make it *real*....

How Computers Work Lecture 13 Page 4



Pipeline Hazards

PROBLEM:

Contents of a register WRITTEN by instruction k is READ by instruction k+1... before its stored in RF! EG:

```

ADD(r1, r2, r3)
CMPLEC(r3, 100, r0)

```

fails since CMPLEC sees "stale" <r3>.

Time →

How Computers Work Lecture 13 Page 6

ADD(r1,r2,r3)	IF	RF	ALU	WB					
CMPLEC(r3,100,r0)		IF	RF	ALU	WB				
			IF	RF	ALU	WB			
CMPLEC(r3,100,r0)			IF	RF	ALU	WB			

↓ (from row 2, column 2 to row 4, column 2)
→ (from row 1, column 4 to row 2, column 3)
→ (from row 2, column 4 to row 4, column 3)
→ (from row 4, column 3 to row 4, column 4)

SOLUTIONS:

- "Program around it".
 - ... document weirdo semantics, declare it a software problem.
 - Breaks sequential semantics!
 - Costs code efficiency.

EXAMPLE: Rewrite

ADD(r1, r2, r3)	as	ADD(r1, r2, r3)
CMPLEC(r3, 100, r0)		MULC(r1, 100, r4)
MULC(r1, 100, r4)		SUB(r1, r2, r5)
SUB(r1, r2, r5)		CMPLEC(r3, 100, r0)

HOW OFTEN can we do this?

How Computers Work Lecture 13 Page 7

SOLUTIONS:

- Stall the pipeline.
 - Freeze IF, RF stages for 2 cycles, inserting NOPs into ALU IR...

ADD(r1,r2,r3)	IF	RF	ALU	WB					
NOP		IF	RF	ALU	WB				
NOP			IF	RF	ALU	WB			
CMPLEC(r3,100,r0)			IF	RF	ALU	WB			

↓ (from row 1, column 4 to row 4, column 4)
→ (from row 1, column 4 to row 2, column 3)
→ (from row 2, column 4 to row 4, column 3)
→ (from row 4, column 3 to row 4, column 4)

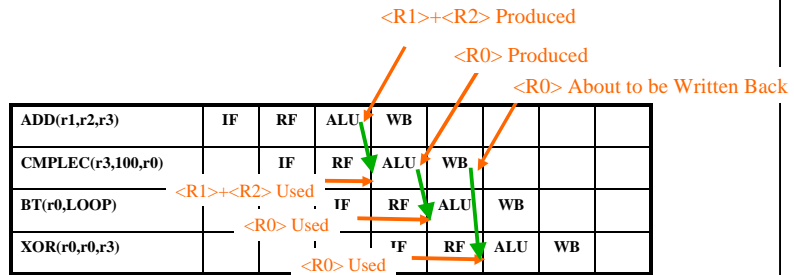
DRAWBACK: SLOW

How Computers Work Lecture 13 Page 8

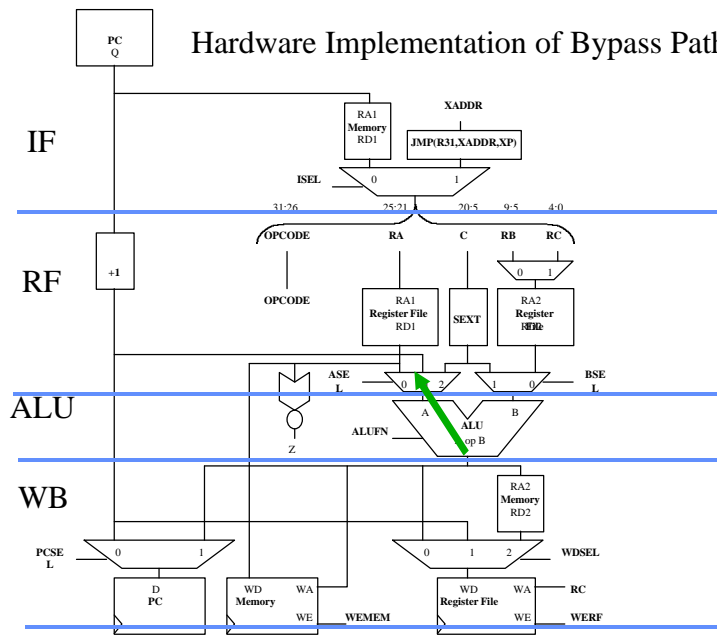
SOLUTIONS:

3. Bypass Paths.

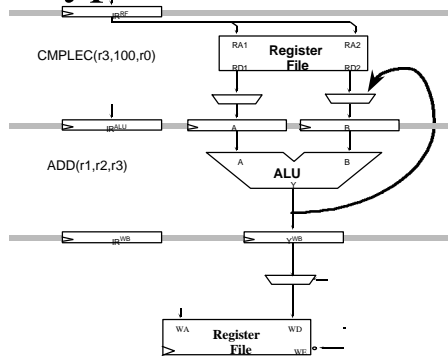
Add extra data paths & control logic to re-route data in problem cases.



