

U. Wisconsin CS/ECE 552
Introduction to Computer Architecture

Prof. Karu Sankaralingam

Pipelining (Chapter 6)

www.cs.wisc.edu/~karu/courses/cs552

Slides combined and enhanced by Karu Sankaralingam from work by Falsafi, Hill, Marculescu, Nagle, Patterson, Roth, Rutenbar, Schmidt, Shen, Sohi, Sorin, Thottethodi, Vijaykumar, & Wood

Outline

- Pipelining
 - What? Basic concepts
 - Overlapping execution
 - Latency vs. throughput
 - Why? Performance implications
 - Speedup
 - CPI, cycletime
 - How? Implementation challenges


CS/ECE 552

(2)

Sankaralingam

Pipelining

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

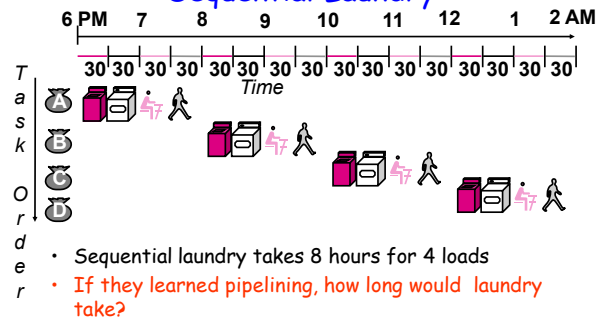


CS/ECE 552

(3)

Sankaralingam

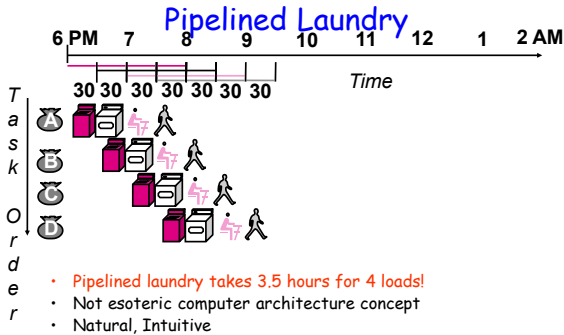
Sequential Laundry



CS/ECE 552

(4)

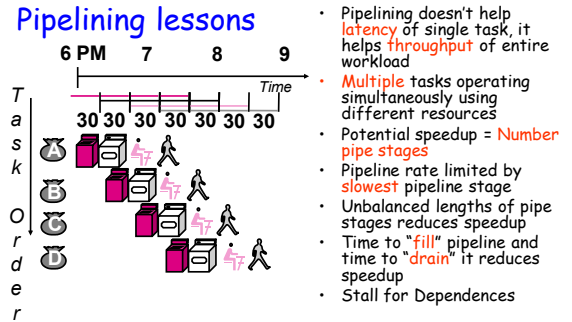
Sankaralingam



CS/ECE 552

(5)

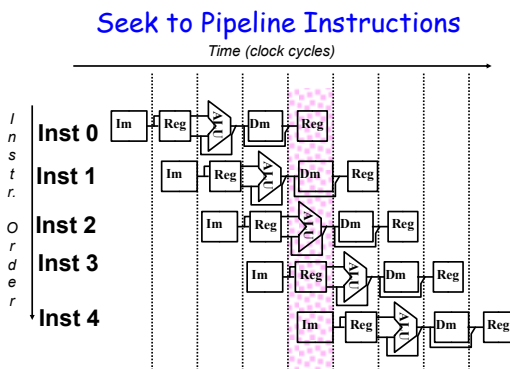
Sankaralingam



CS/ECE 552

(6)

Sankaralingam




CS/ECE 552

(7)

Sankaralingam

The Stages of Load




- Ifetch: Instruction Fetch
 - Fetch the instruction from the Instruction Memory
- Reg/Dec: Registers Fetch and Instruction Decode
- Exec: Calculate the memory address
- Mem: Read the data from the Data Memory
- Wr: Write the data back to the register file

CS/ECE 552

(8)

Sankaralingam

Pipelined Execution Representation



- Ideal speedup = 5, but ...

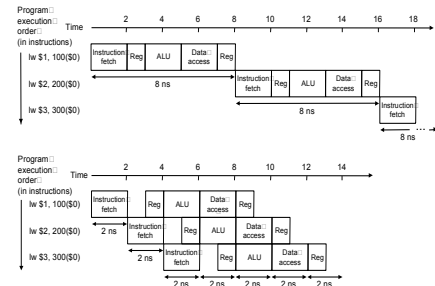
CS/ECE 552

(9)

Sankaralingam

Pipelining

- Improve performance by increasing instruction throughput



Ideal speedup is number of stages in the pipeline. Do we achieve this?

CS/ECE 552

(10)

Sankaralingam

Non-uniform stages

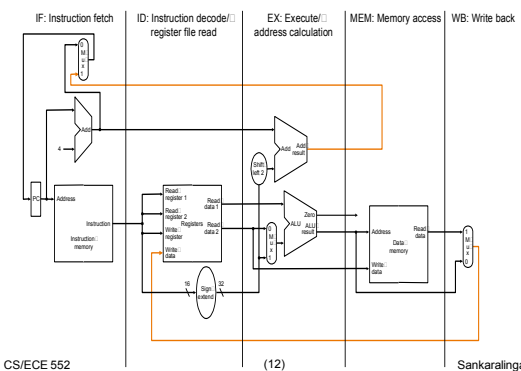
$$\text{Maximum Speedup} \leq \frac{\text{Number of stages}}{\text{Time for longest stage}}$$

CS/ECE 552

(11)

Sankaralingam

Pipeline Forecast: Single-Cycle Datapath




CS/ECE 552

(12)

Sankaralingam

Pipeline Forecast: Pipelined Datapath



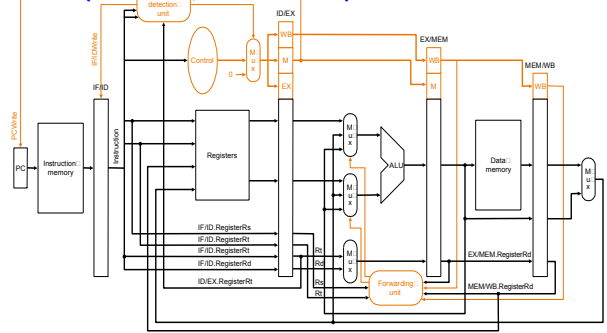
- Pipeline datapath with registers

CS/ECE 552

(13)

Sankaralingam

Pipeline Forecast: Pipelined Control



CS/ECE 552

(14)

Sankaralingam

Pipeline Forecast: Big Picture

- Datapath similar to single-cycle datapath
- Partition datapath with **pipeline latches** (D-FFs)
- Naïve Control
 - Generate single-cycle control signals
 - Pass control signals through pipeline latches
 - Apply control signals at appropriate stage/cycle
- Truth is more complex (instruction interact)

CS/ECE 552

(15)

Sankaralingam

Instructions vs. Laundry

- Definitely:
 - Need to maintain "illusion" of sequential execution
 - Execution is actually overlapped.
- Pipeline Hazards
 - **structural hazards**: attempt to use the same resource two different ways at the same time
 - E.g., combined washer/dryer would be a structural hazard or folder busy doing something else (watching TV)
 - **data hazards**: attempt to use item before it is ready
 - E.g., one sock of pair in dryer and one in washer; can't fold until get sock from washer through dryer
 - instruction depends on result of prior instruction still in the pipeline
 - **control hazards**: attempt to make a decision before condition is evaluated
 - E.g., washing football uniforms and need to get proper detergent level; need to see after dryer before next load in
 - branch instructions

CS/ECE 552

(16)

Sankaralingam

Hazards

- Structural hazards
 - Two instructions need the same hardware
- Data Hazards
 - Data not ready
- Control Hazards
 - Which instruction to fetch? Not known.

CS/ECE 552

(17)

Sankaralingam

Hazards

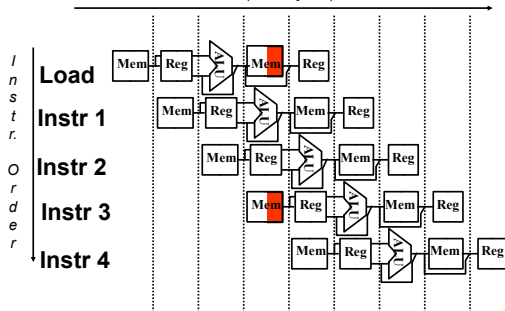
- Can always resolve hazards by **waiting**
 - pipeline control must detect the hazard
 - take action (or delay action) to resolve hazards
- Delays
 - Pipeline stalls/bubbles
 - Reduce speedup

CS/ECE 552

(18)

Sankaralingam

Single Memory: Structural Hazard



CS/ECE 552

(19)

Sankaralingam

Structural Hazards

- If 1.3 memory accesses per instruction
 - How?
 - 1 per instruction for instruction fetch
 - Fraction for data load/store
 - Depends on instruction mix
 - 20% load + 10% store
 - 15% load + 15% store
- CPI is atleast 1.3 (otherwise memory is used more than 100%)

CS/ECE 552

(20)

Sankaralingam

Data Hazards

```

add r1, r2, r3
sub r4, r1, r3
and r6, r1, r7
or r8, r1, r9
xor r10, r1, r11
    
```

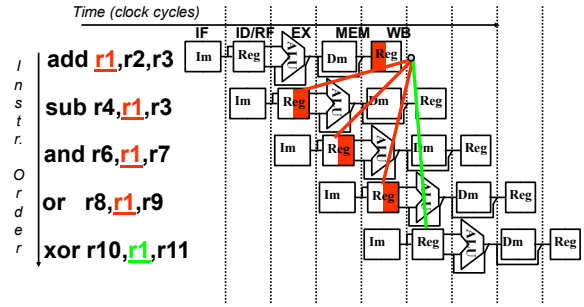
CS/ECE 552

(21)

Sankaralingam

Hazards on r1

- Dependencies backwards in time

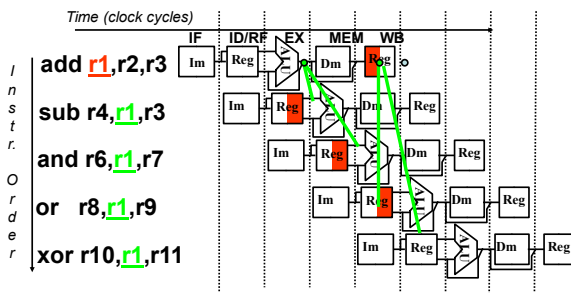


CS/ECE 552

(22)

Sankaralingam

Data Hazard Solution




CS/ECE 552

(23)

Sankaralingam

Forwarding (a.k.a. bypassing)



- Can't solve with forwarding:
 - Must delay/stall instruction dependent on loads


CS/ECE 552

(24)

Sankaralingam

Control Hazard: Solutions

- Stall: wait until decision is clear
 - Its possible to move up decision to 2nd stage by adding hardware to check registers as being read



- Impact: 2 clock cycles per branch instruction
=> slow


CS/ECE 552

(25)

Sankaralingam

Control Hazard: Solutions

- Predict: guess one direction then back up if wrong
 - Predict not taken



- Impact: 1 clock cycles per branch instruction if right, 2 if wrong (right - 50% of time)
- More dynamic scheme: history of 1 branch (- 90%)


CS/ECE 552

(26)

Sankaralingam

Control Hazard: Solutions

- Redefine branch behavior (takes place after next instruction)
"delayed branch"



- Impact: 0 clock cycles per branch instruction if can find instruction to put in "slot" (- 50% of time)
- As launch more instruction per clock cycle, less useful

CS/ECE 552

(27)

Sankaralingam

STOPPED HERE

CS/ECE 552

(28)

Sankaralingam

Pipelined Processor Design

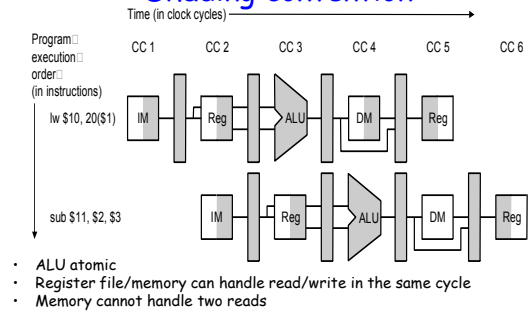
- Designing a pipelined processor
 - Associate resources with states
 - Resources not necessarily atomic
 - Register reads and writes can happen in the same cycle
 - Writes in first half of cycle and reads in the second half
 - Assert appropriate controls in each stage
 - Make sure all necessary information is carried through inter-pipestage flip-flops

CS/ECE 552

(29)

Sankaralingam

Shading convention



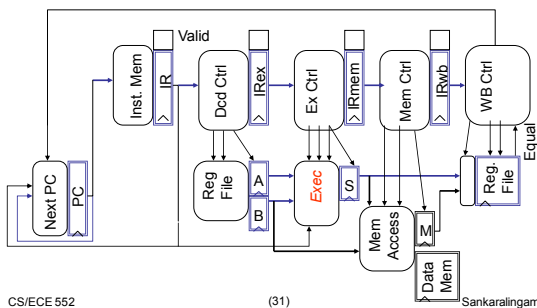
CS/ECE 552

(30)

Sankaralingam

Pipelined Processor (slide version)

- What happens if we start a new instruction every cycle?

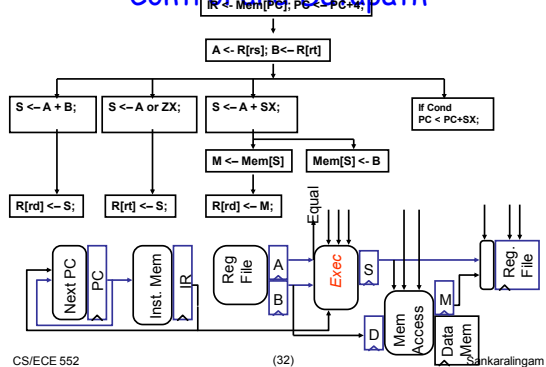


CS/ECE 552

(31)

Sankaralingam

Control and Datapath




CS/ECE 552

(32)

Sankaralingam

Pipelining the Load Instruction



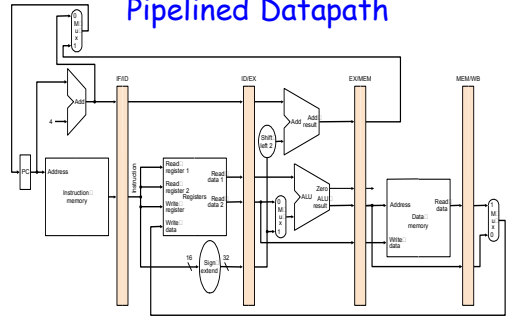
- The five independent functional units in the pipeline datapath are:
 - Instruction Memory for the *Ifetch* stage
 - Register File's Read ports (bus A and busB) for the *Reg/Dec* stage
 - ALU for the *Exec* stage
 - Data Memory for the *Mem* stage
 - Register File's *Write* port (bus W) for the *Wr* stage

CS/ECE 552

(33)

Sankaralingam

Pipelined Datapath




- Pipeline datapath with registers

CS/ECE 552

(34)

Sankaralingam

Instruction Fetch

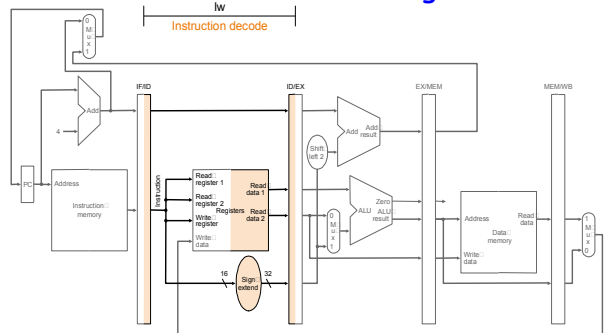


CS/ECE 552

(35)

Sankaralingam

Instruction Decode/Reg Read

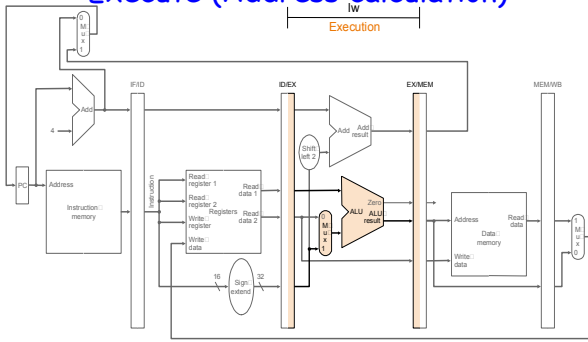


CS/ECE 552

(36)

Sankaralingam

Execute (Address calculation)




CS/ECE 552

(37)

Sankaralingam

Memory Access

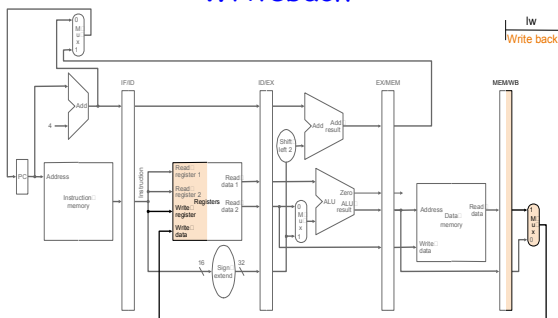


CS/ECE 552

(38)

Sankaralingam

Writeback



CS/ECE 552

(39)

Sankaralingam

Pipeline registers

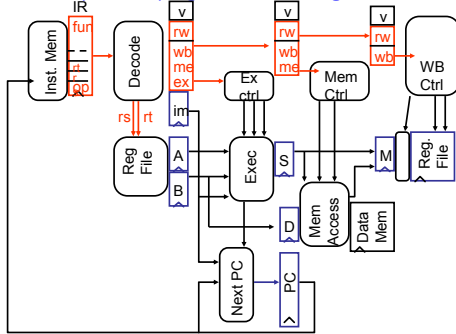
- Length of Pipeline registers
- Book says
 - IF/ID : IR (32), PC+4 (32) : **64** bits
 - ID/EX: IR (32), PC+4 (32) + RegA + RegB : **128** bits
 - EX/MEM: ALUout(32) + zero(1) + PC+4+SX(imm) (32) : **97**
 - MEM/WB: ALUout (32) + MemData(32) : **64**
- Corrections:
 - ALUout and MemData
 - Destination register (5 bits)
 - Other control bits (IR not going through)

CS/ECE 552

(40)

Sankaralingam

Recap: Carrying State in Registers




CS/ECE 552

(41)

Sankaralingam

The Four Stages of R-type




- Ifetch: Instruction Fetch
 - Fetch the instruction from the Instruction Memory
- Reg/Dec: Registers Fetch and Instruction Decode
- Exec:
 - ALU operates on the two register operands
 - Update PC
- Wr: Write the ALU output back to the register file

CS/ECE 552

(42)

Sankaralingam

Pipelining R-type and Loads



- We have pipeline conflict or structural hazard:
 - Two instructions try to write to the register file at the same time!
 - Only one write port


CS/ECE 552

(43)

Sankaralingam

Key observation

- Each functional unit can only be used **once** per instruction
- Each functional unit must be used at the **same** stage for all instructions:
 - Load uses Register File's Write Port during its **5th** stage



- R-type uses Register File's Write Port during its **4th** stage


- **2 ways to solve this pipeline hazard.**

CS/ECE 552

(44)

Sankaralingam

Soln.1: Insert "Bubble"



- Insert a "bubble" into the pipeline to prevent 2 writes at the same cycle
 - The control logic can be complex.
 - Lose instruction fetch and issue opportunity.
- No instruction is started in Cycle 6!

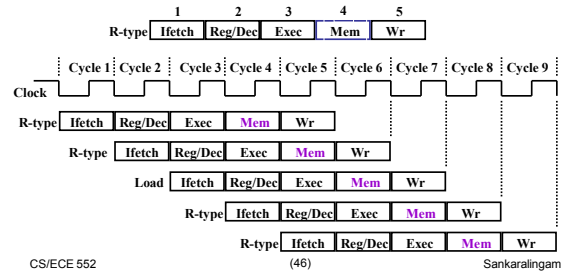
CS/ECE 552

(45)

Sankaralingam

Soln.2: Delay R-type's Write

- Delay R-type's register write by one cycle:
 - Now R-type instructions also use Reg File's write port at Stage 5
 - Mem stage is a **NOOP** stage: nothing is being done.

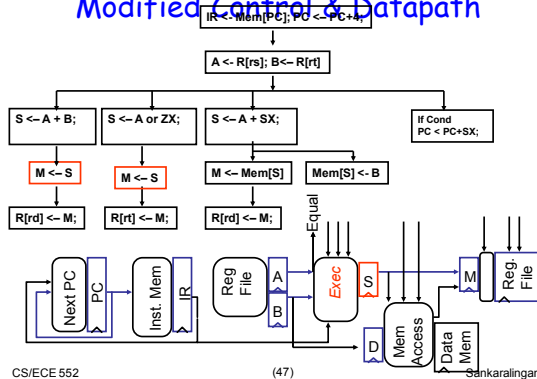


CS/ECE 552

(46)

Sankaralingam

Modified Control & Datapath




CS/ECE 552

(47)

Sankaralingam

Four Stages of Store



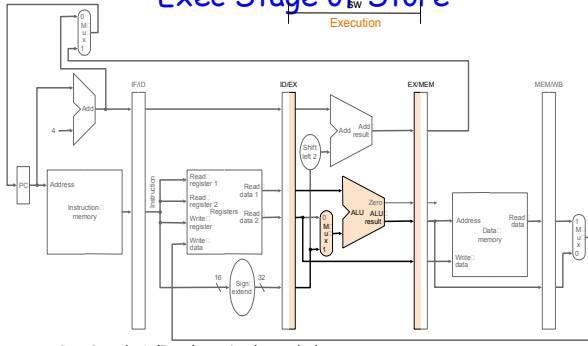
- Ifetch: Instruction Fetch
 - Fetch the instruction from the Instruction Memory
- Reg/Dec: Registers Fetch and Instruction Decode
- Exec: Calculate the memory address
- Mem: Write the data into the Data Memory

CS/ECE 552

(48)

Sankaralingam

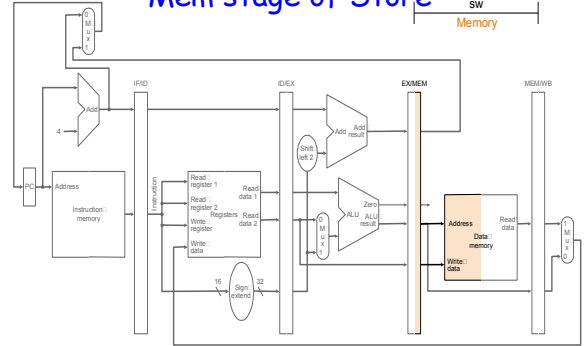
Exec Stage of Store



• Reg R and SX(Imm) are both needed (49)

Sankaralingam

Mem stage of Store




CS/ECE 552

(50)

Sankaralingam

Three Stages of Beq



- Ifetch: Instruction Fetch
 - Fetch the instruction from the Instruction Memory
- Reg/Dec:
 - Registers Fetch and Instruction Decode
- Exec:
 - compares the two register operand,
 - select correct branch target address
 - latch into PC
- Four stages as in book
 - Assume one delay slot (shadow instruction)
 - One "predict not taken" instruction

CS/ECE 552

(51)

Sankaralingam

Visualizing the pipeline

```

10 lw    r1, r2(35)
14 addl  r2, r2, 3
20 sub   r3, r4, r5
24 beq   r6, r7, 100
30 ori   r8, r9, 17
34 add   r10, r11, r12

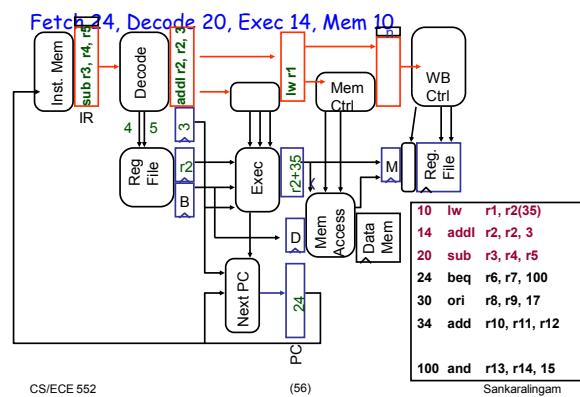
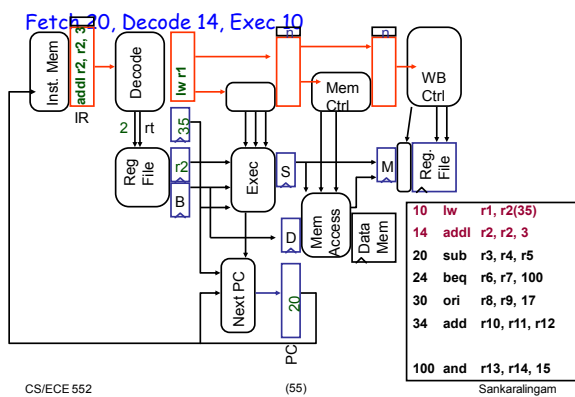
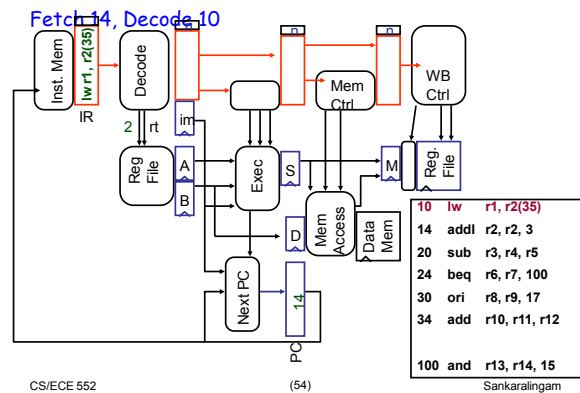
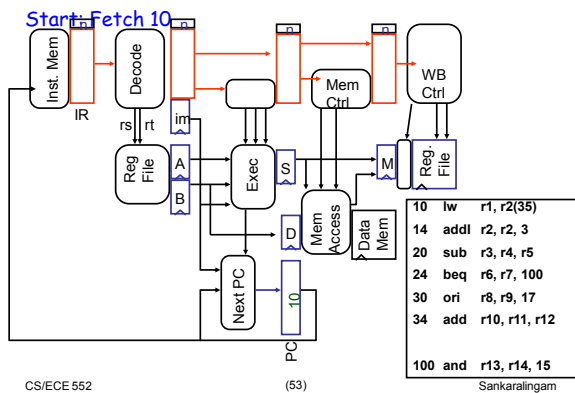
100 and  r13, r14, 15
    
```

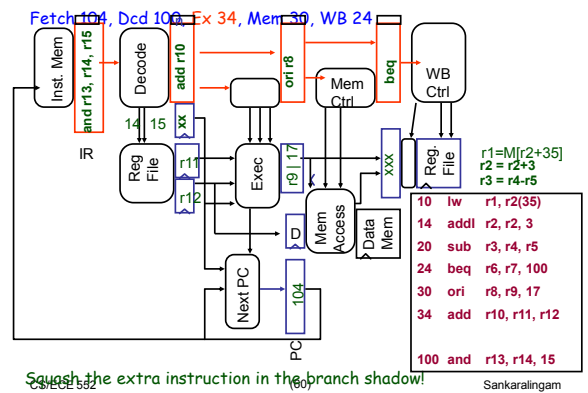
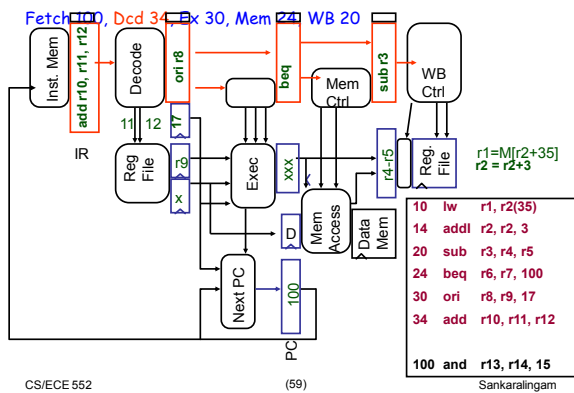
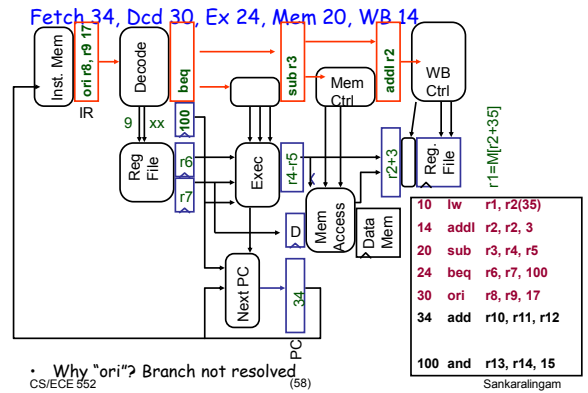
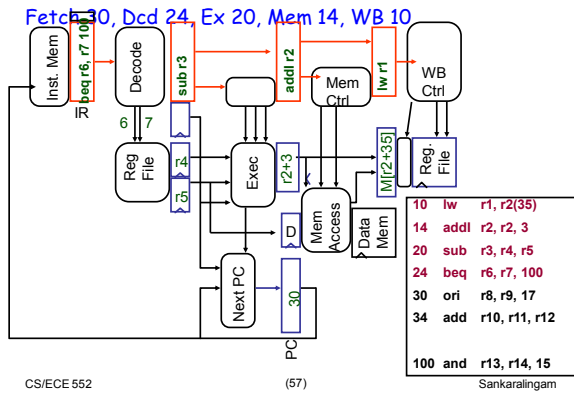
these addresses are octal

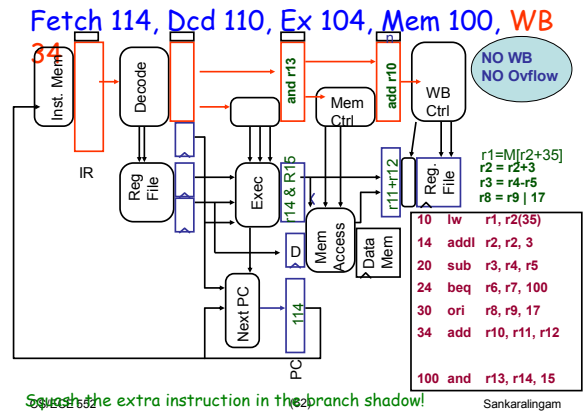
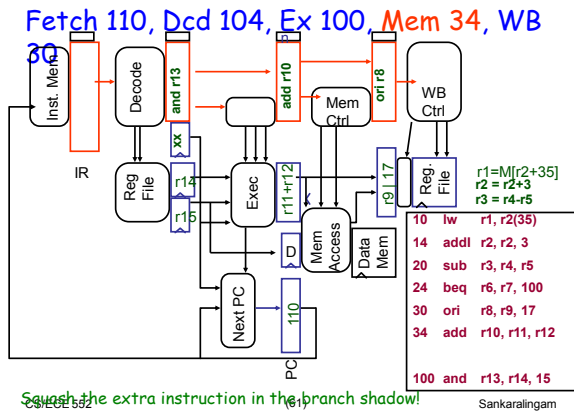
CS/ECE 552

(52)

Sankaralingam



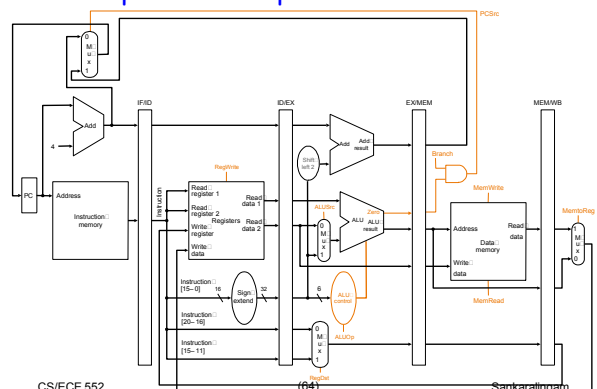




Outline

- Pipeline Datapath
 - Under simplifying assumptions
 - All independent instructions
 - Talked about bypassing/forwarding
 - Datapath not capable of handling forwarding yet
 - Walk-through
 - Delayed branches
- Control
 - Most complicated (read irregular/unstructured)
 - Gets more complicated for pipelined processors
 - But we'll start with a simple case

Pipelined Datapath with Controls



Pipeline Control vs. Single cycle control

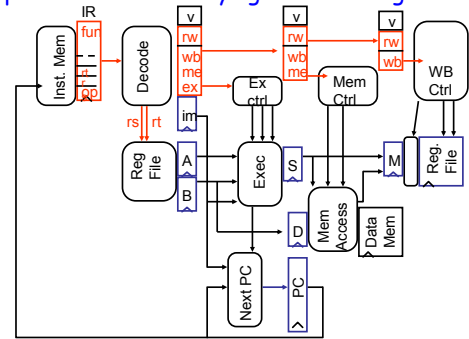
- Similarity
 - Replicated functional units like single-cycle implementation
 - Imem and Dmem
 - Separate adder for PC+ ΔX (Imm) computation
- What about
 - PCWrite? IR-write?
 - Write enable for the pipeline registers?

CS/ECE 552

(65)

Sankaralingam

Pipeline Control: Carrying State in Registers




CS/ECE 552

(66)

Sankaralingam

Focus on Control



- The Main Control generates the control signals during Reg/Dec
 - Control signals for Exec (ExtOp, ALUSrc, ...) are used 1 cycle later
 - Control signals for Mem (MemWr Branch) are used 2 cycles later
 - Control signals for Wr (MemtoReg MemWr) are used 3 cycles later

CS/ECE 552

(67)

Sankaralingam

Generating controls

- We just simplified the problem
 - Reduced pipeline control to single-cycle control (Almost)
 - **Generate** controls once
 - **Consume** (i.e., use and discard) signals as you proceed along the pipeline stages
- Identify Stage of consumption for all control signals

CS/ECE 552

(68)

Sankaralingam

Meaning of controls

- **RegWr** : 1→ write, 0→ no write
- **MemToReg** : 1→ MDR, 0→ ALUOut
- **RegDst** : 1→rd, 0→rt
- **ALUOp<1:0>** : 00→Add,01→Sub,10→'funct'
- **ALUSrc** : 0→RegB, 1→SX(Imm)
- **ExtOp****: --- needed for ORI---
- **Branch**: 0→non-branch inst, 1→ branch
- **MemRead**: 1→memread, 0→ no memread
- **MemWrite**: 1→memwrite, 0→no memwrite
- ALU control abstracted away (as before)
 - Inputs: ALUOp (2 bits), 6 "funct" bits from IR

WB <1:0>
Exec <3:0>
Mem <2:0>

Pipeline Walkthrough with controls

- Use walkthrough worksheets
 - Use code segment shown
 - Fill in controls
 - Interesting stages
 - Controls **generated** in Decode stage
 - Controls **consumed** in subsequent stages

lw \$10, 20(\$1)
sub \$11, \$2, \$3
and \$12, \$4, \$5
or \$13, \$6, \$7
add \$14, \$8, \$9

CS/ECE 552


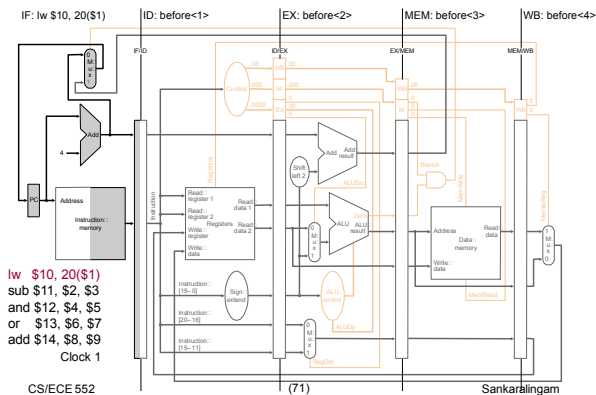
(69)





Sankaralingam

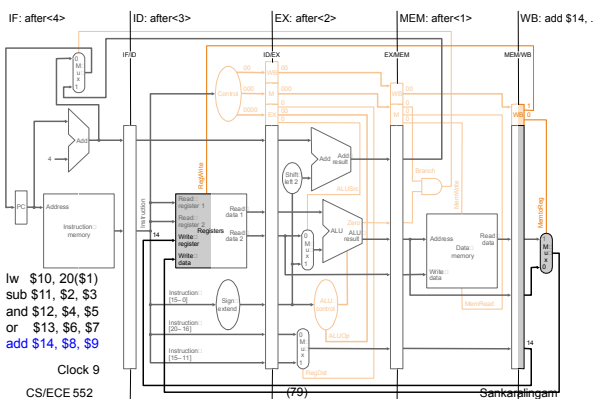


CS/ECE 552

(70)

Sankaralingam







Implementing Pipeline control

- How do we design the control logic block?
 - Similar to single-cycle implementation
 - Derive logic expressions
 - E.g. MemtoReg = lw
 - ALUSrc = lw OR sw
 - RegWrite = R-type OR lw
 - Implement Combinational logic
 - PLA implementation
 - ROM implementation

CS/ECE 552

(80)

Sankaralingam

Pipeline registers

- Preliminary estimates
 - IF/ID : IR (32), PC+4 (32) : **64** bits
 - ID/EX: IR (32), PC+4 (32) + RegA + RegB : **128** bits
 - EX/MEM: ALUout(32) + zero(1) + PC+4+SX(imm) (32) : **97**
 - MEM/WB: ALUout (32) + MemData(32) : **64**
- Corrections:
 - ALUout and MemData
 - Destination register (5 bits)
 - Other control bits (IR not going through)

CS/ECE 552

(81)

Sankaralingam

Implementing Pipeline Control

- Exercise
 - Compute required bit-width of pipeline registers
 - Before the next lecture (**NOT Feb 23rd**)

CS/ECE 552

(82)

Sankaralingam

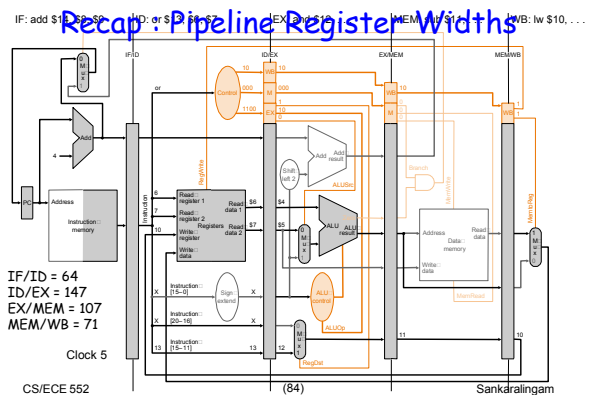
Summary

- Only slightly more complicated than single cycle
 - not really, only because we:
 - ignored complications of forwarding, branch prediction, pipeline bubbles, squashing mispredicted instructions (after branches)
 - Need deeper understanding of hazards
 - need to modify datapath as well

CS/ECE 552

(83)

Sankaralingam



Next

- Pipeline datapath and control assuming independent instructions (no hazards)
- Data hazards
 - Types
 - Detecting RAW hazards
 - Handling RAW hazards (Partial)
 - Datapath
 - Control behavior


CS/ECE 552

(85)

Sankaralingam

Data Hazards

- Challenge: maintain *illusion* of sequential execution
- Types of data hazards
 - RAW, WAR, WAW




CS/ECE 552

(86)

Sankaralingam

Data Hazards



- Avoid some "by design"
 - eliminate **WAR** by always fetching operands early (DCD) in pipe
 - eliminate **WAW** by doing all WBs in order (last stage, static)
- Detect and resolve remaining ones
 - stall or forward (if possible)

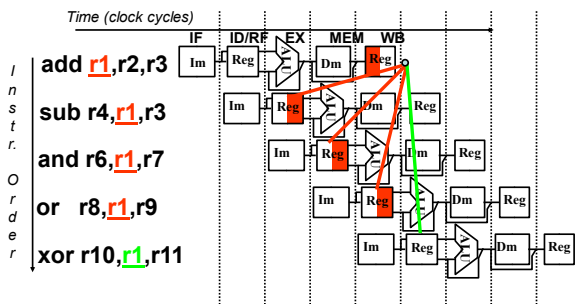
CS/ECE 552

(87)

Sankaralingam

Hazards on r1

- Dependencies backwards in time

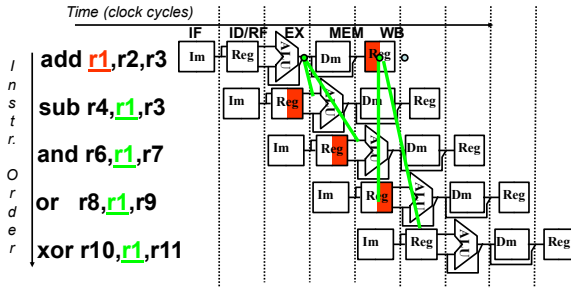


CS/ECE 552

(88)

Sankaralingam

Data Hazard Solution



CS/ECE 552

(89)

Sankaralingam

Handling RAW Hazards

- Pre-requisite for handling RAW hazard
 - Detection!
 - Need to know:
 - Pending writes
 - available results that haven't been written back to registers
 - Operand Reads
 - Later instructions that potentially use these values
 - Instructions may not write to register file (store, branch)

CS/ECE 552

(90)

Sankaralingam

Detecting RAW hazards

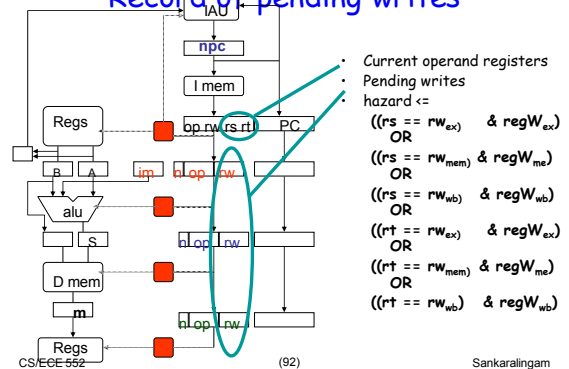
- Suppose instruction i is about to be issued and a predecessor instruction j is in the instruction pipeline.
- A RAW hazard exists on register p if $p \in \text{Rregs}(i) \cap \text{Wregs}(j)$
 - Keep a record of pending writes (for inst's in the pipe) and compare with operand regs of current instruction.
 - When instruction issues, reserve its result register.
 - When on operation completes, remove its write reservation
- A WAW hazard exists on register p if $p \in \text{Wregs}(i) \cap \text{Wregs}(j)$
- A WAR hazard exists on register p if $p \in \text{Wregs}(i) \cap \text{Rregs}(j)$

CS/ECE 552

(91)

Sankaralingam

Record of pending writes



CS/ECE 552

(92)

Sankaralingam

Logic equations for Hazard Detection

- Restatement of equations
- Text book version
 - WB stage is not really a hazard
 - Data is written in first half of cycle, read in 2nd half
 - $EX/MEM.RegisterRd = ID/EX.RegisterRs$
 - $EX/MEM.RegisterRd = ID/EX.RegisterRt$
 - $MEM/WB.RegisterRd = ID/EX.RegisterRs$
 - $MEM/WB.RegisterRd = ID/EX.RegisterRt$

CS/ECE 552

(93)

Sankaralingam

Lookahead: Forwarding datapath

- We know how to detect RAW hazards
- Now,
 - Modify Datapath to enable forwarding
 - Desired control behavior

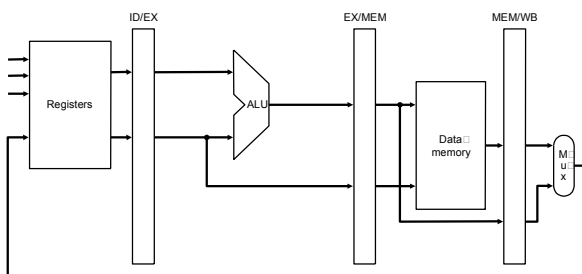
CS/ECE 552

(94)

Sankaralingam

Base Pipelined Datapath

- Simplified representation of pipelined datapath
 - To avoid clutter

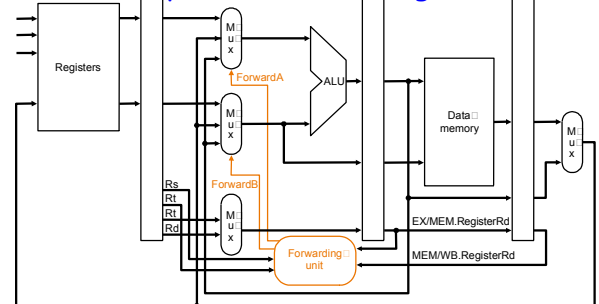


a. No forwarding

(95)

Sankaralingam

Datapath w/Forwarding Unit



b. With forwarding

CS/ECE 552

(96)

Sankaralingam

ForwardA/ForwardB: 01->Mem, 10->EX

Data Hazards and Forwarding: Walkthrough

- Code snippet
 - identify hazards
 - identify forwarding paths

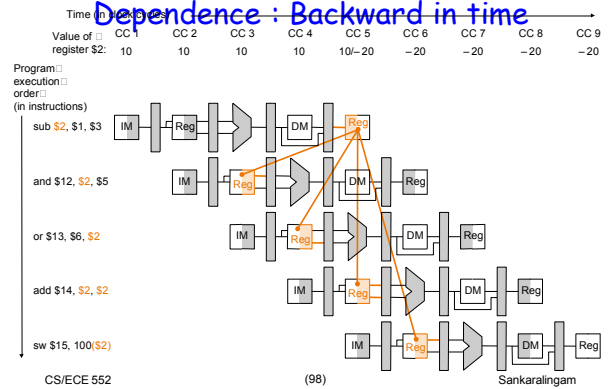
sub \$2, \$1, \$3
and \$4, \$2, \$5
or \$4, \$4, \$2
add \$9, \$4, \$2

CS/ECE 552

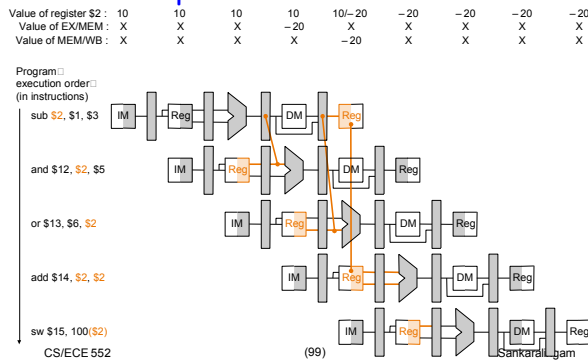
(97)

Sankaralingam


Dependence : Backward in time

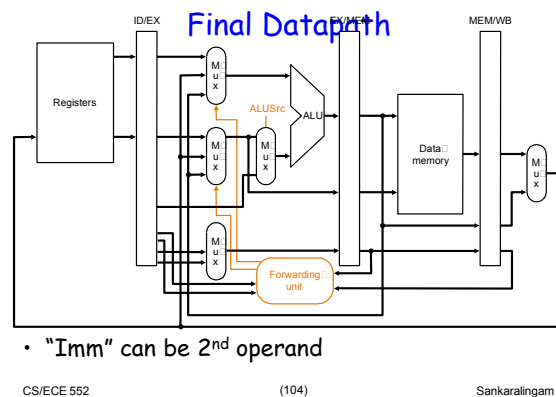





True dependence : Forward in time



Walkthrough





Forwarding Control Behavior

- EX hazard

```
If (EX/MEM.RegWrite AND // not store or branch
    EX/MEM.RegisterRd != 0 AND // Result is used
    EX/MEM.RegisterRd = ID/EX.RegisterRs)
    ForwardA = 10
```

```
If (EX/MEM.RegWrite AND
    EX/MEM.RegisterRd != 0 AND
    EX/MEM.RegisterRd = ID/EX.RegisterRt)
    ForwardB = 10
```

CS/ECE 552

(105)

Sankaralingam

Forwarding Control Behavior

- MEM hazard

```
If (MEM/WB.RegWrite AND
    MEM/WB.RegisterRd != 0 AND
    MEM/WB.RegisterRd = ID/EX.RegisterRs)
    ForwardA = 01
```

```
If (MEM/WB.RegWrite AND
    MEM/WB.RegisterRd != 0 AND
    MEM/WB.RegisterRd = ID/EX.RegisterRt)
    ForwardB = 01
```


- Does this fully meet our requirements ?

CS/ECE 552

(106)

Sankaralingam

Lookahead: RAW hazard with load



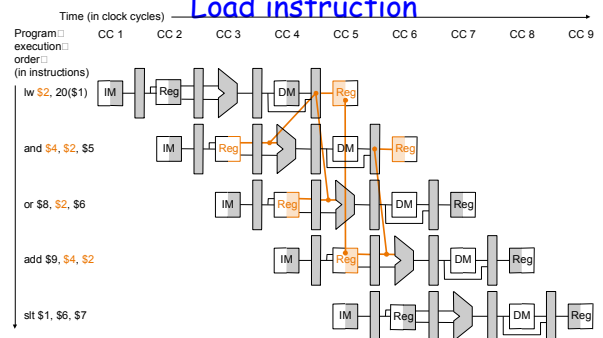
- Forwarding as solution to RAW hazard
 - possible if no (true) dependence going backwards in time
 - True for R-type instructions
 - Data available after EX stage (i.e., at ALUOut)
 - Not true for load instruction

CS/ECE 552

(107)

Sankaralingam

Load instruction



- Replaced "sub" with "lw" in previous code-example

CS/ECE 552

(108)

Sankaralingam

Solution

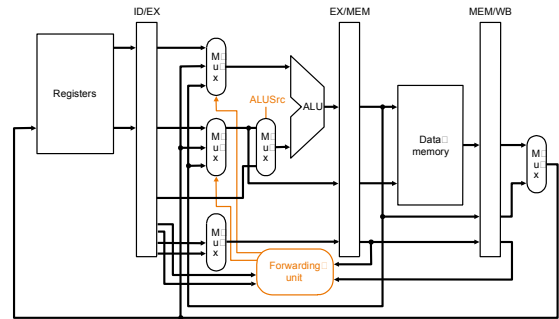
- Catch-all solution for hazards
 - Stall
 - always works, but hurts performance
 - Use as last resort
- Challenge:
 - Modify pipeline implementation to support stalls when hazards are detected

CS/ECE 552

(109)

Sankaralingam

Recap



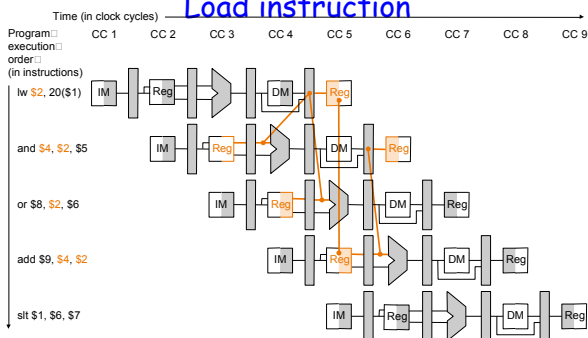
- Designed forwarding unit to solve RAW hazards for R-type instructions

CS/ECE 552

(110)

Sankaralingam

Load instruction



- True backward (in time) dependence

CS/ECE 552

(111)

Sankaralingam

Hazards with load instruction

- True dependencies: backward in time
- Stall the pipeline
- Minor change in terminology
 - If **forwarding** can solve it, it is not a hazard!
 - "**Hazard**" refers only to true backward dependencies in time.

CS/ECE 552

(112)

Sankaralingam

Handling the hazard

- As before
 - Detection
 - Logic equations to detect hazard
 - Actual stalling
 - Datapath/control modifications to achieve stalling

Detection

- Conditions
 - Preceding instruction must read memory
 - **MemRead** must be asserted
 - Destination of preceding instruction (**rt**) must be one of operands of current instruction
- Logic equations- restate above conditions formally
 - If(**ID/EX.MemRead** AND ((**ID/EX.RegRt = IF/ID.RegRs**) OR (**ID/EX.RegRt = IF/ID.RegRt**)))
STALL

lw \$2, 20(\$1)
and \$4, \$2, \$5
or \$4, \$4, \$2
add \$9, \$4, \$2

CS/ECE 552

(113)

Sankaralingam

CS/ECE 552

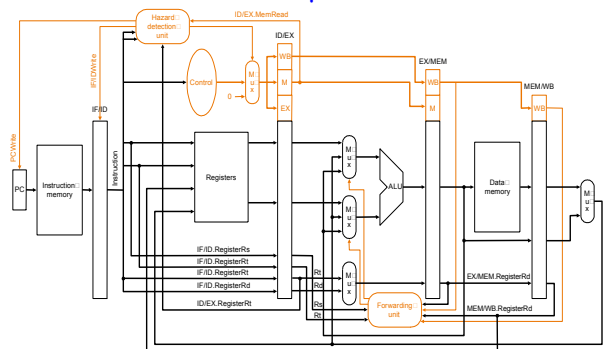
(114)

Sankaralingam

Stalling the pipeline

- Instruction cannot proceed
 - Following instruction must be stalled too.
 - Otherwise state in pipeline registers is overwritten
- Preceding instructions may proceed as usual
- Solution
 - inject NOP into EX/Mem pipeline
 - Prevent writes to PC to IF/ID register

Datapath



CS/ECE 552

(115)


Sankaralingam

CS/ECE 552


(116)

Sankaralingam


Walk-through (1 of 6)




Walk-through (2 of 6)




Walk-through (3 of 6)




Walk-through (4 of 6)



Walk-through (5 of 6)



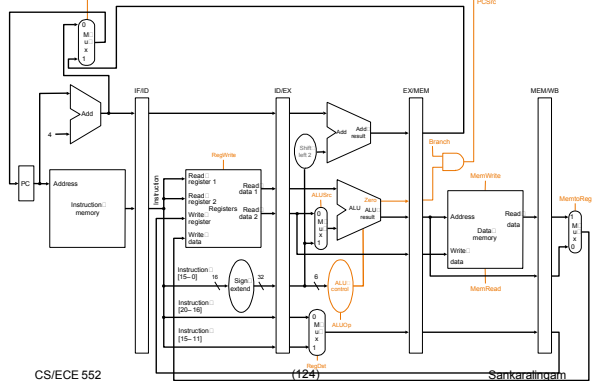
Walk-through (6 of 6)



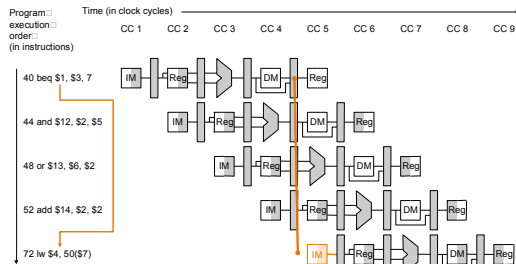
RAW Hazard with Loads: Summary

- True backward dependencies in time
 - Need to stall
- Stall achieved by
 - Detecting hazard (remember logic equation)
 - Inserting NOP (all EX/MEM/WB controls set to 0)
 - Preventing IF/ID register and PC from being overwritten
- Next Branch/Control Hazards

When conditional branches resolved?



Branch Hazards



- Branch resolved in the MEM stage
- If taken,
 - $PC \leftarrow PC + 4 + SX(Imm \cdot 4)$
 - $40 + 4 + 7 \cdot 4 = 72$

CS/ECE 552

(125)

Sankaralingam

Control/Branch Hazards


- Branch resolved in the MEM stage
 - But next instruction has to be fetched in the next cycle
 - Reduce the penalty by moving decision earlier in pipeline
 - Need additional **comparator** ($r1=r2?$) and **adder** ($PC+4+SX(Imm) \cdot 4$)
 - Reduced penalty from 3 cycles to 1 cycle

CS/ECE 552

(126)

Sankaralingam

Datapath for branch hazards



CS/ECE 552

(127)

Sankaralingam

Eliminate 1-cycle stall?


- Two solutions
 - Predict branch is always not taken
 - More sophisticated prediction schemes
 - Delay slots
 - Compiler's problem
- Walkthrough example for solution #1
 - Predict not taken

CS/ECE 552


(128)

Sankaralingam

Walkthrough (1 of 2)




Walkthrough (2 of 2)



Dynamic Branch Prediction


- Better than static prediction
 - Branches are predictable
 - ~90% of program execution time is spent in ~10% of code (inner loops)
 - Think of a program loop of N iterations
 - Taken $N-1$ times
 - Not taken last time

Dynamic Branch Prediction



- How does hardware "learn" branch behavior?
- Store each branch instruction's history ***
 - If a branch was taken "recently", predict taken
 - One bit saturating counter
 - Two bit counters

Branch Prediction



- Store each branch's history ***
 - Not really
- Keep a small table indexed by program counter
- PC is large (32 bit number)
- Mapping to number of table entries
 - E.g. 16-entry branch prediction table
 - Mapping: use last 4 bits of PC
- Problem: Multiple branches may map to same entry in table -- Aliasing

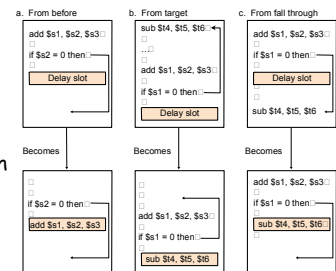
CS/ECE 552

(133)

Sankaralingam

"Easy way"★ to hide branch hazard delay

- Delayed branch
 - Instruction after branch always executes
 - Find an independent instruction from before the branch
 - Find instructions from Taken (target) OR from Not Taken (fall-through) code section
- ★ For Architects



CS/ECE 552

(134)

Sankaralingam

Big picture


- Iron law: Insts/prog * CPI * cycletime
- With pipelining:
 - CPI ~ 1 (with ideal memory, good branch prediction and few data hazards)
 - Cycletime : determined by critical path of one stage

CS/ECE 552

(135)

Sankaralingam

The Design Space



- Summary of design trade-offs
- Can we do better?
 - CPI < 1?
 - i.e., IPC > 1?

CS/ECE 552

(136)

Sankaralingam

Superscalar Processor

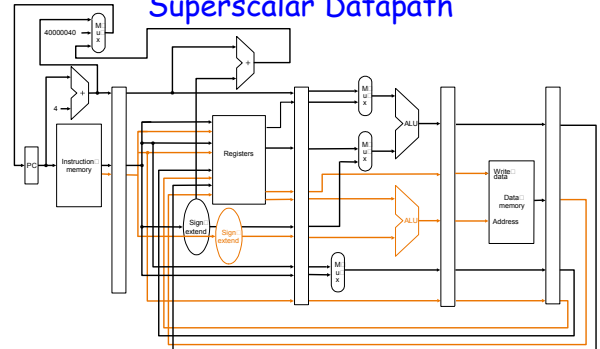
- What does it mean?
 - Scalar processors (operate on scalar quantities)
 - Vector (operate on vectors)
 - Superscalar: multiple scalar operations in one cycle
 - More than one instruction per cycle

CS/ECE 552

(137)

Sankaralingam

Superscalar Datapath



- Replicate datapath elements

CS/ECE 552

(138)

Sankaralingam

Dynamic Scheduling

- No need to suffer hazards if other useful work can be achieved
- Load Hazard results in pipeline stall
 - But other instructions are ready
 - "Oh! But we cannot execute instructions out of order" - Not really

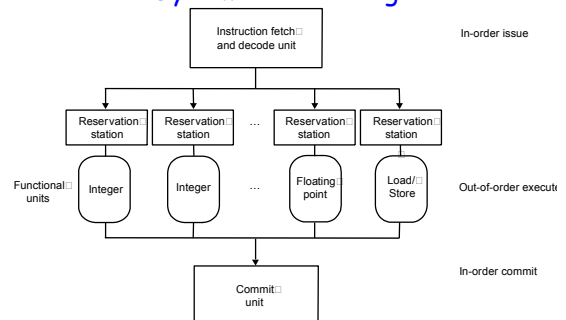
```
lw    $t0, 20($s2)
addu  $t1, $t0, $t2
sub    $s4, $s4, $t3
slli   $t5, $s4, $t3
```

CS/ECE 552

(139)

Sankaralingam

Dynamic Scheduling



- Instructions can execute when operands are ready
- Instructions can "commit" when all preceding instructions have committed

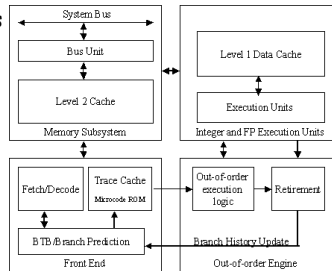
CS/ECE 552

(140)

Sankaralingam

Pentium 4 on 0.18 micron

- 42 million transistors
- 3GHz
- Several parts are clocked at half the speed
- Inorder front-end, out-of-order execution, in order retire




CS/ECE 552

(141)

Sankaralingam

Pentium 4 pipeline



- Pipeline too much; c.f., Core2

CS/ECE 552

(142)

Sankaralingam