

COMPUTER ORGANIZATION AND DESIGN
The Hardware/Software Interface

Chapter 4
The Processor

Introduction

- CPU performance factors
 - Instruction count
 - Determined by ISA and compiler
 - CPI and Cycle time
 - Determined by CPU hardware
- We will examine two MIPS implementations
 - A simplified version
 - A more realistic pipelined version
- Simple subset, shows most aspects
 - Memory reference: lw, sw
 - Arithmetic/logical: add, sub, and, or, slt
 - Control transfer: beq, j

Chapter 4 — The Processor — 2

Instruction Execution

- PC → instruction memory, fetch instruction
- Register numbers → register file, read registers
- Depending on instruction class
 - Use ALU to calculate
 - Arithmetic result
 - Memory address for load/store
 - Branch target address
 - Access data memory for load/store
 - PC ← target address or PC + 4

Chapter 4 — The Processor — 3

CPU Overview

Chapter 4 — The Processor — 4

Multiplexers

Can't just join wires together

- Use multiplexers

Chapter 4 — The Processor — 5

Control

Chapter 4 — The Processor — 6

Logic Design Basics

- Information encoded in binary
 - Low voltage = 0, High voltage = 1
 - One wire per bit
 - Multi-bit data encoded on multi-wire buses
- Combinational element
 - Operate on data
 - Output is a function of input
- State (sequential) elements
 - Store information

Chapter 4 — The Processor — 7

Combinational Elements

- AND-gate
 - $Y = A \& B$
- Adder
 - $Y = A + B$
- Multiplexer
 - $Y = S ? I1 : I0$
- Arithmetic/Logic Unit
 - $Y = F(A, B)$

Chapter 4 — The Processor — 8

Sequential Elements

- Register: stores data in a circuit
 - Uses a clock signal to determine when to update the stored value
 - Edge-triggered: update when Clk changes from 0 to 1

Chapter 4 — The Processor — 9

Sequential Elements

- Register with write control
 - Only updates on clock edge when write control input is 1
 - Used when stored value is required later

Chapter 4 — The Processor — 10

Clocking Methodology

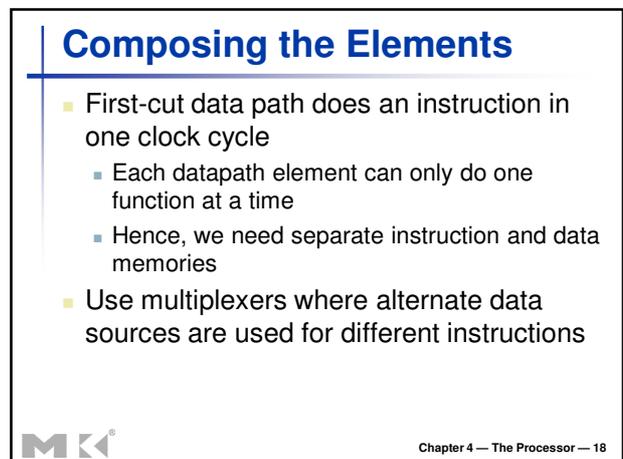
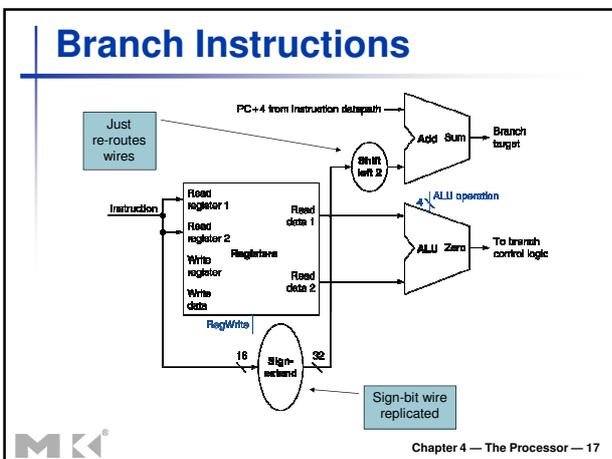
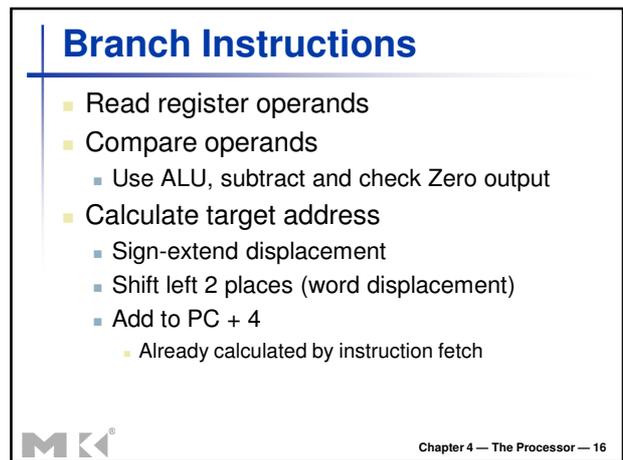
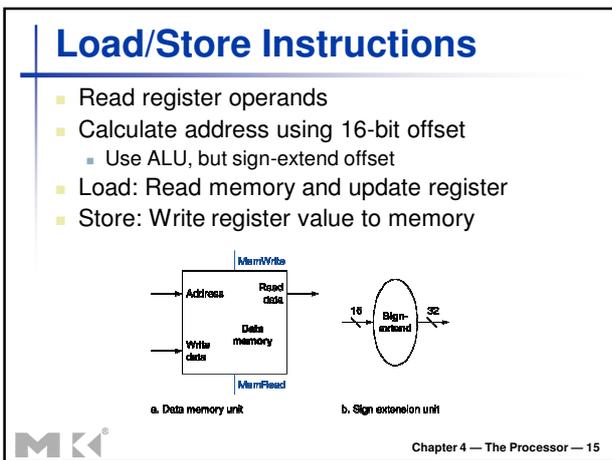
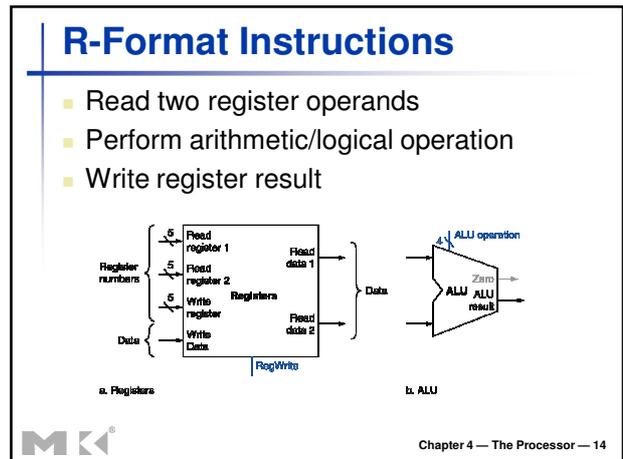
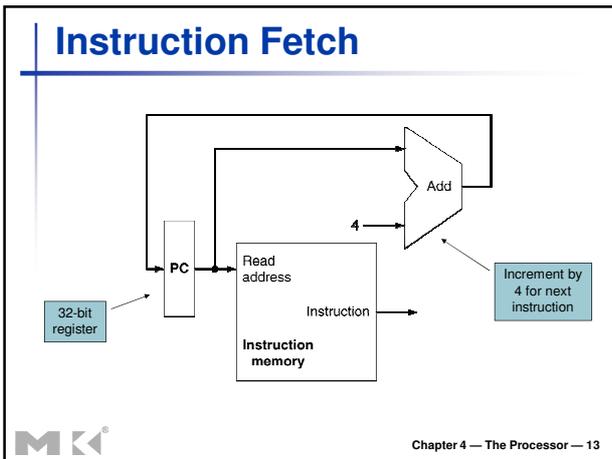
- Combinational logic transforms data during clock cycles
 - Between clock edges
 - Input from state elements, output to state element
 - Longest delay determines clock period