Hardware Implementation of Bypass Paths

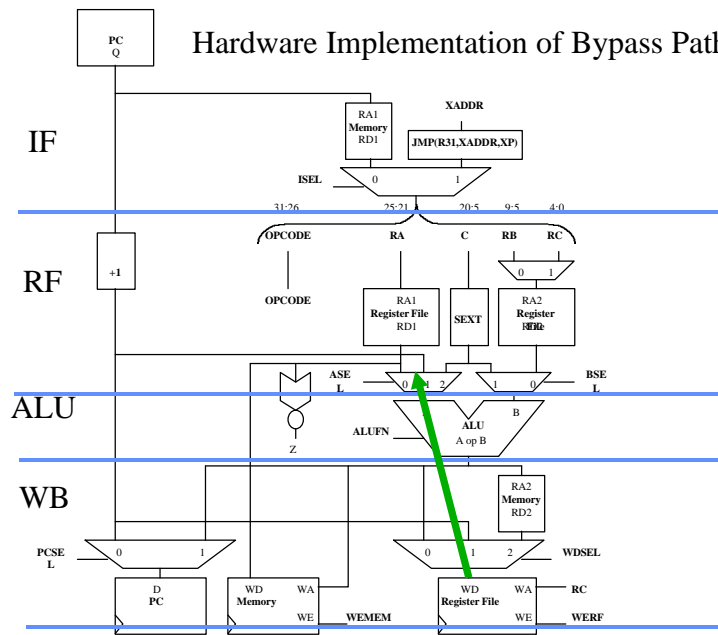


Bypass Paths - I



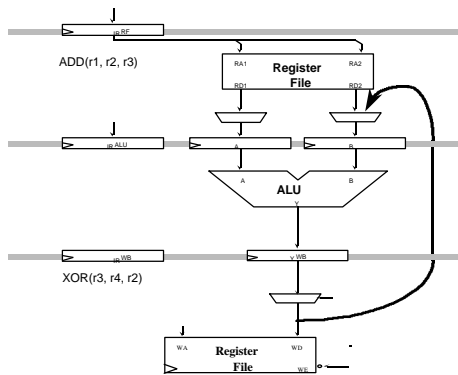
How Computers Work Lecture 13 Page 11

Hardware Implementation of Bypass Paths



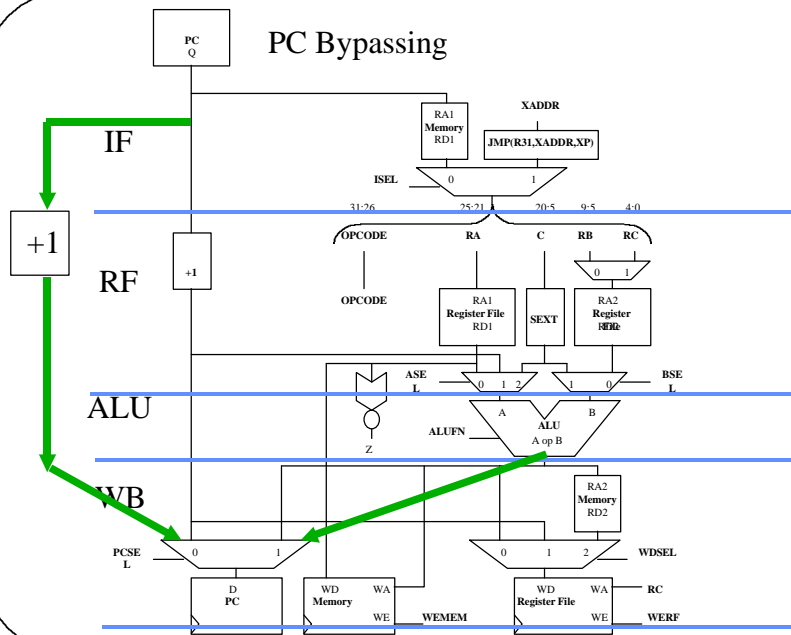
How Computers Work Lecture 13 Page 12

Bypass Paths - II



How Computers Work Lecture 13 Page 13

PC Bypassing



How Computers Work Lecture 13 Page 14

BRANCH DELAY SLOTS

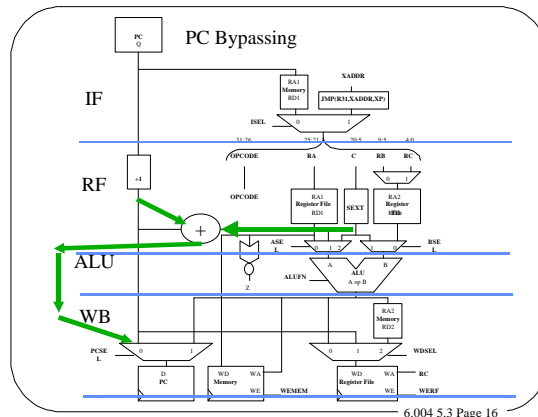
PROBLEM: One (or more) following instructions have been pre-fetched by the time a branch is taken.

