

CS4100: 計算機結構

Pipelining

國立清華大學資訊工程學系
九十三學年度第一學期

Adapted from Prof. D. Patterson's class notes
Copyright 1998, 2000 UCB

Outline

- ◆ An overview of pipelining
- ◆ A pipelined datapath
- ◆ Pipelined control
- ◆ Data hazards and forwarding
- ◆ Data hazards and stalls
- ◆ Branch hazards
- ◆ Exceptions
- ◆ Superscalar and dynamic pipelining

Pipelining Is Natural!

◆ Laundry example:

Ann, Brian, Cathy, Dave
each have one load of
clothes to wash, dry,
and fold

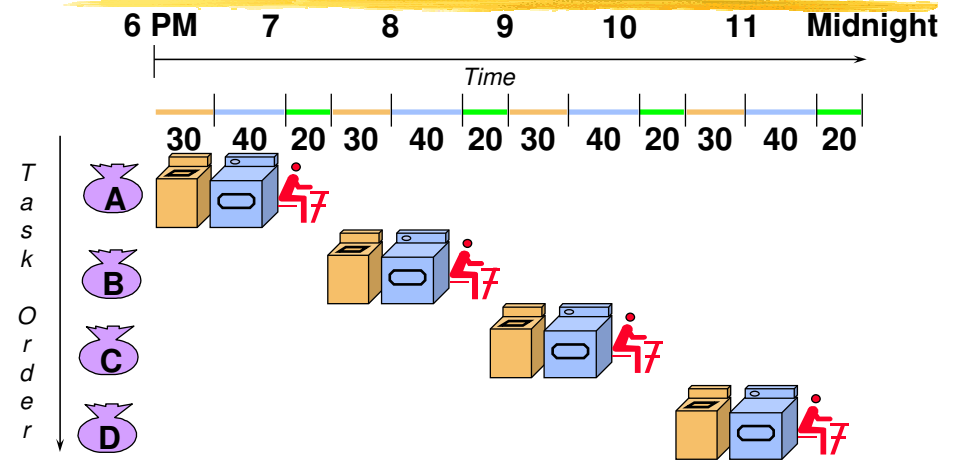
Washer takes 30 minutes

Dryer takes 40 minutes

“Folder” takes 20 minutes



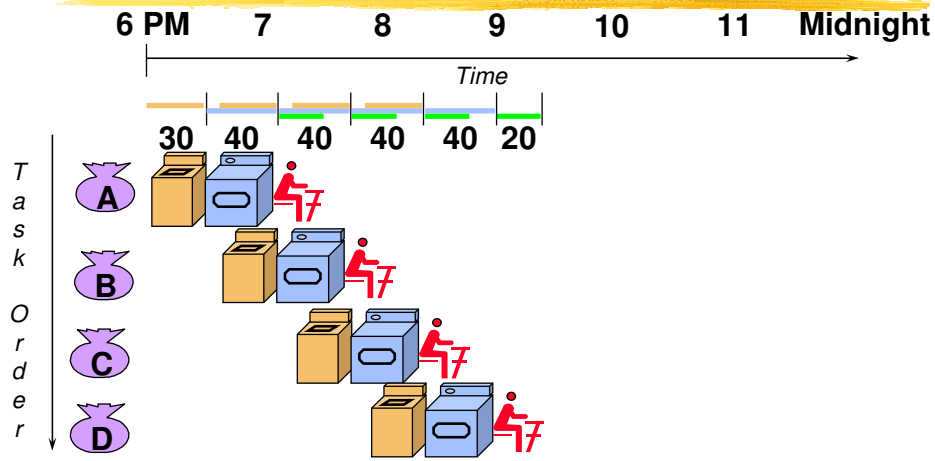
Sequential Laundry



◆ Sequential laundry takes 6 hours for 4 loads


◆ If they learned pipelining, how long would it take?

Pipelined Laundry: Start ASAP



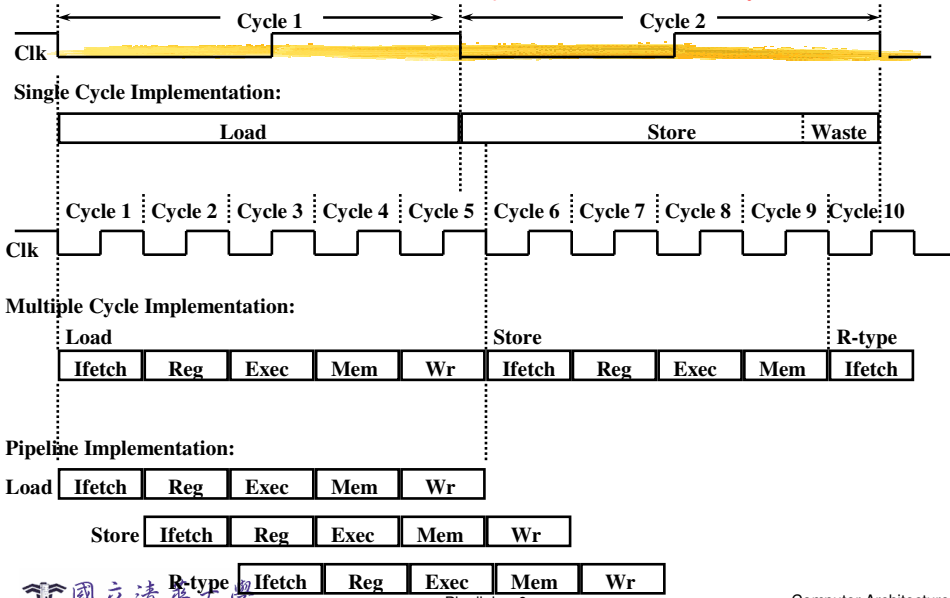
◆ Pipelined laundry takes 3.5 hours for 4 loads

Pipelining Lessons



- ◆ Doesn't help **latency** of single task, but **throughput** of entire
- ◆ Pipeline rate limited by **slowest** stage
- ◆ **Multiple** tasks working at same time using different resources
- ◆ Potential speedup = **Number pipe stages**
- ◆ Unbalanced stage length; time to "fill" & "drain" the pipeline reduce speedup
- ◆ Stall for dependences

Single-, Multi-Cycle, vs. Pipeline



Pipelining MIPS Execution

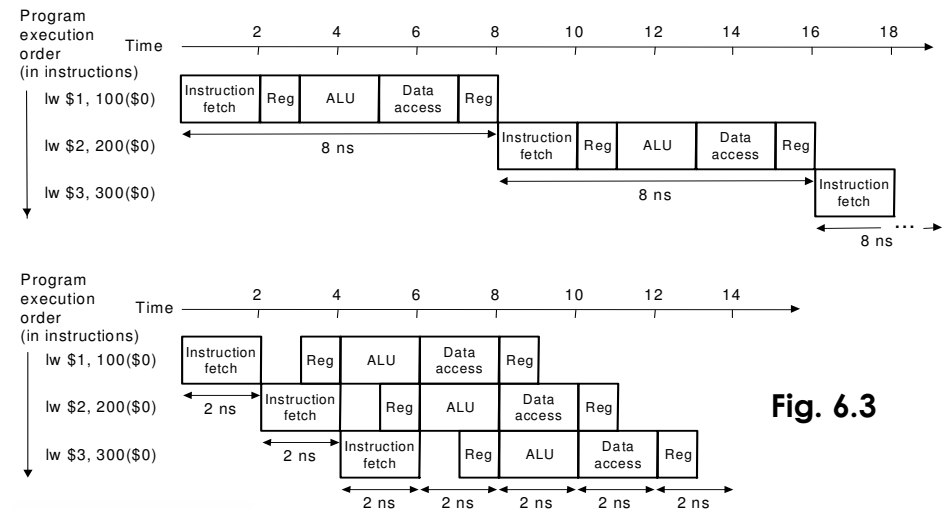
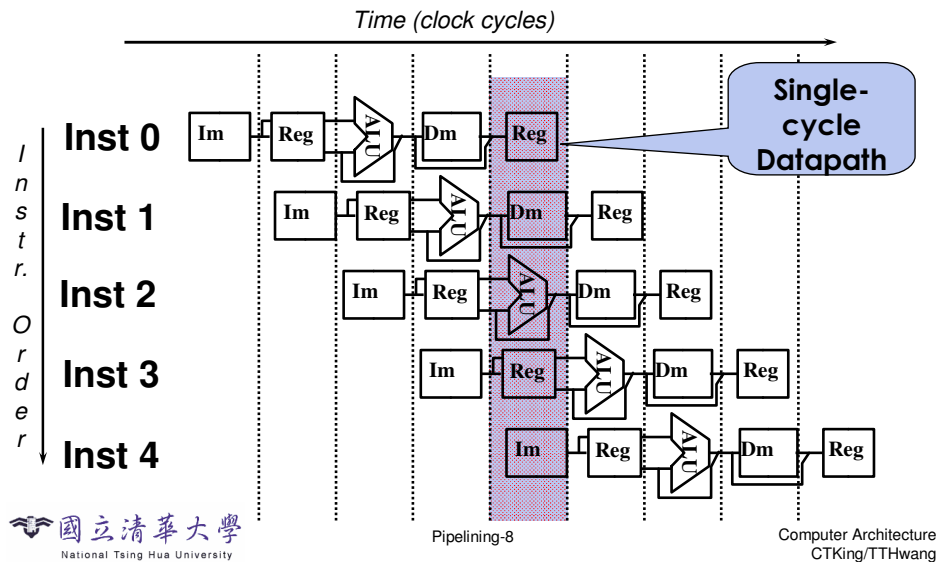


Fig. 6.3

Why Pipeline? Because the Resources Are There!



Outline

- ◆ An overview of pipelining
- ◆ A pipelined datapath
- ◆ Pipelined control
- ◆ Data hazards and forwarding
- ◆ Data hazards and stalls
- ◆ Branch hazards
- ◆ Exceptions
- ◆ Superscalar and dynamic pipelining

Designing a Pipelined Processor

- ◆ Examine the datapath and control diagram
 - Starting with single- or multi-cycle datapath?
 - Single- or multi-cycle control?
- ◆ Partition datapath into stages:
 - IF (instruction fetch), ID (instruction decode and register file read), EX (execution or address calculation), MEM (data memory access), WB (write back)
- ◆ Associate resources with states
- ◆ Ensure that flows do not conflict, or figure out how to resolve
- ◆ Assert control in appropriate stage

Use Multicycle Execution Steps

Step name	Action for R-type instructions	Action for memory-reference instructions	Action for branches	Action for jumps
Instruction fetch	IR = Memory[PC] PC = PC + 4			
Instruction decode/register fetch	A = Reg [IR[25-21]] B = Reg [IR[20-16]] ALUOut = PC + (sign-extend (IR[15-0]) << 2)			
Execution, address computation, branch/jump completion	ALUOut = A op B	ALUOut = A + sign-extend (IR[15-0])	if (A == B) then PC = ALUOut	PC = PC [31-28] (IR[25-0] << 2)
Memory access or R-type completion	Reg [IR[15-11]] = ALUOut	Load: MDR = Memory[ALUOut] or Store: Memory [ALUOut] = B		
Memory read completion		Load: Reg[IR[20-16]] = MDR		

But, use single-cycle datapath ...

Split Single-cycle Datapath

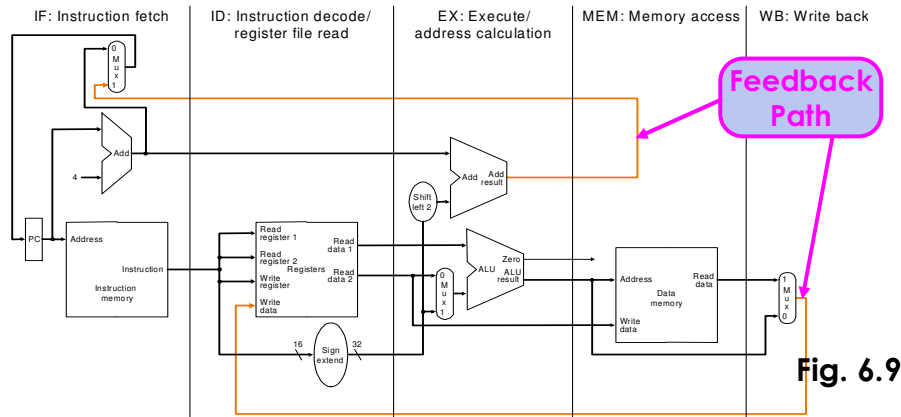


Fig. 6.9

What to add to split the datapath into stages?

Add Pipeline Registers

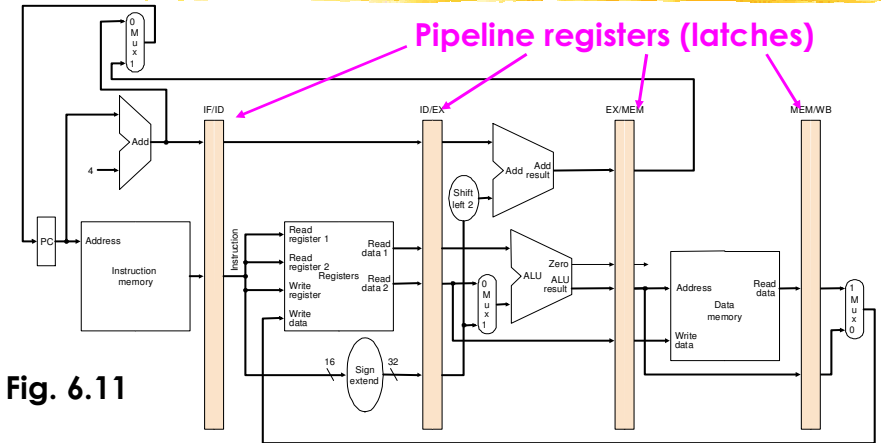



Fig. 6.11


Use registers between stages to carry data and control

Consider load



- ◆ IF: Instruction Fetch
 - Fetch the instruction from the Instruction Memory
- ◆ ID: Instruction Decode
 - Registers fetch and instruction decode
- ◆ EX: Calculate the memory address
- ◆ MEM: Read the data from the Data Memory
- ◆ WB: Write the data back to the register file

Pipelining load



- ◆ 5 functional units in the pipeline datapath are:
 - Instruction Memory for the Ifetch stage
 - Register File's Read ports (busA and busB) for the Reg/Dec stage
 - ALU for the Exec stage
 - Data Memory for the MEM stage
 - Register File's Write port (busW) for the WB stage

IF Stage of load

◆ $IR = mem[PC]; PC = PC + 4$

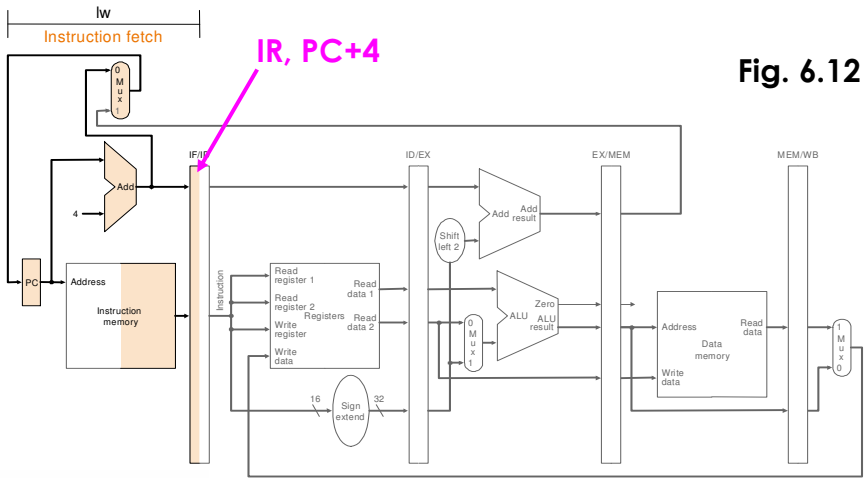


Fig. 6.12

ID Stage of load

◆ $A = Reg[IR[25-21]]; B = Reg[IR[20-16]]; ALUout = PC + (sign-ext(IR[15-0]) \ll 2)$

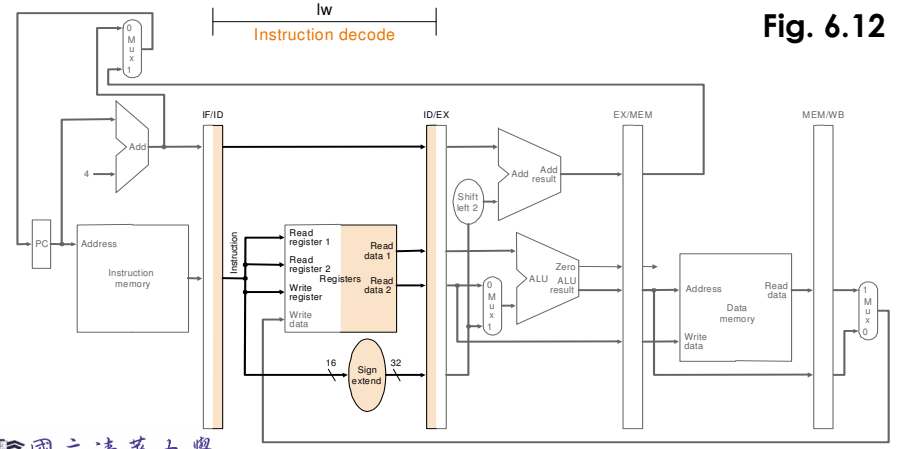


Fig. 6.12

EX Stage of load

◆ $ALUout = A + sign-ext(IR[15-0])$

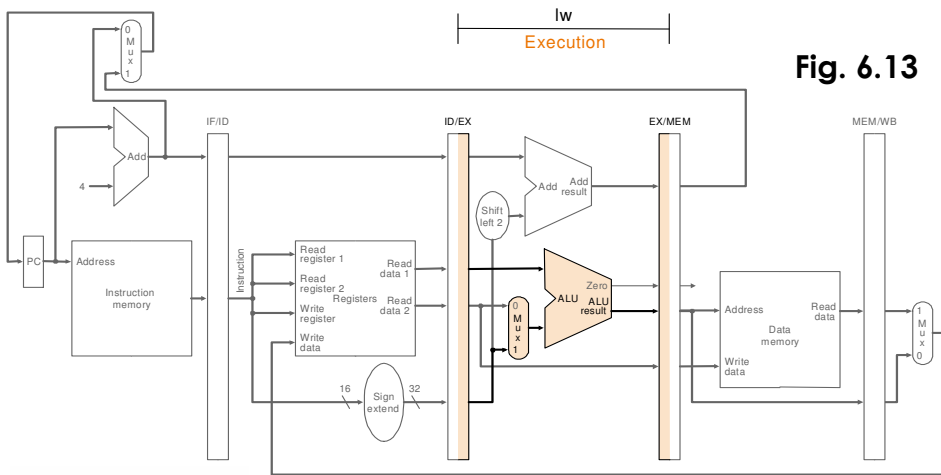


Fig. 6.13

MEM State of load

◆ $MDR = mem[ALUout]$

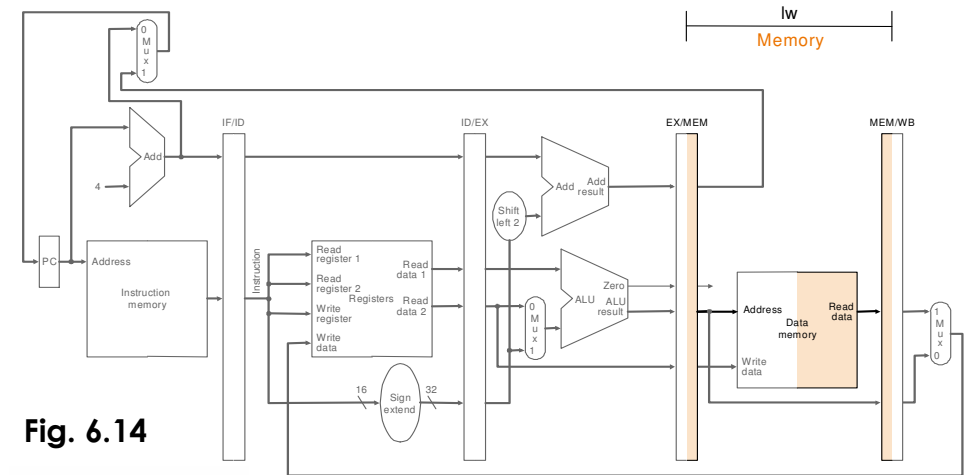


Fig. 6.14

WB Stage of load

◆ $Reg[IR[20-16]] = MDR$

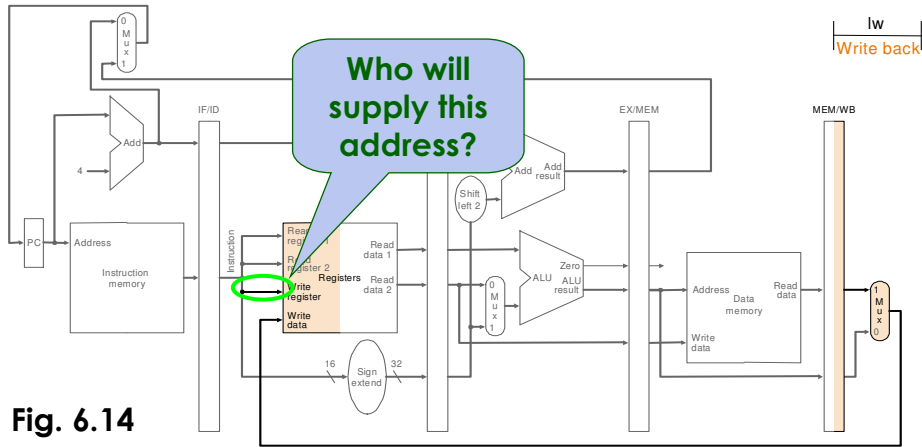



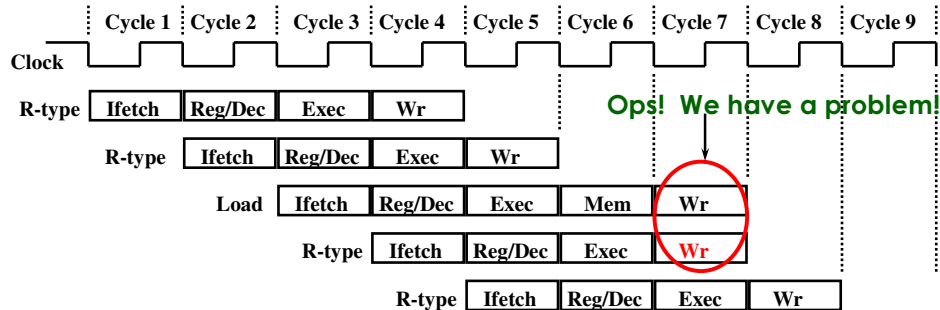
Fig. 6.14

The Four Stages of R-type



- ◆ IF: fetch the instruction from the Instruction Memory
- ◆ ID: registers fetch and instruction decode
- ◆ EX: ALU operates on the two register operands
- ◆ WB: write ALU output back to the register file

Pipelining R-type and load




◆ We have a **structural hazard**:

- Two instructions try to write to the register file at the same time!
- Only one write port


Important 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



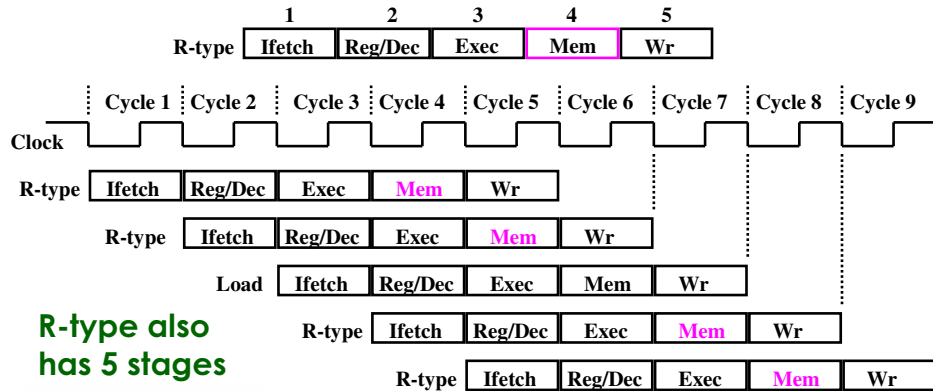
- R-type uses Register File's write port during its **4th** stage



Several ways to solve: forwarding, adding pipeline bubble, making instructions same length


Solution: Delay R-type's Write

- ◆ Delay R-type's register write by one cycle:
 - R-type also use Reg File's write port at Stage 5
 - MEM is a NOP stage: nothing is being done.



R-type also has 5 stages

The Four Stages of store




- ◆ IF: fetch the instruction from the Instruction Memory
- ◆ ID: registers fetch and instruction decode
- ◆ EX: calculate the memory address
- ