

Chapter 4B


The Processor

Control Hazards

- Branch determines flow of control
 - Fetching next instruction depends on branch outcome
 - Pipeline can't always fetch correct instruction
 - Next: still at ID stage of branch before calculating the branch target address
- In MIPS pipeline
 - Need to compare registers and compute target early in the pipeline
 - Add hardware to do it in ID stage


Branch Instruction

When branch decision is made in MEM,




Branch Instruction

When branch decision is made in EXE,




Stall for a branch (2 cycles)




Branch Instruction

When branch decision is made in ID,



Stall for a branch (1 cycles)




Control Hazards

- When the flow of instruction addresses is not sequential (i.e., $PC = PC + 4$); incurred by change of flow instructions
 - Conditional branches (beq, bne)
 - Unconditional branches (j, jal, jr)
 - Exceptions
- Possible approaches
 - Stall (impacts CPI)
 - Delayed branch (requires compiler support)
 - Move decision point as early in the pipeline as possible, thereby reducing the number of stall cycles
 - Predict and hope for the best !
- Control hazards occur less frequently than data hazards, but there is *nothing* as effective against control hazards as forwarding is for data hazards

Stall on Branch

- Wait until branch outcome determined before fetching next instruction



Two “Types” of Stalls

- **Bubble**: **nop** instruction (or bubble) inserted between two instructions in the pipeline (as done for **load-use** situations)
- **Flush** (or instruction **squashing**) : an instruction in the pipeline is **replaced** with a nop instruction (as done for instructions located **sequentially after j instructions**)
 - Zero the control bits for the instruction to be flushed

Exceptions

- Detect Exception
 - Capture exception PC, exception cause
- Flush pipeline
- Begin fetching at new PC

Delayed Branch

- If the branch hardware has been moved to the ID stage, then we can eliminate all branch stalls with **delayed branches** which are defined as **always executing the next sequential instruction after the branch instruction** – the branch takes effect **after** that next instruction
 - MIPS compiler moves an instruction to immediately after the branch that is not affected by the branch (a **safe** instruction) thereby **hiding** the branch delay
- With deeper pipelines, the branch delay grows requiring more than one delay slot
 - Delayed branches have lost popularity compared to more expensive but more flexible (dynamic) hardware branch prediction
 - Growth in available transistors has made hardware branch prediction relatively cheaper

Branch Delay Slots

- Since we need to have a dead cycle anyway, let's put a useful instruction there
- Advantage:
 - Do more useful work
 - Potentially get rid of all stalls
- Disadvantage:
 - Exposes microarchitecture to ISA
 - Deeper pipelines require more delay slots

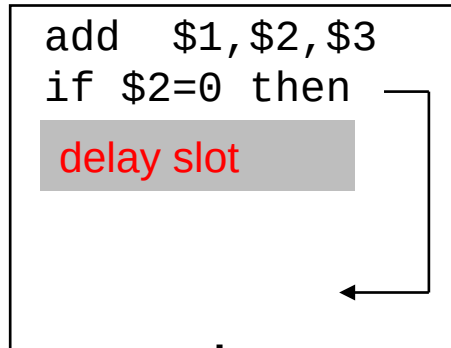
```
ADD $t2,$t3,$t4  
BNEZ $t5,_loop  
NOP
```



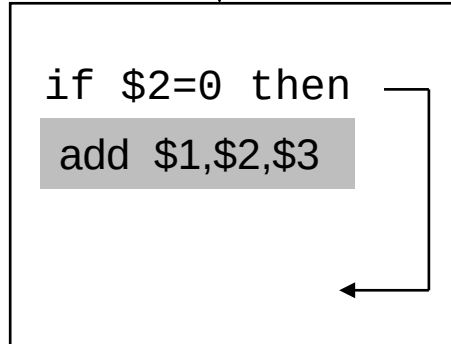
```
BNEZ $t5,_loop  
ADD $t2,$t3,$t4
```

Scheduling Branch Delay Slots

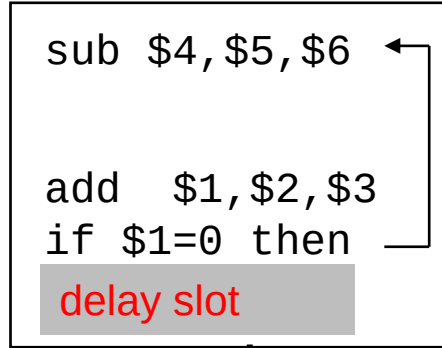
A. From before branch



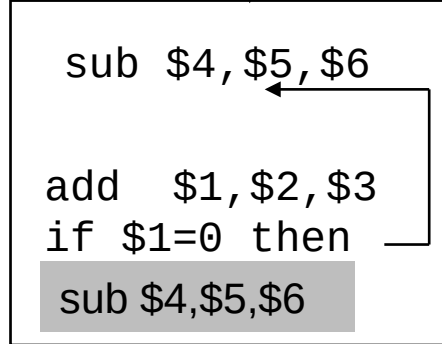
becomes



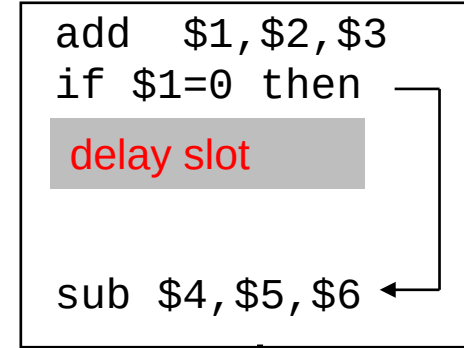
B. From branch target



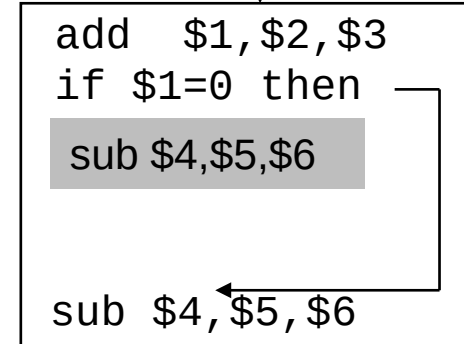
becomes



C. From fall through



becomes



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

Branch Delay Slots: A Solution?

- Not really
 - Exposes stuff to the ISA that is better kept hidden
 - Not scalable as the pipeline changes

Branch Prediction




- Longer pipelines can't readily determine branch outcome early
 - Stall penalty becomes unacceptable
- Predict outcome of branch
 - No stall if prediction is correct
 - Only stall if prediction is wrong
- In MIPS pipeline
 - Can predict branches **not taken**
 - Fetch instruction after branch, with no delay

Static Branch Prediction


- Resolve branch hazards by assuming a given outcome and proceeding without waiting to see the actual branch outcome
- **Predict not taken** – always predict branches will not be taken, continue to fetch from the sequential instruction stream, **only when branch is taken does the pipeline stall**
 - If taken, **flush** instructions after the branch (earlier in the pipeline)
 - in IF and ID stages if branch logic in EX – **two stalls**
 - in IF stage if branch logic in ID – **one stall**
 - ensure that those flushed instructions haven't changed the machine state – automatic in the MIPS pipeline since machine state changing operations are at the tail end of the pipeline (MemWrite (in MEM) or RegWrite (in WB))
 - restart the pipeline at the branch destination

MIPS with Predict Not Taken


Prediction correct



Prediction incorrect



Flushing if Taken



- To flush the IF stage instruction, assert `IF.Flush` to zero the instruction field of the IF/ID pipeline register (transforming it into a nop)

Branching Structures

- Predict not taken works well for “top of the loop” branching structures

- But such loops have jumps at the bottom of the loop to return to the top of the loop - and incur the jump stall overhead

```
Loop: beq $1,$2,Out
      1st loop instr
      .
      .
      .
      last loop instr
      j Loop
Out:  fall out instr
```

- Predict not taken doesn't work well for “bottom of the loop” branching structures

```
Loop: 1st loop instr
      2nd loop instr
      .
      .
      .
      last loop instr
      bne $1,$2,Loop
      fall out instr
```

More-Realistic Branch Prediction

- Static branch prediction
 - Based on typical branch behavior
 - Example: **loop** and if-statement branches
 - Predict backward branches taken
 - Predict forward branches not taken
- Dynamic branch prediction
 - Hardware measures actual branch behavior
 - e.g., record recent history of each branch
 - Assume future behavior will continue the trend
 - When wrong, stall while re-fetching, and update history

Will be explained later again!

Pipeline Summary

The BIG Picture

- Pipelining improves performance by increasing instruction throughput
 - Executes multiple instructions in parallel
 - Each instruction has the same latency
- Subject to hazards
 - Structure, data, control
- Instruction set design affects complexity of pipeline implementation


MIPS Pipelined Datapath

IF: Instruction fetch

ID: Instruction decode/
register file readEX: Execute/
address calculation


MEM: Memory access

WB: Write back



Pipeline registers


- Need registers between stages
 - To hold information produced in previous cycle




Pipeline Operation

- Cycle-by-cycle flow of instructions through the pipelined datapath
 - “Single-clock-cycle” pipeline diagram
 - Shows pipeline usage in a single cycle
 - Highlight resources used
 - c.f. “multi-clock-cycle” diagram
 - Graph of operation over time
- We’ll look at “single-clock-cycle” diagrams for **load & store**


IF for Load, Store, ...




ID for Load, Store, ...




EX for Load




MEM for Load




WB for Load




Corrected Datapath for Load




EX for Store



MEM for Store




WB for Store





Multi-Cycle Pipeline Diagram

■ Traditional form




Single-Cycle Pipeline Diagram

■ State of pipeline in a given cycle




Another Pipeline Diagram

- Form showing resource usage




Pipelined Control (Simplified)



- Control signals derived from instruction
 - As in single-cycle implementation




Pipelined Control



Data Hazards in ALU Instructions

- Consider this sequence:
 sub \$2, \$1, \$3
 and \$12, \$2, \$5
 or \$13, \$6, \$2
 add \$14, \$2, \$2
 sw \$15, 100(\$2)
- We can resolve hazards with forwarding
 - How do we detect when to forward?