POSSIBLE SOLUTIONS:

1. "Program around it". Either
 - 1a. Follow each BR with 2 NOP instructions; or
 - 1b. Make your compiler clever enough to move USEFUL instructions following branches.
2. Make pipeline "annul" instructions following branches which are taken, eg by disabling WERF and WEMEM and PCSEL.

How Computers Work Lecture 13 Page 15

Can we shorten the number of delay slots? A: Yes (by 1)



6.004 5.3 Page 16

How Computers Work Lecture 13 Page 16

Load Delays

Consider LOADS:

Can we fix all these problems using our previous bypass paths?

LD(r1, 0, r4)
 ADD(r1, r4, r5)
 XOR(r3, r4, r6)

ANSWER: No - only 1 (XOR) !

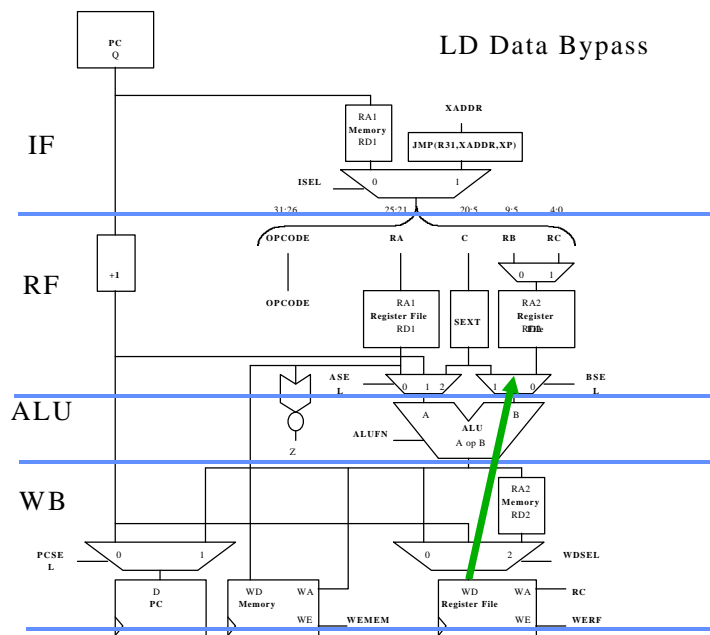
For a LD instruction fetched during clock *i*, data isn't returned from memory until late into cycle *i* + 3 !

LD(r1, 0, r4)	IF	RF	ALU	WB				
ADD(r1, r4, r5)		RF	RF	ALU	WB			
XOR(r3, r4, r6)		RF	RF	ALU	WB			

Note: In the original image, orange arrows labeled 'R4 Needed' point from the RF stage of ADD and XOR to the RF stage of LD. A green arrow labeled 'LD Data Available' points from the WB stage of LD to the ALU stage of ADD and XOR.

How Computers Work Lecture 13 Page 17

LD Data Bypass



How Computers Work Lecture 13 Page 18

Load Delays - II

Load Timing Problems:

LD(r1, 0, r4) Problem 1
ADD(r1, r4, r5) Problem 2
XOR(r3, r4, r6)

Can relegate both problems to Compiler.

Alternatively, fix Problem 2 using

Bypass Paths

and fix Problem 1 using

NOPs / Stalls

How Computers Work Lecture 13 Page 19

Load Problems - III

But, but, what about **FASTER** processors?

FACT: Processors will become fast relative to memories!

Do we just lengthen the cycle time?