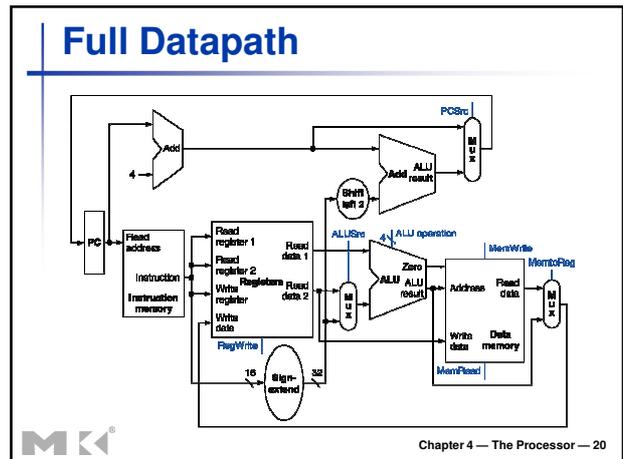
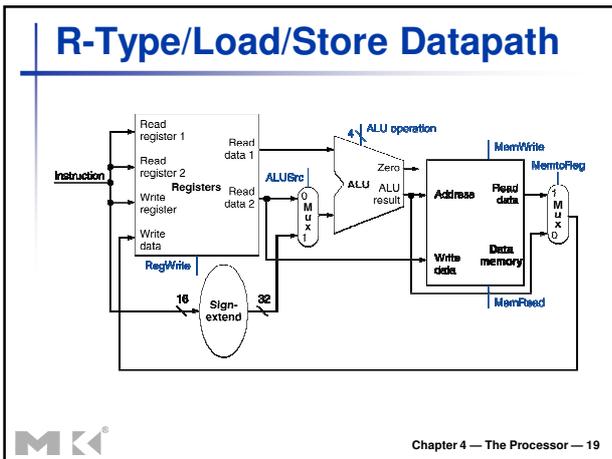
Chapter 4 — The Processor — 11

Building a Datapath

- Datapath
 - Elements that process data and addresses in the CPU
 - Registers, ALUs, mux's, memories, ...
- We will build a MIPS datapath incrementally
 - Refining the overview design

Chapter 4 — The Processor — 12





ALU Control

- ALU used for
 - Load/Store: F = add
 - Branch: F = subtract
 - R-type: F depends on funct field

| ALU control | Function |
|-------------|------------------|
| 0000 | AND |
| 0001 | OR |
| 0010 | add |
| 0110 | subtract |
| 0111 | set-on-less-than |
| 1100 | NOR |

Chapter 4 — The Processor — 21

ALU Control

- Assume 2-bit ALUOp derived from opcode
 - Combinational logic derives ALU control

| opcode | ALUOp | Operation | funct | ALU function | ALU control |
|--------|-------|------------------|--------|------------------|-------------|
| lw | 00 | load word | XXXXXX | add | 0010 |
| sw | 00 | store word | XXXXXX | add | 0010 |
| beq | 01 | branch equal | XXXXXX | subtract | 0110 |
| R-type | 10 | add | 100000 | add | 0010 |
| | | subtract | 100010 | subtract | 0110 |
| | | AND | 100100 | AND | 0000 |
| | | OR | 100101 | OR | 0001 |
| | | set-on-less-than | 101010 | set-on-less-than | 0111 |

Chapter 4 — The Processor — 22

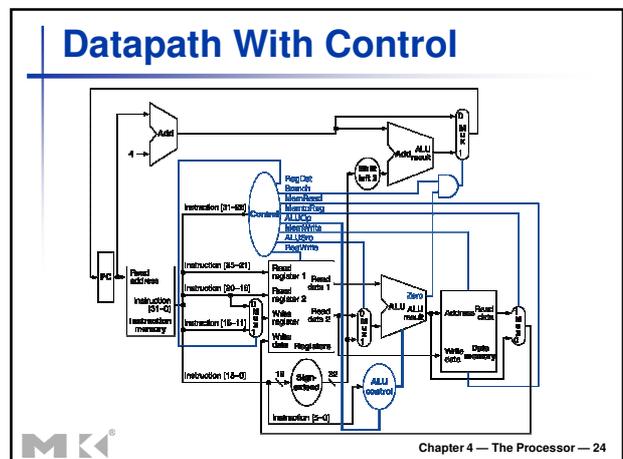
The Main Control Unit

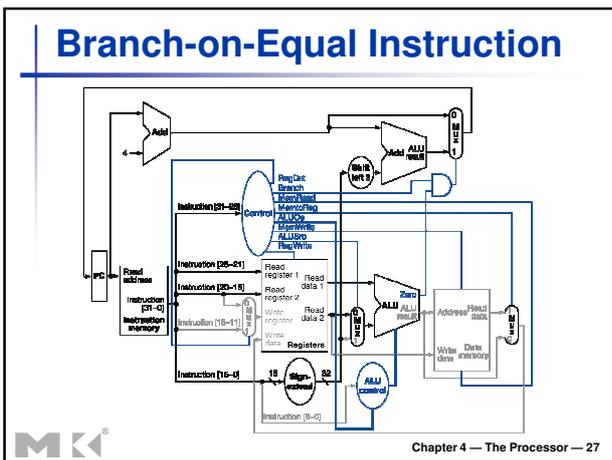
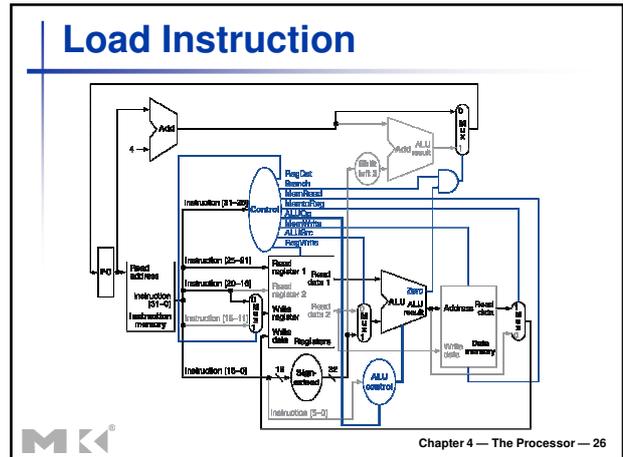
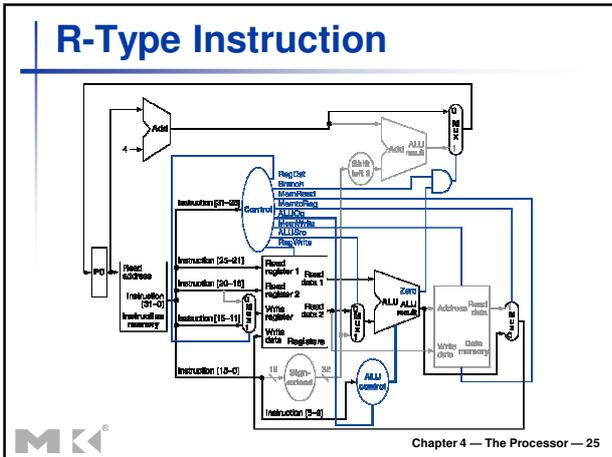
- Control signals derived from instruction

| Instruction Type | opcode | rs | rt | rd | shamt | funct |
|------------------|----------|----|----|---------|-------|-------|
| R-type | 0 | rs | rt | rd | shamt | funct |
| Load/Store | 35 or 43 | rs | rt | address | | |
| Branch | 4 | rs | rt | address | | |

Legend:
 - opcode: always read
 - rs: read, except for load
 - rt: write for R-type and load
 - rd: sign-extend and add

Chapter 4 — The Processor — 23



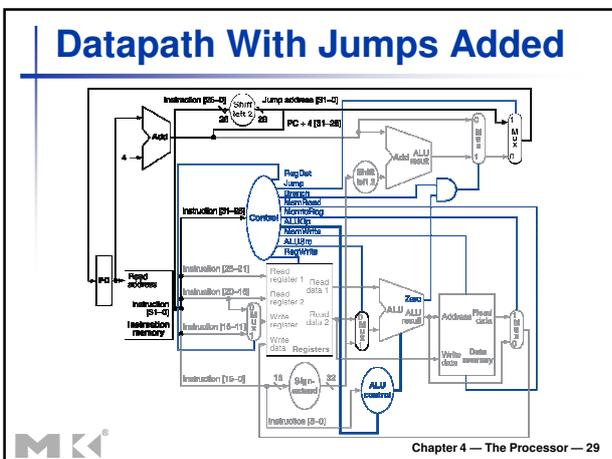


Implementing Jumps

| | | |
|------|-------|---------|
| Jump | 2 | address |
| | 31:26 | 25:0 |

- Jump uses word address
- Update PC with concatenation of
 - Top 4 bits of old PC
 - 26-bit jump address
 - 00
- Need an extra control signal decoded from opcode

Chapter 4 — The Processor — 28



Performance Issues

- Longest delay determines clock period
 - Critical path: load instruction
 - Instruction memory → register file → ALU → data memory → register file
- Not feasible to vary period for different instructions
- Violates design principle
 - Making the common case fast
- We will improve performance by pipelining

Chapter 4 — The Processor — 30

Pipelining Analogy

- Pipelined laundry: overlapping execution
 - Parallelism improves performance

- Four loads:
 - Speedup = $8/3.5 = 2.3$
- Non-stop:
 - Speedup = $2n/0.5n + 1.5 \approx 4$
 - = number of stages

Chapter 4 — The Processor — 31

MIPS Pipeline

- Five stages, one step per stage
 - IF: Instruction fetch from memory
 - ID: Instruction decode & register read
 - EX: Execute operation or calculate address
 - MEM: Access memory operand
 - WB: Write result back to register

Chapter 4 — The Processor — 32

Pipeline Performance

- Assume time for stages is
 - 100ps for register read or write
 - 200ps for other stages
- Compare pipelined datapath with single-cycle datapath

| Instr | Instr fetch | Register read | ALU op | Memory access | Register write | Total time |
|----------|-------------|---------------|--------|---------------|----------------|------------|
| lw | 200ps | 100 ps | 200ps | 200ps | 100 ps | 800ps |
| sw | 200ps | 100 ps | 200ps | 200ps | | 700ps |
| R-format | 200ps | 100 ps | 200ps | | 100 ps | 600ps |
| beq | 200ps | 100 ps | 200ps | | | 500ps |

Chapter 4 — The Processor — 33

Pipeline Performance

Chapter 4 — The Processor — 34

Pipeline Speedup

- If all stages are balanced
 - i.e., all take the same time
 - Time between instructions_{pipelined} = $\frac{\text{Time between instructions}_{\text{nonpipelined}}}{\text{Number of stages}}$
- If not balanced, speedup is less
- Speedup due to increased throughput
 - Latency (time for each instruction) does not decrease

Chapter 4 — The Processor — 35

Pipelining and ISA Design

- MIPS ISA designed for pipelining
 - All instructions are 32-bits
 - Easier to fetch and decode in one cycle
 - c.f. x86: 1- to 17-byte instructions
 - Few and regular instruction formats
 - Can decode and read registers in one step
 - Load/store addressing
 - Can calculate address in 3rd stage, access memory in 4th stage
 - Alignment of memory operands
 - Memory access takes only one cycle

Chapter 4 — The Processor — 36

Hazards

- Situations that prevent starting the next instruction in the next cycle
- Structure hazards
 - A required resource is busy
- Data hazard
 - Need to wait for previous instruction to complete its data read/write
- Control hazard
 - Deciding on control action depends on previous instruction



Chapter 4 — The Processor — 37

Structure Hazards

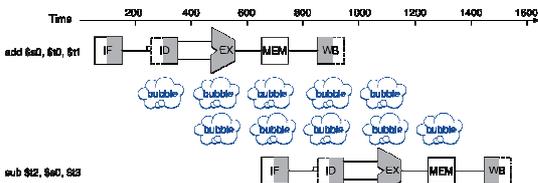
- Conflict for use of a resource
- In MIPS pipeline with a single memory
 - Load/store requires data access
 - Instruction fetch would have to *stall* for that cycle
 - Would cause a pipeline “bubble”
- Hence, pipelined datapaths require separate instruction/data memories
 - Or separate instruction/data caches



Chapter 4 — The Processor — 38

Data Hazards

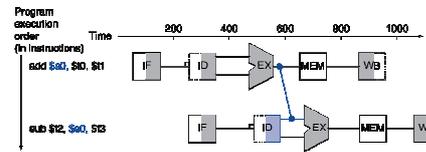
- An instruction depends on completion of data access by a previous instruction
 - add \$s0, \$t0, \$t1
 - sub \$t2, \$s0, \$t3



Chapter 4 — The Processor — 39

Forwarding (aka Bypassing)

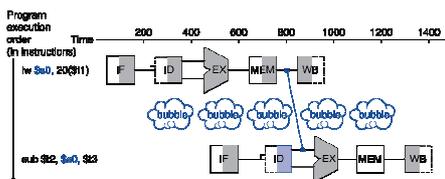
- Use result when it is computed
 - Don't wait for it to be stored in a register
 - Requires extra connections in the datapath



Chapter 4 — The Processor — 40

Load-Use Data Hazard

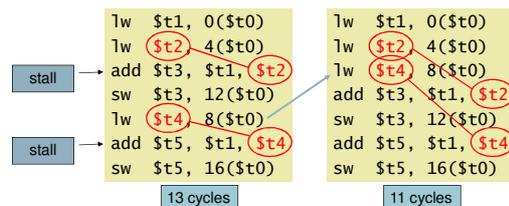
- Can't always avoid stalls by forwarding
 - If value not computed when needed
 - Can't forward backward in time!



Chapter 4 — The Processor — 41

Code Scheduling to Avoid Stalls

- Reorder code to avoid use of load result in the next instruction
- C code for $A = B + E$; $C = B + F$;



