

Advanced Pipelining CS740

October 30, 2007

Topics

- Data Hazards
 - Stalling and Forwarding
 - Systematic testing of hazard-handling logic
- Control Hazards
 - Stalling, Predict not taken
- Exceptions
- Multicycle Instructions

Alpha ALU Instructions

RR-type instructions (addq, subq, xor, bis, cmplt): $rc \leftarrow ra \text{ funct } rb$

Op	ra	rb	000 0	funct	rc
31-26	25-21	20-16	15-13 12	11-5	4-0

RI-type instructions (addq, subq, xor, bis, cmplt): $rc \leftarrow ra \text{ funct } ib$

Op	ra	ib	1	funct	rc
31-26	25-21	20-13	12	11-5	4-0

Encoding

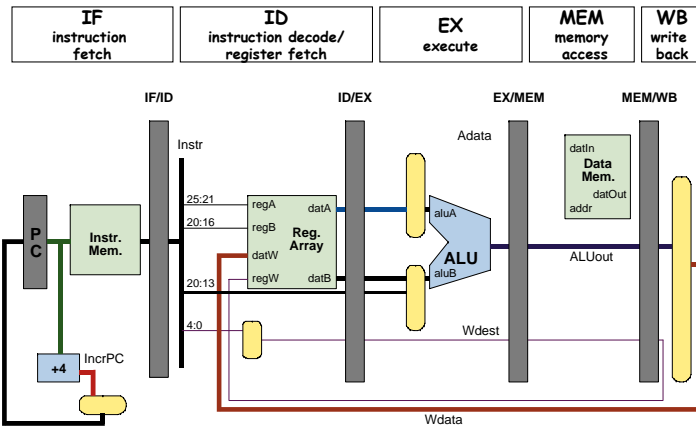
- ib is 8-bit unsigned literal

Operation	Op field	funct field
addq	0x10	0x20
subq	0x10	0x29
bis	0x11	0x20
xor	0x11	0x40
cmoveq	0x11	0x24
cmplt	0x11	0x4D

- 2 -

CS 740 F '07

Pipelined ALU Instruction Datapath




- 3 -

CS 740 F '07

Data Hazards in Alpha Pipeline

Problem

- Registers read in ID, and written in WB
- Must resolve conflict between instructions competing for register array
 - Generally do write back in first half of cycle, read in second
- But what about intervening instructions?
- E.g., suppose initially \$2 is zero:



- 4 -

CS 740 F '07


Handling Hazards by Stalling

Idea

- Delay instruction until hazard eliminated
- Put "bubble" into pipeline
- Dynamically generated NOP

Pipe Register Operation

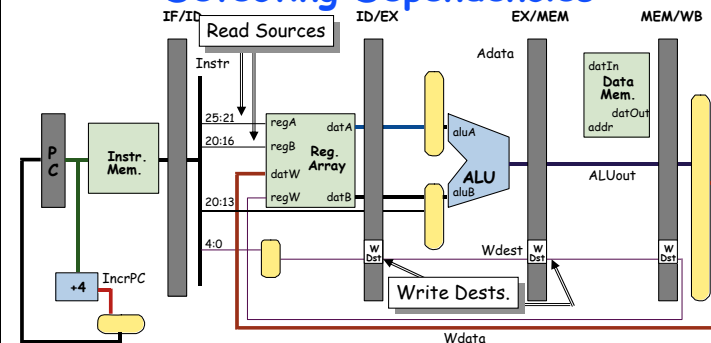
- "Transfer" (normal operation) indicates should transfer next state to current
- "Stall" indicates that current state should not be changed
- "Bubble" indicates that current state should be set to 0
 - Stage logic designed so that 0 is like NOP
 - [Other conventions possible]



- 5 -

CS 740 F'07

Detecting Dependencies



Pending Register Reads

- By instruction in ID
- ID_in.IR[25:21]: Operand A
- ID_in.IR[20:16]: Operand B
- Only for RR


Pending Register Writes

- EX_in.WDst: Destination register of instruction in EX
- MEM_in.WDst: Destination register of instruction in MEM

- 6 -

CS 740 F'07

Implementing Stalls



Stall Control Logic

- Determines which stages to stall, bubble, or transfer on next update

Rule:

- Stall in ID if either pending read matches either pending write
- Also stall IF; bubble EX

Effect

- Instructions with pending writes allowed to complete before instruction allowed out of ID