Dependencies & Forwarding



Detecting the Need to Forward


- Pass register numbers along pipeline
 - e.g., ID/EX.RegisterRs = register number for Rs sitting in ID/EX pipeline register
- ALU operand register numbers in EX stage are given by
 - ID/EX.RegisterRs, ID/EX.RegisterRt
- Data hazards when
 - 1a. EX/MEM.RegisterRd = ID/EX.RegisterRs
 - 1b. EX/MEM.RegisterRd = ID/EX.RegisterRt
 - 2a. MEM/WB.RegisterRd = ID/EX.RegisterRs
 - 2b. MEM/WB.RegisterRd = ID/EX.RegisterRt

Fwd from
EX/MEM
pipeline reg

Fwd from
MEM/WB
pipeline reg

Another Pipeline Diagram


- Form showing resource usage



Detecting the Need to Forward

- But only if forwarding instruction will write to a register!
 - EX/MEM.RegWrite, MEM/WB.RegWrite
- And **only if Rd for that instruction is not \$zero**
 - EX/MEM.RegisterRd $\neq$ 0,
MEM/WB.RegisterRd $\neq$ 0

Forwarding Paths



b. With forwarding

Forwarding Conditions

■ EX hazard

- if (EX/MEM.RegWrite and (EX/MEM.RegisterRd $\neq$ 0)
and (EX/MEM.RegisterRd = ID/EX.RegisterRs))

ForwardA = 10

- if (EX/MEM.RegWrite and (EX/MEM.RegisterRd $\neq$ 0)
and (EX/MEM.RegisterRd = ID/EX.RegisterRt))

ForwardB = 10

■ MEM hazard

- if (MEM/WB.RegWrite and (MEM/WB.RegisterRd $\neq$ 0)
and (MEM/WB.RegisterRd = ID/EX.RegisterRs))

ForwardA = 01

- if (MEM/WB.RegWrite and (MEM/WB.RegisterRd $\neq$ 0)
and (MEM/WB.RegisterRd = ID/EX.RegisterRt))

ForwardB = 01


Double Data Hazard

- Consider the sequence:
 - add **\$1**, \$1, \$2
 - add **\$1**, **\$1**, \$3
 - add \$1, **\$1**, \$4
- Both hazards occur
 - Want to use the most recent
 - EX hazard condition does not need to revise
- Revise MEM hazard condition
 - Only fwd if EX hazard condition isn't true

Revised Forwarding Condition

- MEM hazard
 - if (MEM/WB.RegWrite and (MEM/WB.RegisterRd $\neq$ 0)
and not (EX/MEM.RegWrite and (EX/MEM.RegisterRd $\neq$ 0)
and (EX/MEM.RegisterRd = ID/EX.RegisterRs))
and (MEM/WB.RegisterRd = ID/EX.RegisterRs))
ForwardA = 01
 - if (MEM/WB.RegWrite and (MEM/WB.RegisterRd $\neq$ 0)
and not (EX/MEM.RegWrite and (EX/MEM.RegisterRd $\neq$ 0)
and (EX/MEM.RegisterRd = ID/EX.RegisterRt))
and (MEM/WB.RegisterRd = ID/EX.RegisterRt))
ForwardB = 01


Datapath with Forwarding



Control for Forwarding Muxes


Mux control	Source	Explanation
ForwardA = 00	ID/EX	The first ALU operand comes from the register file.
ForwardA = 10	EX/MEM	The first ALU operand is forwarded from the prior ALU result.
ForwardA = 01	MEM/WB	The first ALU operand is forwarded from data memory or an earlier ALU result.
ForwardB = 00	ID/EX	The second ALU operand comes from the register file.
ForwardB = 10	EX/MEM	The second ALU operand is forwarded from the prior ALU result.
ForwardB = 01	MEM/WB	The second ALU operand is forwarded from data memory or an earlier ALU result.

Load-Use Data Hazard



Another Pipeline Diagram

- Form showing resource usage




Load-Use Hazard Detection

- Check when using instruction is decoded in ID stage
- ALU operand register numbers in ID stage are given by
 - IF/ID.RegisterRs, IF/ID.RegisterRt
- Load-use hazard when
 - ID/EX.MemRead and
 - ((ID/EX.RegisterRt = IF/ID.RegisterRs) or (ID/EX.RegisterRt = IF/ID.RegisterRt))
- If detected, stall and insert bubble


How to Stall the Pipeline

- Force control values in ID/EX register to 0
 - EX, MEM and WB do nop (no-operation)
- Prevent update of PC and IF/ID register
 - Using instruction is decoded again
IF/IDWrite=0
 - Following instruction is fetched again
PCWrite=0
 - 1-cycle stall allows MEM to read data for lw
 - Can subsequently forward to EX stage

Datapath with Hazard Detection




Stall/Bubble in the Pipeline




Another Pipeline Diagram

- Form showing resource usage



Another Pipeline Diagram

- Form showing resource usage



Stalls and Performance

The BIG Picture

- Stalls reduce performance
 - But are required to get correct results
- Compiler can arrange code to avoid hazards and stalls
 - Requires knowledge of the pipeline structure