Chapter 4 — The Processor — 42

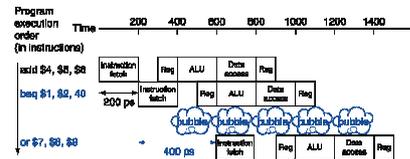
Control Hazards

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



Stall on Branch

- Wait until branch outcome determined before fetching next instruction

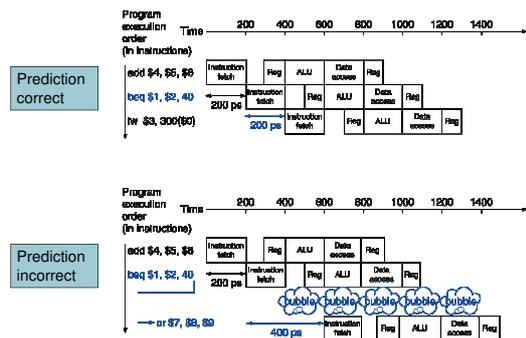


Branch Prediction

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



MIPS with Predict Not Taken



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

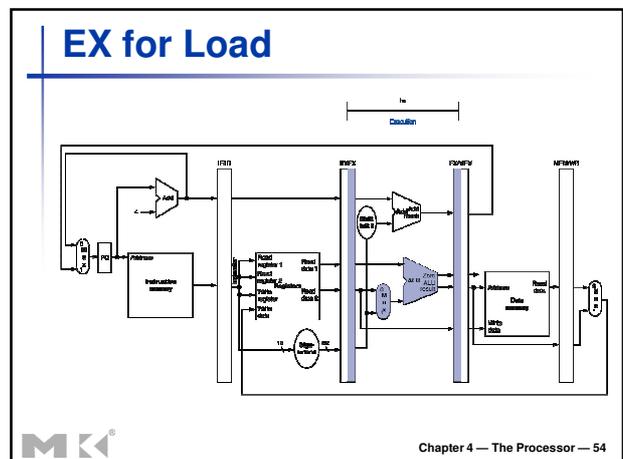
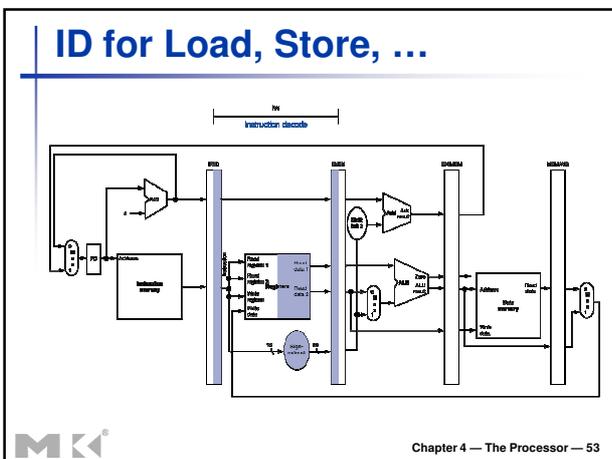
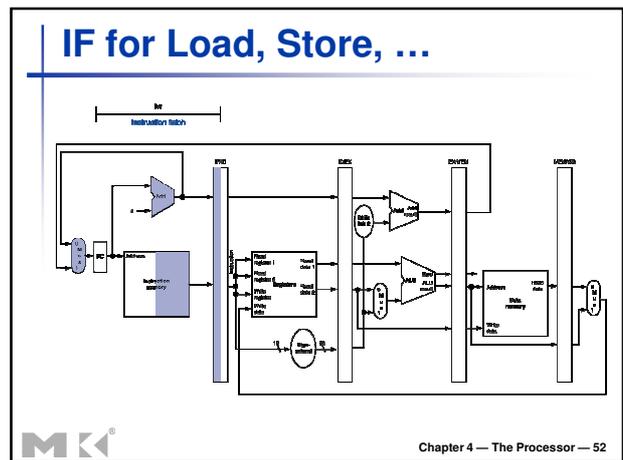
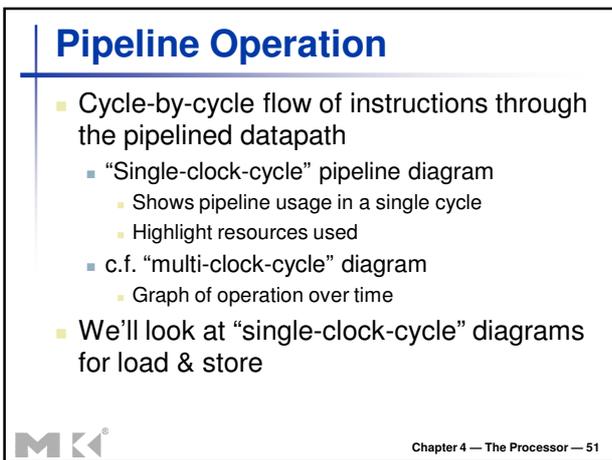
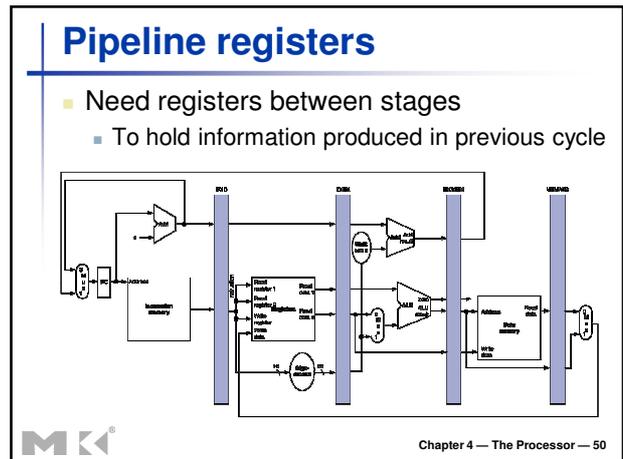
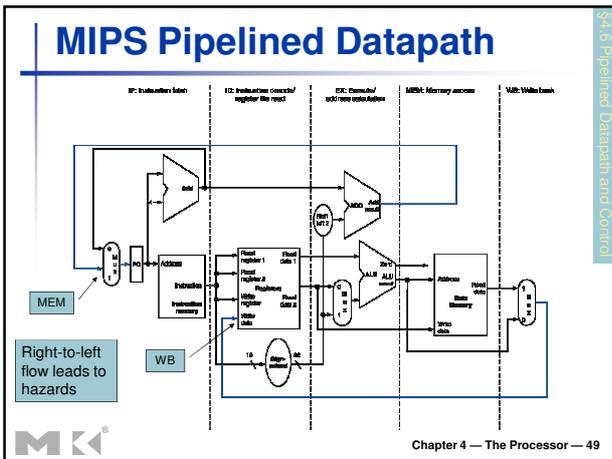


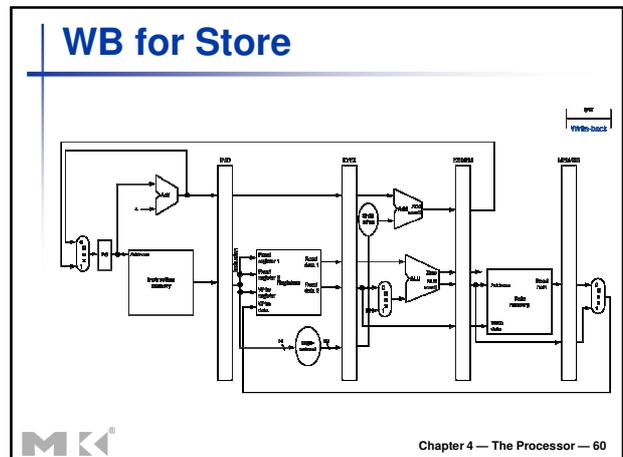
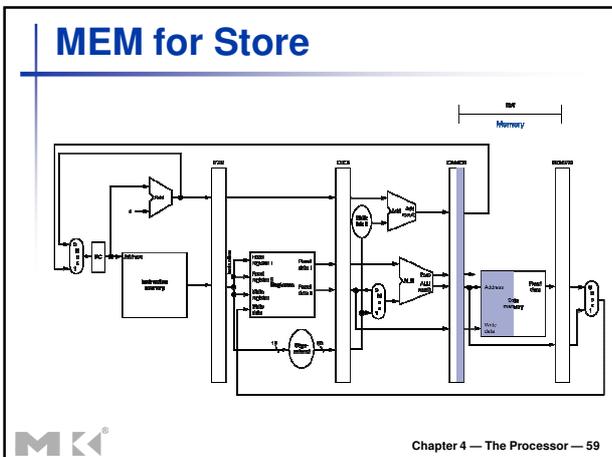
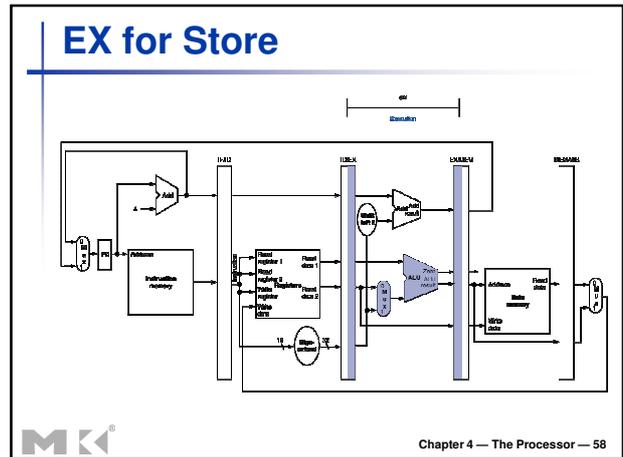
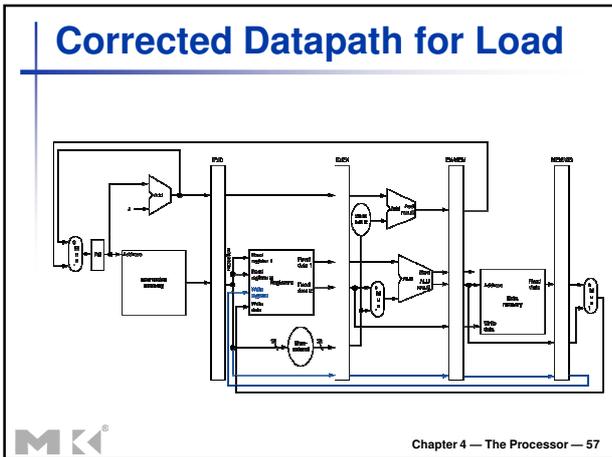
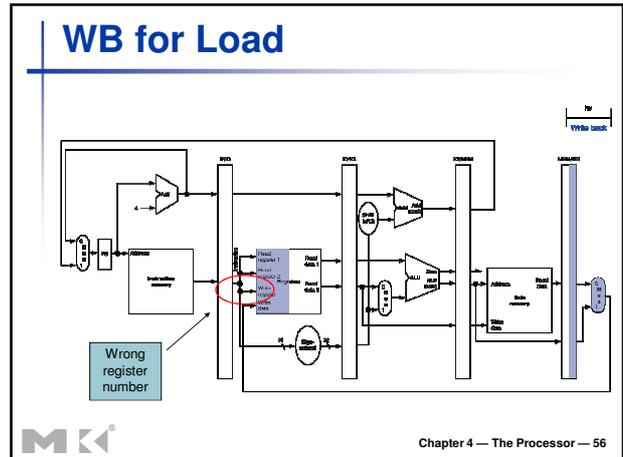
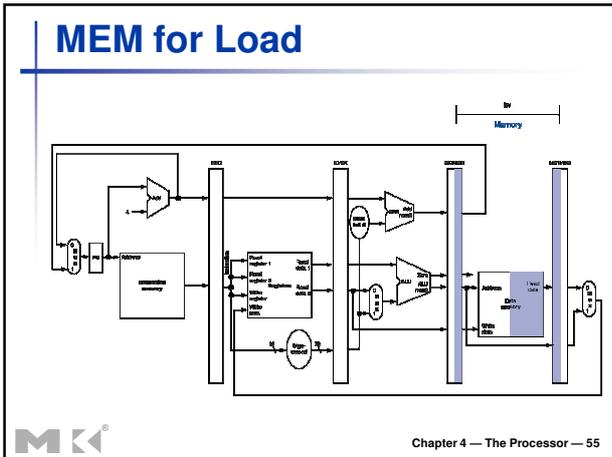
Pipeline Summary