- 7 -

CS 740 F'07

Stalling for Data Hazards

Operation

- First instruction progresses unimpeded
- Second waits in ID until first hits WB
- Third waits in IF until second allowed to progress



- 8 -

CS 740 F'07

Observations on Stalling

Good

- Relatively simple hardware
- Only penalizes performance when hazard exists

Bad

- As if placed NOPs in code
 - Except that does not waste instruction memory

Reality

- Some problems can only be dealt with by stalling
 - Instruction cache miss
 - Data cache miss
- Otherwise, want technique with better performance

- 9 -

CS 740 F'07

Forwarding (Bypassing)

Observation

- ALU data generated at end of EX
 - Steps through pipe until WB
- ALU data consumed at beginning of EX

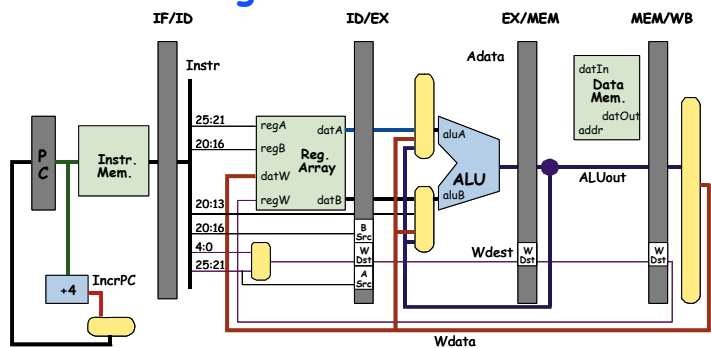
Idea

- Expedite passing of previous instruction result to ALU
 - By adding **extra data pathways and control**

- 10 -

CS 740 F'07

Forwarding for ALU Instructions



Operand Destinations

- ALU input A
 - Register EX_in.ASrc
- ALU input B
 - Register EX_in.BSrc

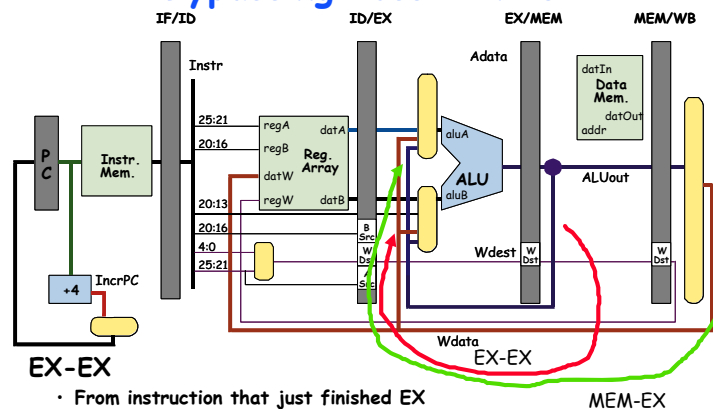
- 11 -

Operand Sources

- MEM_in.ALUout
 - Pending write to MEM_in.WDst
- WB_in.ALUout
 - Pending write to WB_in.WDst

CS 740 F'07

Bypassing Possibilities



EX-EX

- From instruction that just finished EX

MEM-EX

- From instruction that finished EX two cycles earlier


- 12 -

CS 740 F'07

Bypassing Data Hazards

Operation

- First instruction progresses down pipeline
- When in MEM, forward result to second instruction (in EX)
 - EX-EX forwarding
- When in WB, forward result to third instruction (in EX)
 - MEM-EX forwarding



- 13 -

CS 740 F'07

Load & Store Instructions

Load: $Ra \leftarrow Mem[Rb + offset]$

Op	ra	rb	offset
31-26	25-21	20-16	15-0

Store: $Mem[Rb + offset] \leftarrow Ra$

Op	ra	rb	offset
31-26	25-21	20-16	15-0

ID: Instruction decode/register fetch

- Store: $A \leftarrow Register[IR[25:21]]$
- B $\leftarrow Register[IR[20:16]]$

MEM: Memory

- Load: $Mem-Data \leftarrow DMemory[ALUOutput]$
- Store: $DMemory[ALUOutput] \leftarrow A$

WB: Write back

- Load: $Register[IR[25:21]] \leftarrow Mem-Data$

- 14 -

CS 740 F'07

Analysis of Data Transfers

Data Sources

- Available after EX
 - ALU Result Reg-Reg Result
- Available after MEM
 - Read Data Load result
 - ALU Data Reg-Reg Result passing through MEM stage

Data Destinations

- ALU A input Need in EX
 - Reg-Reg or Reg-Immediate Operand
- ALU B input Need in EX
 - Reg-Reg Operand
 - Load/Store Base
- Write Data Need in MEM
 - Store Data

- 15 -

CS 740 F'07

Some Hazards with Loads & Stores

Data Generated by Load

Load-Store Data

```
ldq $1, 8($2)
stq $1, 16($2)
```

Load-ALU

```
ldq $1, 8($2)
addq $2, $1, $2
```

Load-Store (or Load) Addr.

```
ldq $1, 8($2)
stq $2, 16($1)
```

Data Generated by Store

Store-Load Data

```
stq $1, 8($2)
ldq $3, 8($2)
```

Not a concern for us

Data Generated by ALU

ALU-Store (or Load) Addr

```
addq $1, $3, $2
stq $3, 8($2)
```

ALU-Store Data

```
addq $2, $3, $1
stq $1, 16($2)
```

- 16 -

CS 740 F'07


MEM-MEM Forwarding

Condition

- Data generated by load instruction
 - Register WB_in.WDst
- Used by immediately following store
 - Register MEM_in.ASrc

Load-Store Data

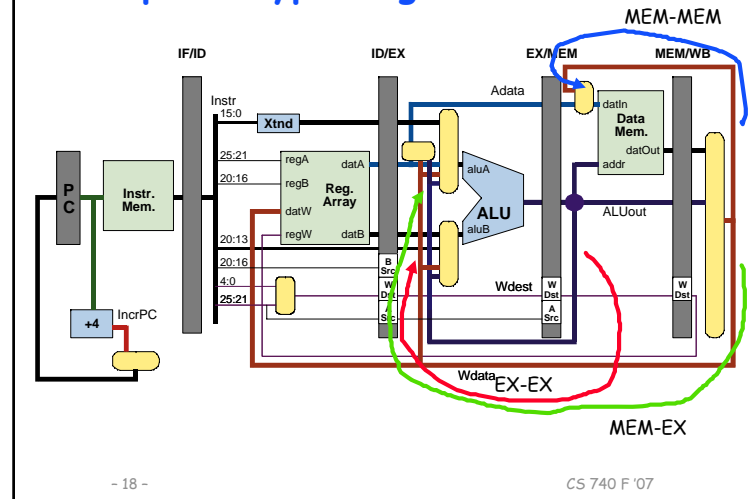
```
ldq $1, 8($2)
stq $1, 16($2)
```



- 17 -

CS 740 F'07

Complete Bypassing for ALU & L/S



- 18 -

CS 740 F'07

Impact of Forwarding

Single Remaining Unsolved Hazard Class


- Load followed by ALU operation
 - Including address calculation

Load-ALU

```
ldq $1, 8($2)
addq $2, $1, $2
```



With 1 Cycle Stall



- 19 -

CS 740 F'07

Methodology for characterizing and Enumerating Data Hazards

OP	writes	reads
RR	rc	ra, rb
RI	rc	ra
Load	ra	rb
Store		ra, rb

The space of data hazards (from a program-centric point of view) can be characterized by 3 independent axes:

3 possible write regs (axis 1):
RR.rc, RI.rc, Load.ra

6 possible read regs (axis 2):
RR.ra, RR.rb, RI.ra, Load.ra, Store.ra, Store.rb

A dependent read can be a distance of either 1 or 2 from the corresponding write (axis 3):

distance 2 hazard: RR.rc/RR.ra/2

distance 1 hazard: RR.rc/RR.rb/1

addq \$31, 63, \$2

addq \$31, \$2, \$3

addu \$2, \$31, \$4

- 20 -

CS 740 F'07

Enumerating data hazards

reads distance = 1

writes	RR,ra	RR,rb	RI,ra	L,rb	S,ra	S,rb
RR,rc						
RI,rc						
L,ra						

reads distance = 2

writes	RR,ra	RR,rb	RI,ra	L,rb	S,ra	S,rb
RR,rc						
RI,rc						
L,ra						

Testing Methodology

- 36 cases to cover all interactions between RR, RI, Load, & Store
- Would need to consider more read source and write destinations when add other instruction types

- 21 -

CS 740 F '07

Branch Instructions

Cond. Branch: $PC \leftarrow \text{Cond}(Ra) ? PC + 4 + \text{disp} * 4 : PC + 4$

Op	ra	disp
31-26	25-21	20-0

Sources

- PC, Ra

Destinations

- PC

Branch [Subroutine] (br, bsr): $Ra \leftarrow PC + 4; PC \leftarrow PC + 4 + \text{disp} * 4$

Op	ra	disp
31-26	25-21	20-0

Sources

- PC

Destinations

- PC, Ra

- 22 -

CS 740 F '07

New Data Hazards

Branch Uses Register Data

- Generated by ALU instruction
- Read from register in ID

Handling

- Same as other instructions with register data source
- Bypass
 - EX-EX
 - MEM-EX

ALU-Branch

```
addq $2, $3, $1
beq $1, targ
```

Distant ALU-Branch

```
addq $2, $3, $1
bis $31, $31, $31
beq $1, targ
```

Load-Branch

```
lw $1, 8($2)
beq $1, targ
```

- 23 -

CS 740 F '07

Jump Instructions

jmp, jsr, ret: $Ra \leftarrow PC + 4; PC \leftarrow Rb$

Op	ra	rb	Hint
31-26	25-21	20-16	15-0

Sources

- PC, Rb

Destinations


- PC, Ra

- 24 -

CS 740 F '07

Branch on equal

beq: $PC \leftarrow Ra == 0 ? PC + 4 + disp * 4 : PC + 4$



IF: Instruction fetch

- $IR \leftarrow IMemory[PC]$
- $incrPC \leftarrow PC + 4$

ID: Instruction decode/register fetch

- $A \leftarrow Register[IR[25:21]]$

Ex: Execute

- $Target \leftarrow incrPC + SignExtend(IR[20:0]) \ll 2$
- $Z \leftarrow (A == 0)$

MEM: Memory

- $PC \leftarrow Z ? Target : incrPC$

WB: Write back

- nop

- 29 -

CS 740 F '07

Branch Example

Desired Behavior

- Take branch at 0x00
- Execute target 0x18
 - $PC + 4 + disp \ll 2$
 - $PC = 0x00$
 - $disp = 5$

Displacement

Branch Code (demo08.O)

```

0x0: e7e00005 beq  r31, 0x18      # Take
0x4: 43e7f401 addq r31, 0x3f, r1  # (Skip)
0x8: 43e7f402 addq r31, 0x3f, r2  # (Skip)
0xc: 43e7f403 addq r31, 0x3f, r3  # (Skip)
0x10: 43e7f404 addq r31, 0x3f, r4 # (Skip)
0x14: 47ff041f bis  r31, r31, r31
0x18: 43e7f405 addq r31, 0x3f, r5 # (Target)
0x1c: 47ff041f bis  r31, r31, r31
0x20: 00000000 call_pal          halt
    
```

- 30 -

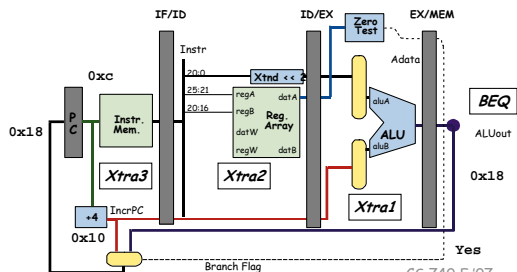
CS 740 F '07

Branch Hazard Example

```

0x0: beq  r31, 0x18      # Take
0x4: addq r31, 0x3f, r1  # Xtra1
0x8: addq r31, 0x3f, r2  # Xtra2
0xc: addq r31, 0x3f, r3  # Xtra3
0x10: addq r31, 0x3f, r4 # Xtra4
0x18: addq r31, 0x3f, r5 # Target
    
```

- With BEQ in Mem stage



- 31 -

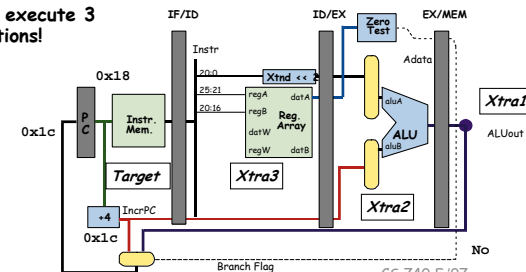
CS 740 F '07

Branch Hazard Example (cont.)

```

0x0: beq  r31, 0x18      # Take
0x4: addq r31, 0x3f, r1  # Xtra1
0x8: addq r31, 0x3f, r2  # Xtra2
0xc: addq r31, 0x3f, r3  # Xtra3
0x10: addq r31, 0x3f, r4 # Xtra4
0x18: addq r31, 0x3f, r5 # Target
    
```

- One cycle later
- Problem: Will execute 3 extra instructions!



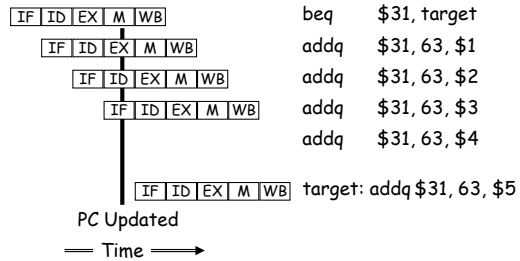
- 32 -

CS 740 F '07

Branch Hazard Pipeline Diagram

Problem

- Instruction fetched in IF, branch condition set in MEM




- 33 -

CS 740 F'07

Stall Until Resolve Branch

- Detect when branch in stages ID or EX
- Stop fetching until resolve
 - Stall IF. Inject bubble into ID



- 34 -

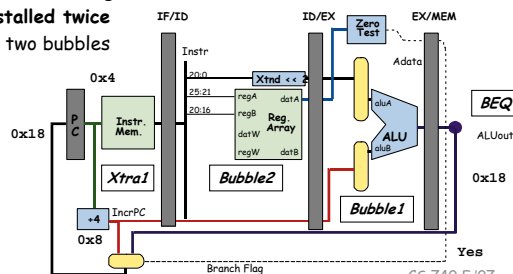
CS 740 F'07

Stalling Branch Example

```

0x0: beq    r31, 0x18    # Take
0x4: addq   r31, 0x3f, r1 # Xtra1
0x8: addq   r31, 0x3f, r2 # Xtra2
0xc: addq   r31, 0x3f, r3 # Xtra3
0x10: addq  r31, 0x3f, r4 # Xtra4
0x18: addq  r31, 0x3f, r5 # Target
    
```

- With BEQ in Mem stage
- Will have stalled twice
 - Injects two bubbles

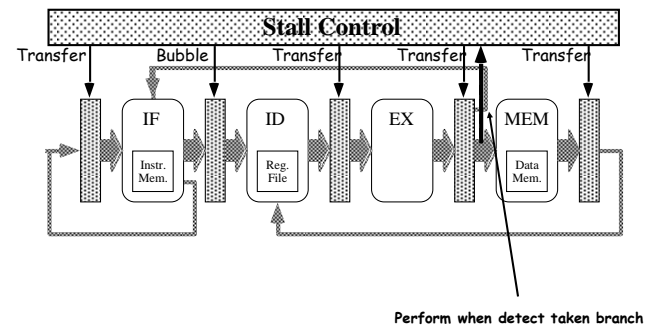


- 35 -

CS 740 F'07

Taken Branch Resolution

- When branch taken, still have instruction Xtra1 in pipe
- Need to flush it when detect taken branch in Mem
 - Convert it to bubble



- 36 -

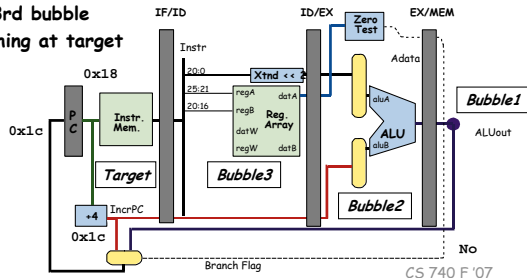
CS 740 F'07

Taken Branch Resolution Example

```

0x0: beq    r31, 0x18    # Take
0x4: addq   r31, 0x3f, r1 # Xtra1
0x8: addq   r31, 0x3f, r2 # Xtra2
0xc: addq   r31, 0x3f, r3 # Xtra3
0x10: addq  r31, 0x3f, r4 # Xtra4
0x18: addq  r31, 0x3f, r5 # Target
    
```

- When branch taken
- Generate 3rd bubble
- Begin fetching at target



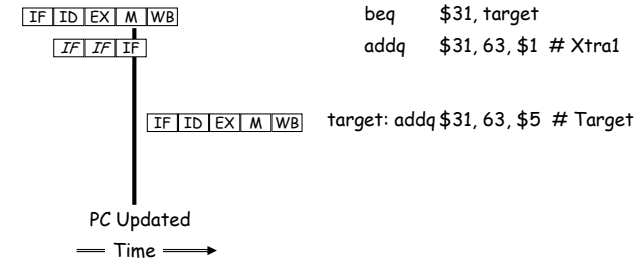
- 37 -

CS 740 F '07

Taken Branch Pipeline Diagram

Behavior

- Instruction Xtra1 held in IF for two extra cycles
- Then turn into bubble as enters ID




- 38 -

CS 740 F '07

Not Taken Branch Resolution

- [Stall two cycles with not-taken branches as well]
- When branch not taken, already have instruction Xtra1 in pipe
- Let it proceed as usual



- 39 -

CS 740 F '07

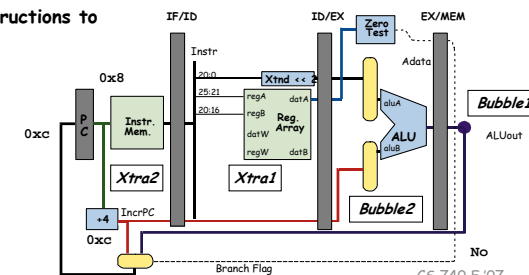
Not Taken Branch Resolution Example

demo09.O

```

0x0: bne    r31, 0x18    # Don't Take
0x4: addq   r31, 0x3f, r1 # Xtra1
0x8: addq   r31, 0x3f, r2 # Xtra2
0xc: addq   r31, 0x3f, r3 # Xtra3
0x10: addq  r31, 0x3f, r4 # Xtra4
    
```

- Branch not taken
- Allow instructions to proceed



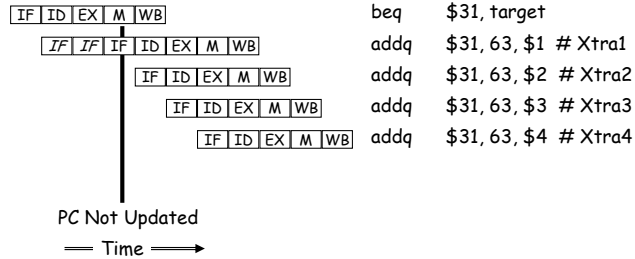
- 40 -

CS 740 F '07

Not Taken Branch Pipeline Diagram

Behavior

- Instruction Xtra1 held in IF for two extra cycles
- Then allowed to proceed



- 41 -

CS 740 F '07

Analysis of Stalling

Branch Instruction Timing

- 1 instruction cycle
- 3 extra cycles when taken
- 2 extra cycles when not taken

Performance Impact

- Branches 16% of instructions in SpecInt92 benchmarks
- 67% branches are taken
- Adds $0.16 * (0.67 * 3 + 0.33 * 2) = 0.43$ cycles to CPI
 - Average number of cycles per instruction
 - Serious performance impact

- 42 -


CS 740 F '07

Fetch & Cancel When Taken

- Instruction does not cause any updates until MEM or WB stages
- Instruction can be "cancelled" from pipe up through EX stage
 - Replace with bubble

Strategy

- Continue fetching under assumption that branch not taken
- If decide to take branch, cancel undesired ones



- 43 -

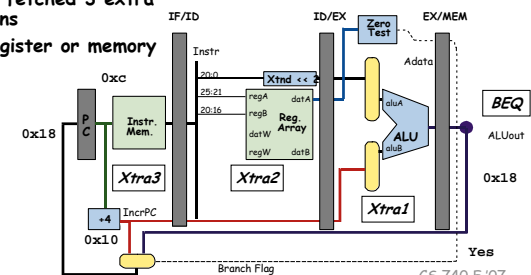
CS 740 F '07

Canceling Branch Example

```

0x0: beq    r31, 0x18    # Take
0x4: addq   r31, 0x3f, r1 # Xtra1
0x8: addq   r31, 0x3f, r2 # Xtra2
0xc: addq   r31, 0x3f, r3 # Xtra3
0x10: addq  r31, 0x3f, r4 # Xtra4
0x18: addq  r31, 0x3f, r5 # Target
  
```

- With BEQ in Mem stage
- Will have fetched 3 extra instructions
- But no register or memory updates



- 44 -

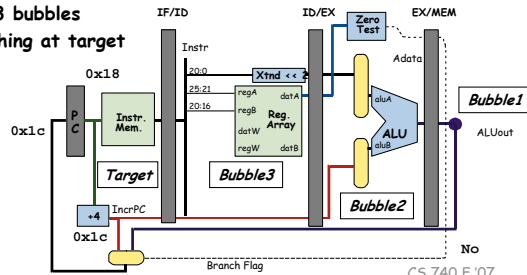
CS 740 F '07

Canceling Branch Resolution Example

```

0x0: beq    r31, 0x18    # Take
0x4: addq   r31, 0x3f, r1 # Xtra1
0x8: addq   r31, 0x3f, r2 # Xtra2
0xc: addq   r31, 0x3f, r3 # Xtra3
0x10: addq  r31, 0x3f, r4 # Xtra4
0x18: addq  r31, 0x3f, r5 # Target
    
```

- When branch taken
- Generate 3 bubbles
- Begin fetching at target



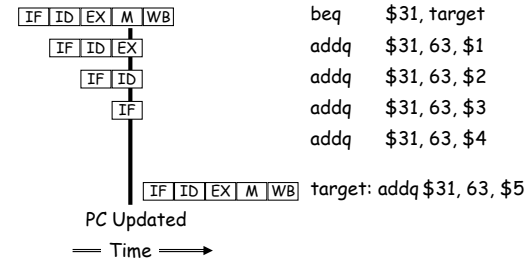
- 45 -

CS 740 F '07

Canceling Branch Pipeline Diagram

Operation

- Process instructions assuming branch will not be taken
- When *is* taken, cancel 3 following instructions



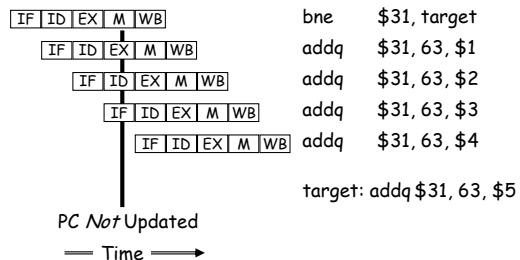
- 46 -

CS 740 F '07

Noncanceling Branch Pipeline Diagram

Operation

- Process instructions assuming branch will not be taken
- If really isn't taken, then instructions flow unimpeded



- 47 -

CS 740 F '07

Branch Prediction Analysis

Our Scheme Implements "Predict Not Taken"

- But 67% of branches are taken
- Impact on CPI: $0.16 * 0.67 * 3.0 = 0.32$
 - Still not very good

Alternative Schemes


- Predict taken
 - Would be hard to squeeze into our pipeline
 - » Can't compute target until ID
- Backwards taken, forwards not taken
 - Predict based on sign of displacement
 - Exploits fact that loops usually closed with backward branches

- 48 -

CS 740 F '07

Exceptions


An *exception* is a transfer of control to the OS in response to some *event* (i.e. change in processor state)



- 49 -

CS 740 F '07

Issues with Exceptions



A1: What kinds of events can cause an exception?

A2: When does the exception occur?

B1: How does the handler determine the location and cause of the exception?

B2: Are exceptions allowed within exception handlers?

C1: Can the user process restart?

C2: If so, where?

- 50 -

CS 740 F '07

Internal (CPU) Exceptions

Internal exceptions occur as a result of events generated by executing instructions.

Execution of a `CALL_PAL` instruction.

- allows a program to transfer control to the OS

Errors during instruction execution

- arithmetic overflow, address error, parity error, undefined instruction

Events that require OS intervention

- virtual memory page fault

- 51 -

CS 740 F '07

External (I/O) exceptions

External exceptions occur as a result of events generated by devices external to the processor.

I/O interrupts

- hitting `^C` at the keyboard
- arrival of a packet
- arrival of a disk sector

Hard reset interrupt

- hitting the reset button

Soft reset interrupt

- hitting `ctl-alt-delete` on a PC

- 52 -

CS 740 F '07

Exception handling (hardware tasks)


Recognize event(s)

Associate one event with one instruction.

- external event: pick any instruction
- multiple internal events: typically choose the earliest instruction.
- multiple external events: prioritize
- multiple internal and external events: prioritize

Create Clean Break in Instruction Stream

- Complete all instructions before excepting instruction
- Abort excepting and all following instructions
 - this clean break is called a "precise exception"



- 53 -

CS 740 F '07

Exception handling (hardware tasks)

Set status registers

- **Exception Address:** the EXC_ADDR register
 - external exception: address of instruction about to be executed
 - internal exception: address of instruction causing the exception
 - » except for arithmetic exceptions, where it is the following instruction
- **Cause of the Exception:** the EXC_SUM and FPCR registers
 - was the exception due to division by zero, integer overflow, etc.
- **Others**
 - which ones get set depends on CPU and exception type

Disable interrupts and switch to kernel mode

Jump to common exception handler location

- 54 -

CS 740 F '07

Exception handling (software tasks)

Deal with event

(Optionally) resume execution

- using special REI (return from exception or interrupt) instruction
- similar to a procedure return, but restores processor to user mode as a side effect.

Where to resume execution?

- usually re-execute the instruction causing exception

- 55 -

CS 740 F '07

Precise vs. Imprecise Exceptions

In the Alpha architecture:


- arithmetic exceptions may be *imprecise* (similar to the CRAY-1)
 - motivation: simplifies pipeline design, helping to increase performance
- all other exceptions are precise

Imprecise exceptions:

- all instructions before the excepting instruction complete
- the excepting instruction and instructions after it may or may not complete

What if precise exceptions are needed?

- insert a TRAPB (trap barrier) instruction immediately after
 - stalls until certain that no earlier insts take exceptions



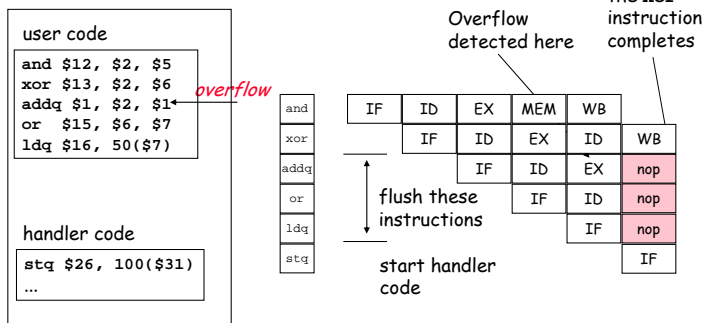
In the remainder of our discussion, assume for the sake of simplicity that all Alpha exceptions are precise.

- 56 -

CS 740 F '07

Example: Integer Overflow

(This example illustrates a *precise* version of the exception.)



- 57 -

CS 740 F'07

Multicycle instructions

Alpha 21264 Execution Times:

• Measured in clock cycles

Operation	Integer	FP-Single	FP-Double
add / sub	1	4	4
multiply	8-16	4	4
divide	N / A	10	23


H&P Dynamic Instruction Counts:

Operation	Integer	FP Benchmarks	
	Benchmarks	Integer	FP
add / sub	14%	11%	14%
multiply	< 0.1%	< 0.1%	13%
divide	< 0.1%	< 0.1%	1%

- 58 -

CS 740 F'07

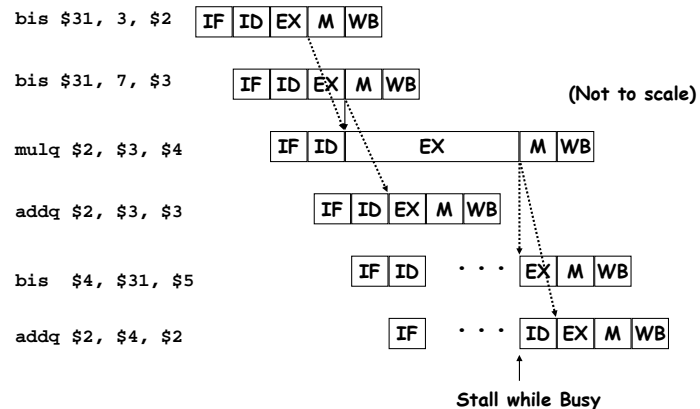
Pipeline Revisited



- 59 -

CS 740 F'07

Multiply Timing Example



- 60 -

CS 740 F'07

Conclusion

Pipeline Characteristics for Multi-cycle Instructions

- **In-order issue**
 - Instructions fetched and decoded in program order
- **Out-of-order completion**
 - Slow instructions may complete after ones that are later in program order

Performance Opportunities

- Transformations such as loop unrolling & software pipelining to expose potential parallelism
- Schedule code to use multiple functional units
 - Must understand idiosyncrasies of pipeline structure