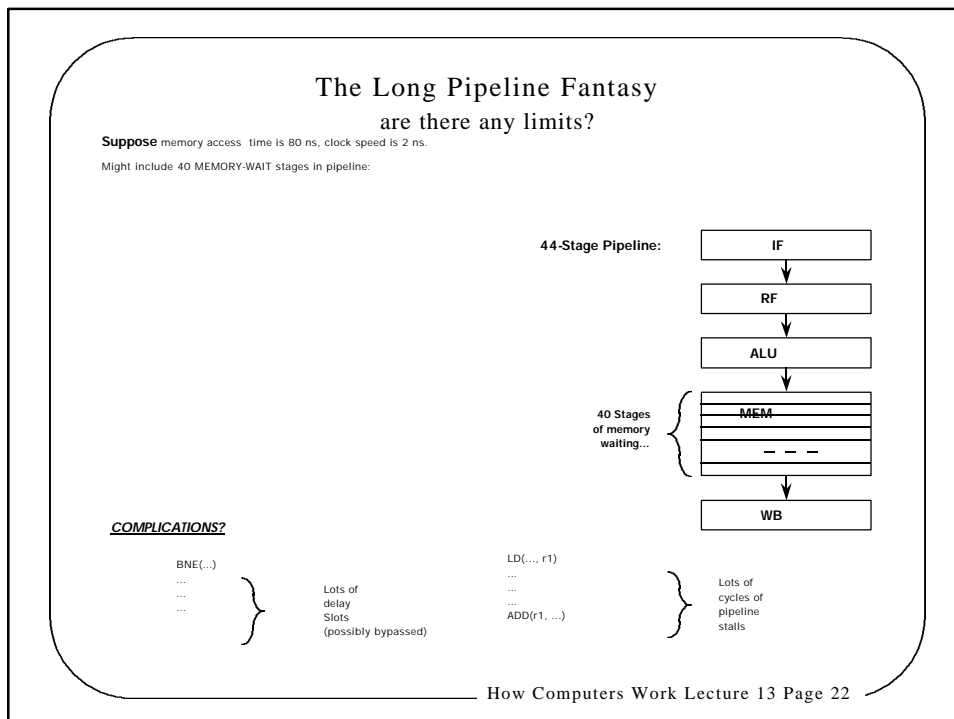
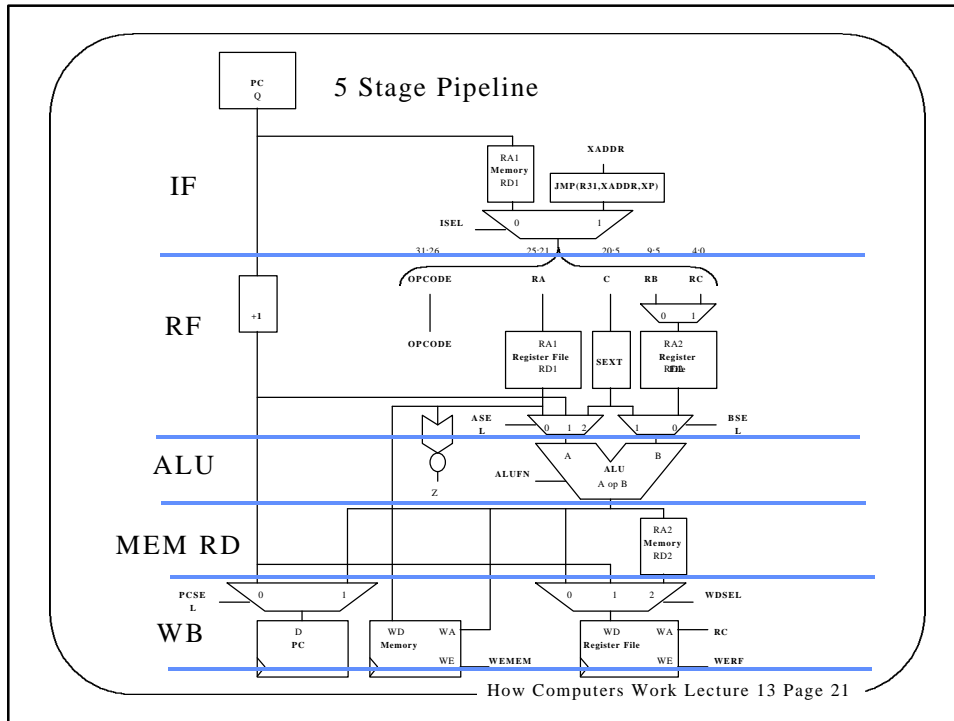
ALTERNATIVE: Longer pipelines.

1. Add "MEMORY WAIT" stages between START of read operation & return of data.
2. Build pipelined memories, so that multiple (say, N) memory transactions can be in progress at once.
3. (Optional). Stall pipeline when the N limit is exceeded.

4-Stage pipeline requires 1 instruction's delay.

5-Stage pipeline requires 2 instruction's delay.

How Computers Work Lecture 13 Page 20



What Have We Learned Today?

- Pipelining improves throughput by lowering clock period
- Pipelining cannot improve latency
- Data Hazards can be fixed with
 - Re-Programming, NOPs, Bypass Paths
- Branch Hazards can be fixed with
 - Re-Programming, NOPs, Annulment
- Memory Hazards can be fixed with
 - Re-Programming, NOPs, (1) Bypass Path
- As in Karate, balance is important ... too much pipelining is
BAD

How Computers Work Lecture 13 Page 23

Next Time:

- Implicit Multiple Issue
- Automatic Out-of-Order Execution

How Computers Work Lecture 13 Page 24