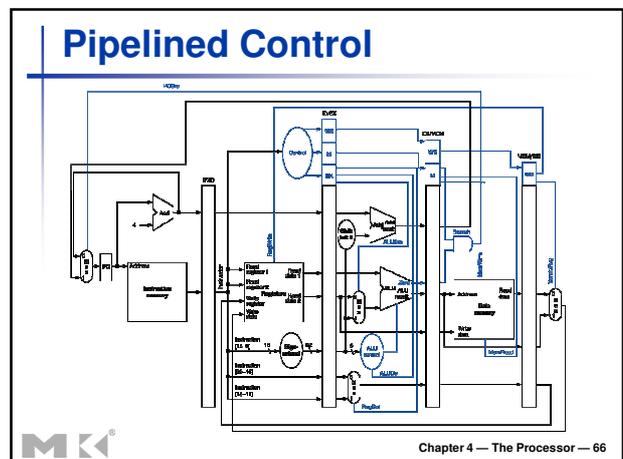
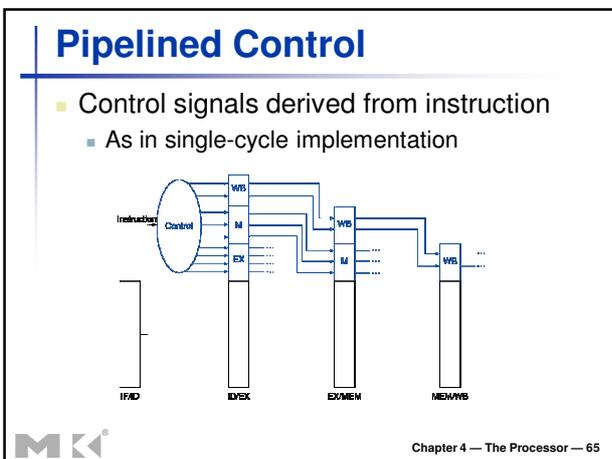
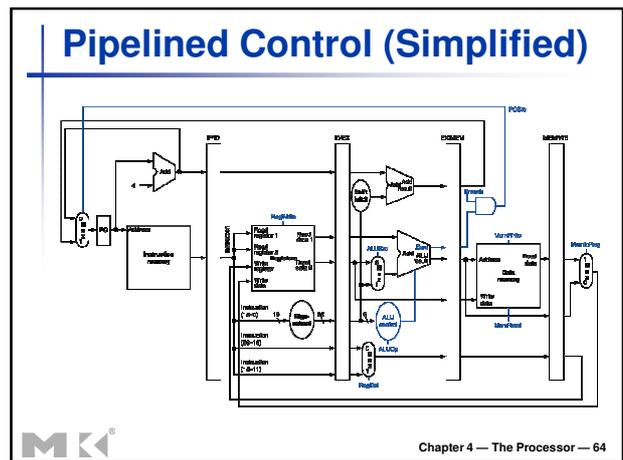
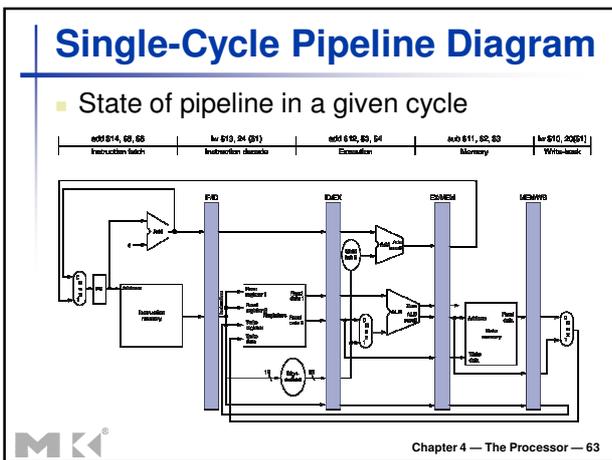
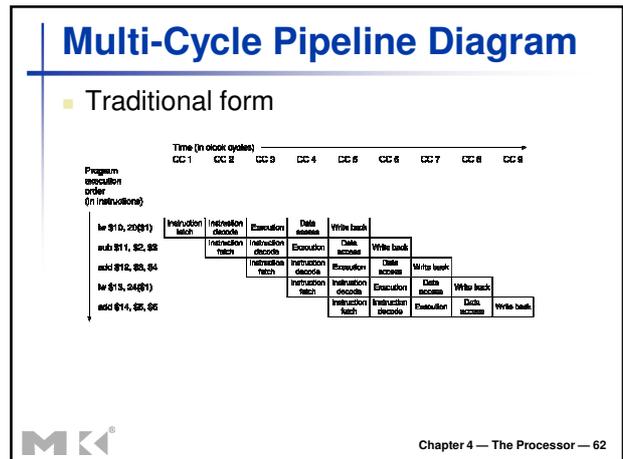
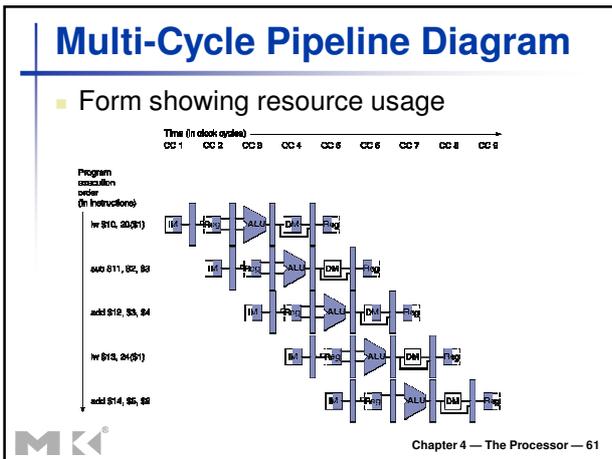
The BIG Picture

- 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









Data Hazards in ALU Instructions

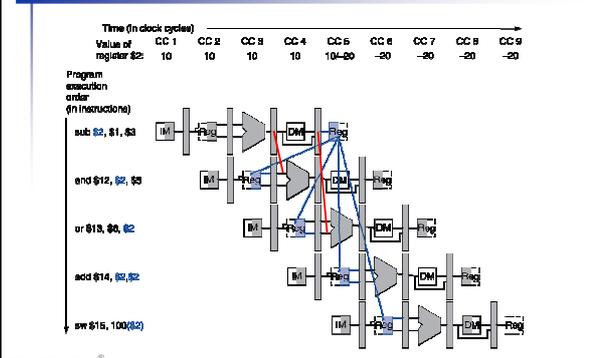
- Consider this sequence:
 - sub \$2, \$1, \$3
 - and \$12, \$2, \$5
 - or \$13, \$6, \$2
 - add \$14, \$2, \$2
 - sw \$15, 100(\$2)
- We can resolve hazards with forwarding
 - How do we detect when to forward?

§4.7 Data Hazards: Forwarding vs. Stalling



Chapter 4 — The Processor — 67

Dependencies & Forwarding



Time (in clock cycles)

| Value of register \$2 | CC 1 | CC 2 | CC 3 | CC 4 | CC 5 | CC 6 | CC 7 | CC 8 | CC 9 |
|-----------------------|------|------|------|------|------|------|------|------|------|
| | 10 | 10 | 10 | 10 | 10 | 10 | -20 | -20 | -20 |

Program execution order (in instructions)

sub \$2, \$1, \$3

and \$12, \$2, \$5

or \$13, \$6, \$2

add \$14, \$2, \$2

sw \$15, 100(\$2)



Chapter 4 — The Processor — 68

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



Chapter 4 — The Processor — 69

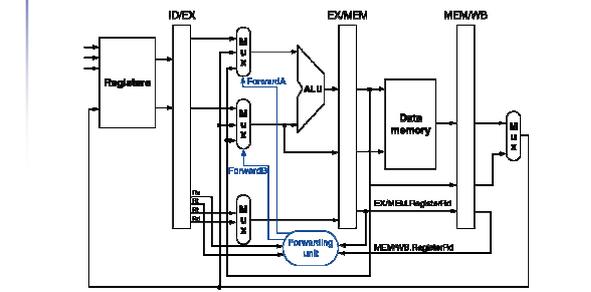
Detecting the Need to Forward

- 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 ≠ 0, MEM/WB.RegisterRd ≠ 0



Chapter 4 — The Processor — 70

Forwarding Paths



b. With forwarding



Chapter 4 — The Processor — 71

Forwarding Conditions

- EX hazard
 - if (EX/MEM.RegWrite and (EX/MEM.RegisterRd ≠ 0) and (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
- 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



Chapter 4 — The Processor — 72

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
- Revise MEM hazard condition
 - Only fwd if EX hazard condition isn't true



Chapter 4 — The Processor — 73

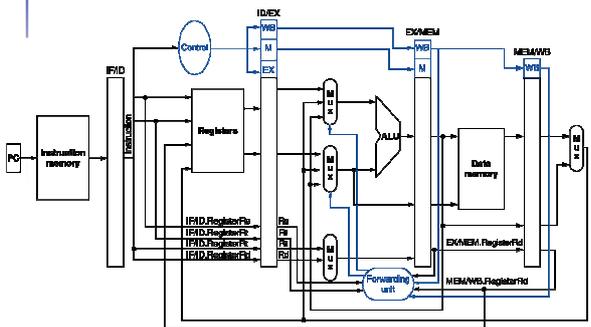
Revised Forwarding Condition

- MEM hazard
 - if (MEM/WB.RegWrite and (MEM/WB.RegisterRd ≠ 0) and not (EX/MEM.RegWrite and (EX/MEM.RegisterRd ≠ 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 ≠ 0) and not (EX/MEM.RegWrite and (EX/MEM.RegisterRd ≠ 0) and (EX/MEM.RegisterRd = ID/EX.RegisterRt)) and (MEM/WB.RegisterRd = ID/EX.RegisterRt)) ForwardB = 01



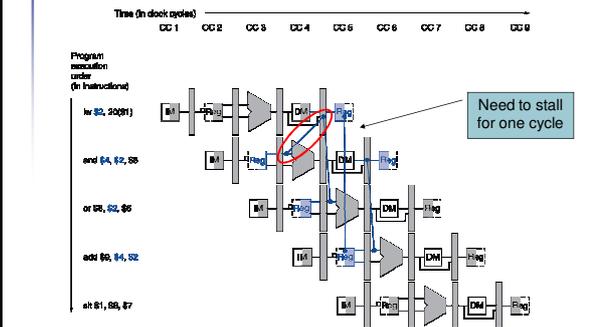
Chapter 4 — The Processor — 74

Datapath with Forwarding



Chapter 4 — The Processor — 75

Load-Use Data Hazard



Chapter 4 — The Processor — 76

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



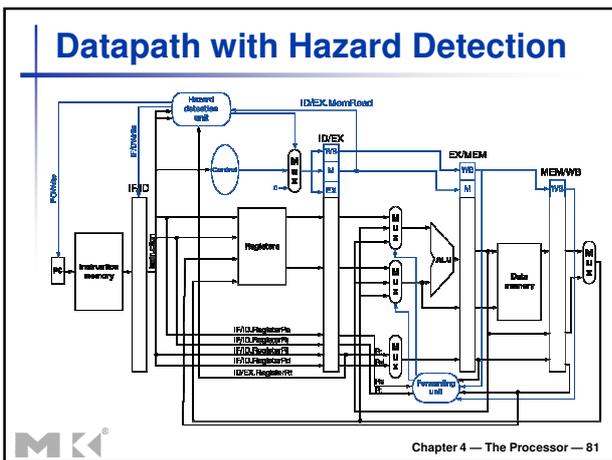
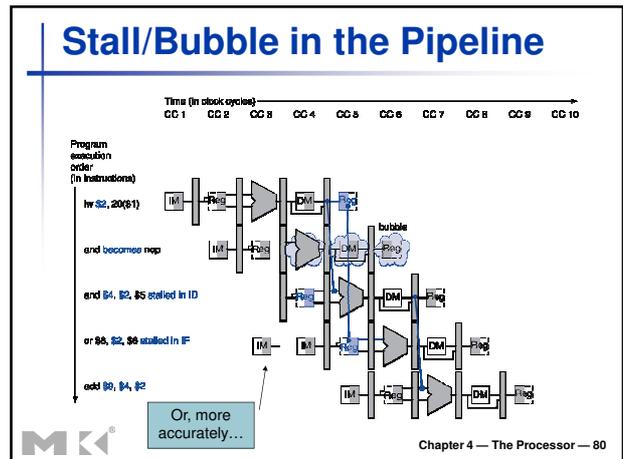
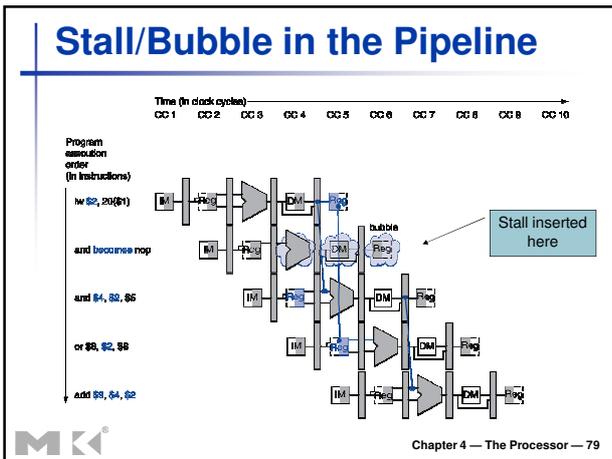
Chapter 4 — The Processor — 77

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
 - Following instruction is fetched again
 - 1-cycle stall allows MEM to read data for 1w
 - Can subsequently forward to EX stage



Chapter 4 — The Processor — 78

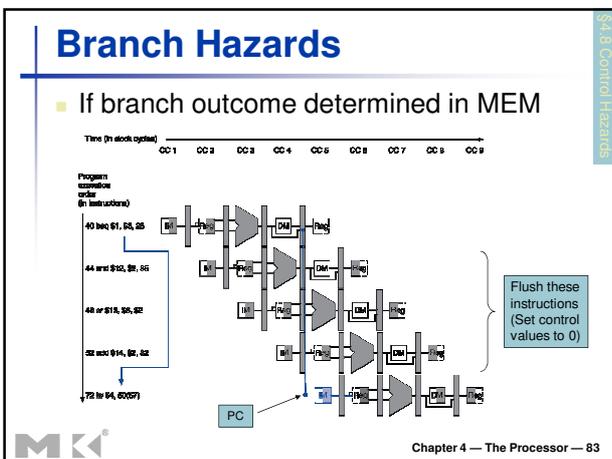


Stalls and Performance

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

Chapter 4 — The Processor — 82



Reducing Branch Delay

- Move hardware to determine outcome to ID stage
 - Target address adder
 - Register comparator
- Example: branch taken


```

36: sub $10, $4, $8
40: beq $1, $3, 7
44: and $12, $2, $5
48: or $13, $2, $6
52: add $14, $4, $2
56: slt $15, $6, $7
...
72: lw $4, 50($7)
            
```

Chapter 4 — The Processor — 84

1-Bit Predictor: Shortcoming

- Inner loop branches mispredicted twice!

```

outer: ...
inner: ...
    beq ..., ..., inner
    ...
    beq ..., ..., outer
    
```

- Mispredict as taken on last iteration of inner loop
- Then mispredict as not taken on first iteration of inner loop next time around

Chapter 4 — The Processor — 91

2-Bit Predictor

- Only change prediction on two successive mispredictions

Chapter 4 — The Processor — 92

Calculating the Branch Target

- Even with predictor, still need to calculate the target address
 - 1-cycle penalty for a taken branch
- Branch target buffer
 - Cache of target addresses
 - Indexed by PC when instruction fetched
 - If hit and instruction is branch predicted taken, can fetch target immediately

Chapter 4 — The Processor — 93

Exceptions and Interrupts

- “Unexpected” events requiring change in flow of control
 - Different ISAs use the terms differently
- Exception
 - Arises within the CPU
 - e.g., undefined opcode, overflow, syscall, ...
- Interrupt
 - From an external I/O controller
- Dealing with them without sacrificing performance is hard

Chapter 4 — The Processor — 94

Handling Exceptions

- In MIPS, exceptions managed by a System Control Coprocessor (COP0)
- Save PC of offending (or interrupted) instruction
 - In MIPS: Exception Program Counter (EPC)
- Save indication of the problem
 - In MIPS: Cause register
 - We'll assume 1-bit
 - 0 for undefined opcode, 1 for overflow
- Jump to handler at 8000 00180

Chapter 4 — The Processor — 95

An Alternate Mechanism

- Vectored Interrupts
 - Handler address determined by the cause
- Example:

| | |
|-------------------|-----------|
| Undefined opcode: | C000 0000 |
| Overflow: | C000 0020 |
| ...: | C000 0040 |
- Instructions either
 - Deal with the interrupt, or
 - Jump to real handler

Chapter 4 — The Processor — 96

Handler Actions

- Read cause, and transfer to relevant handler
- Determine action required
- If restartable
 - Take corrective action
 - use EPC to return to program
- Otherwise
 - Terminate program
 - Report error using EPC, cause, ...



Chapter 4 — The Processor — 97

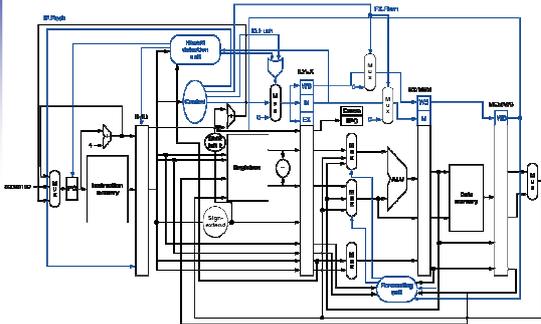
Exceptions in a Pipeline

- Another form of control hazard
- Consider overflow on add in EX stage
 - add \$1, \$2, \$1
 - Prevent \$1 from being clobbered
 - Complete previous instructions
 - Flush add and subsequent instructions
 - Set Cause and EPC register values
 - Transfer control to handler
- Similar to mispredicted branch
 - Use much of the same hardware



Chapter 4 — The Processor — 98

Pipeline with Exceptions



Chapter 4 — The Processor — 99

Exception Properties

- Restartable exceptions
 - Pipeline can flush the instruction
 - Handler executes, then returns to the instruction
 - Refetched and executed from scratch
- PC saved in EPC register
 - Identifies causing instruction
 - Actually PC + 4 is saved
 - Handler must adjust



Chapter 4 — The Processor — 100

Exception Example

- Exception on add in


```

40 sub $11, $2, $4
44 and $12, $2, $5
48 or $13, $2, $6
4C add $1, $2, $1
50 slt $15, $6, $7
54 lw $16, 50($7)
...
            
```
- Handler

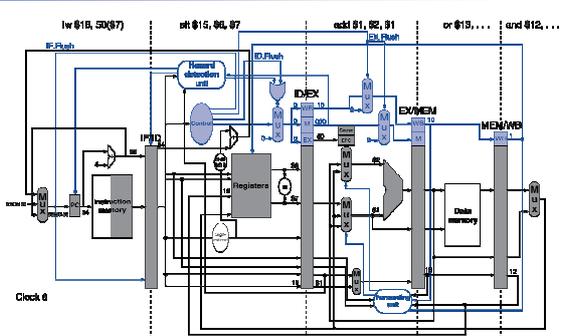

```

80000180 sw $25, 1000($0)
80000184 sw $26, 1004($0)
...
            
```



Chapter 4 — The Processor — 101

Exception Example



Chapter 4 — The Processor — 102

Compiler/Hardware Speculation

- Compiler can reorder instructions
 - e.g., move load before branch
 - Can include “fix-up” instructions to recover from incorrect guess
- Hardware can look ahead for instructions to execute
 - Buffer results until it determines they are actually needed
 - Flush buffers on incorrect speculation



Chapter 4 — The Processor — 109

Speculation and Exceptions

- What if exception occurs on a speculatively executed instruction?
 - e.g., speculative load before null-pointer check
- Static speculation
 - Can add ISA support for deferring exceptions
- Dynamic speculation
 - Can buffer exceptions until instruction completion (which may not occur)



Chapter 4 — The Processor — 110

Static Multiple Issue

- Compiler groups instructions into “issue packets”
 - Group of instructions that can be issued on a single cycle
 - Determined by pipeline resources required
- Think of an issue packet as a very long instruction
 - Specifies multiple concurrent operations
 - ⇒ Very Long Instruction Word (VLIW)



Chapter 4 — The Processor — 111

Scheduling Static Multiple Issue

- Compiler must remove some/all hazards
 - Reorder instructions into issue packets
 - No dependencies with a packet
 - Possibly some dependencies between packets
 - Varies between ISAs; compiler must know!
 - Pad with nop if necessary



Chapter 4 — The Processor — 112

MIPS with Static Dual Issue

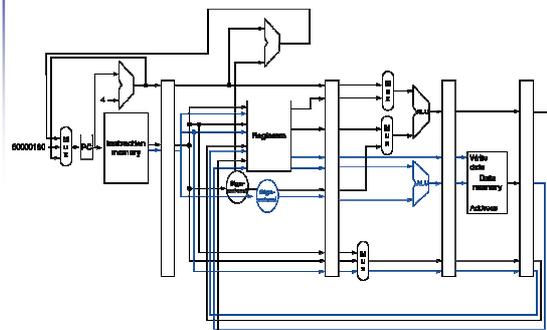
- Two-issue packets
 - One ALU/branch instruction
 - One load/store instruction
 - 64-bit aligned
 - ALU/branch, then load/store
 - Pad an unused instruction with nop

| Address | Instruction type | Pipeline Stages | | | | | |
|---------|------------------|-----------------|----|----|-----|-----|----|
| | | IF | ID | EX | MEM | WB | |
| n | ALU/branch | IF | ID | EX | MEM | WB | |
| n + 4 | Load/store | IF | ID | EX | MEM | WB | |
| n + 8 | ALU/branch | IF | ID | EX | MEM | WB | |
| n + 12 | Load/store | | IF | ID | EX | MEM | WB |
| n + 16 | ALU/branch | | IF | ID | EX | MEM | WB |
| n + 20 | Load/store | | IF | ID | EX | MEM | WB |



Chapter 4 — The Processor — 113

MIPS with Static Dual Issue



Chapter 4 — The Processor — 114

Hazards in the Dual-Issue MIPS

- More instructions executing in parallel
- EX data hazard
 - Forwarding avoided stalls with single-issue
 - Now can't use ALU result in load/store in same packet
 - add \$t0, \$s0, \$s1
 - load \$s2, 0(\$t0)
 - Split into two packets, effectively a stall
- Load-use hazard
 - Still one cycle use latency, but now two instructions
- More aggressive scheduling required



Scheduling Example

- Schedule this for dual-issue MIPS

```

Loop: lw  $t0, 0($s1)    # $t0=array element
      addu $t0, $t0, $s2 # add scalar in $s2
      sw  $t0, 0($s1)    # store result
      addi $s1, $s1, -4  # decrement pointer
      bne $s1, $zero, Loop # branch $s1!=0
    
```

| | ALU/branch | Load/store | cycle |
|-------|------------------------|------------------|-------|
| Loop: | nop | lw \$t0, 0(\$s1) | 1 |
| | addi \$s1, \$s1, -4 | nop | 2 |
| | addu \$t0, \$t0, \$s2 | nop | 3 |
| | bne \$s1, \$zero, Loop | sw \$t0, 4(\$s1) | 4 |

- $IPC = 5/4 = 1.25$ (c.f. peak $IPC = 2$)



Loop Unrolling

- Replicate loop body to expose more parallelism
 - Reduces loop-control overhead
- Use different registers per replication
 - Called "register renaming"
 - Avoid loop-carried "anti-dependencies"
 - Store followed by a load of the same register
 - Aka "name dependence"
 - Reuse of a register name



Loop Unrolling Example

| | ALU/branch | Load/store | cycle |
|-------|------------------------|-------------------|-------|
| Loop: | addi \$s1, \$s1, -16 | lw \$t0, 0(\$s1) | 1 |
| | nop | lw \$t1, 12(\$s1) | 2 |
| | addu \$t0, \$t0, \$s2 | lw \$t2, 8(\$s1) | 3 |
| | addu \$t1, \$t1, \$s2 | lw \$t3, 4(\$s1) | 4 |
| | addu \$t2, \$t2, \$s2 | sw \$t0, 16(\$s1) | 5 |
| | addu \$t3, \$t3, \$s2 | sw \$t1, 12(\$s1) | 6 |
| | nop | sw \$t2, 8(\$s1) | 7 |
| | bne \$s1, \$zero, Loop | sw \$t3, 4(\$s1) | 8 |

- $IPC = 14/8 = 1.75$
 - Closer to 2, but at cost of registers and code size



Dynamic Multiple Issue

- "Superscalar" processors
- CPU decides whether to issue 0, 1, 2, ... each cycle
 - Avoiding structural and data hazards
- Avoids the need for compiler scheduling
 - Though it may still help
 - Code semantics ensured by the CPU



Dynamic Pipeline Scheduling

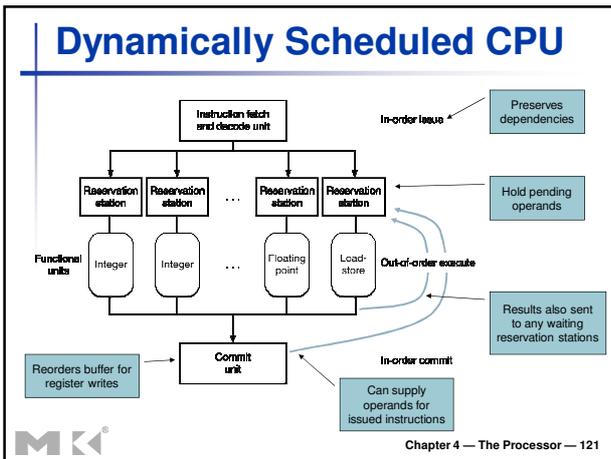
- Allow the CPU to execute instructions out of order to avoid stalls
 - But commit result to registers in order
- Example


```

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

 - Can start sub while addu is waiting for lw





- ### Register Renaming
- Reservation stations and reorder buffer effectively provide register renaming
 - On instruction issue to reservation station
 - If operand is available in register file or reorder buffer
 - Copied to reservation station
 - No longer required in the register; can be overwritten
 - If operand is not yet available
 - It will be provided to the reservation station by a function unit
 - Register update may not be required
- Chapter 4 — The Processor — 122

- ### Speculation
- Predict branch and continue issuing
 - Don't commit until branch outcome determined
 - Load speculation
 - Avoid load and cache miss delay
 - Predict the effective address
 - Predict loaded value
 - Load before completing outstanding stores
 - Bypass stored values to load unit
 - Don't commit load until speculation cleared
- Chapter 4 — The Processor — 123

- ### Why Do Dynamic Scheduling?
- Why not just let the compiler schedule code?
 - Not all stalls are predictable
 - e.g., cache misses
 - Can't always schedule around branches
 - Branch outcome is dynamically determined
 - Different implementations of an ISA have different latencies and hazards
- Chapter 4 — The Processor — 124

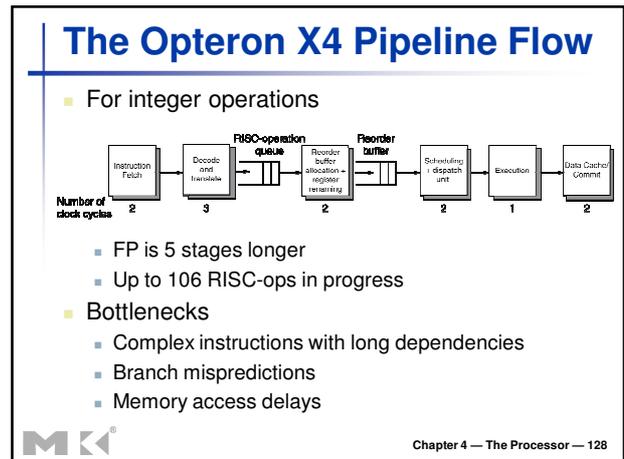
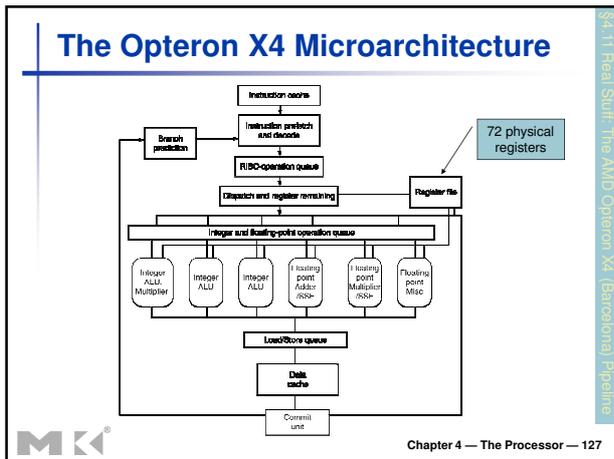
- ### Does Multiple Issue Work?
- The BIG Picture**
- Yes, but not as much as we'd like
 - Programs have real dependencies that limit ILP
 - Some dependencies are hard to eliminate
 - e.g., pointer aliasing
 - Some parallelism is hard to expose
 - Limited window size during instruction issue
 - Memory delays and limited bandwidth
 - Hard to keep pipelines full
 - Speculation can help if done well
- Chapter 4 — The Processor — 125

Power Efficiency

- Complexity of dynamic scheduling and speculations requires power
- Multiple simpler cores may be better

| Microprocessor | Year | Clock Rate | Pipeline Stages | Issue width | Out-of-order/Speculation | Cores | Power |
|----------------|------|------------|-----------------|-------------|--------------------------|-------|-------|
| i486 | 1989 | 25MHz | 5 | 1 | No | 1 | 5W |
| Pentium | 1993 | 66MHz | 5 | 2 | No | 1 | 10W |
| Pentium Pro | 1997 | 200MHz | 10 | 3 | Yes | 1 | 29W |
| P4 Willamette | 2001 | 2000MHz | 22 | 3 | Yes | 1 | 75W |
| P4 Prescott | 2004 | 3600MHz | 31 | 3 | Yes | 1 | 103W |
| Core | 2006 | 2930MHz | 14 | 4 | Yes | 2 | 75W |
| UltraSparc III | 2003 | 1950MHz | 14 | 4 | No | 1 | 90W |
| UltraSparc T1 | 2005 | 1200MHz | 6 | 1 | No | 8 | 70W |

Chapter 4 — The Processor — 126



- ### Fallacies
- Pipelining is easy (!)
 - The basic idea is easy
 - The devil is in the details
 - e.g., detecting data hazards
 - Pipelining is independent of technology
 - So why haven't we always done pipelining?
 - More transistors make more advanced techniques feasible
 - Pipeline-related ISA design needs to take account of technology trends
 - e.g., predicated instructions
- Chapter 4 — The Processor — 129

- ### Pitfalls
- Poor ISA design can make pipelining harder
 - e.g., complex instruction sets (VAX, IA-32)
 - Significant overhead to make pipelining work
 - IA-32 micro-op approach
 - e.g., complex addressing modes
 - Register update side effects, memory indirection
 - e.g., delayed branches
 - Advanced pipelines have long delay slots
- Chapter 4 — The Processor — 130

- ### Concluding Remarks
- ISA influences design of datapath and control
 - Datapath and control influence design of ISA
 - Pipelining improves instruction throughput using parallelism
 - More instructions completed per second
 - Latency for each instruction not reduced
 - Hazards: structural, data, control
 - Multiple issue and dynamic scheduling (ILP)
 - Dependencies limit achievable parallelism
 - Complexity leads to the power wall
- Chapter 4 — The Processor — 131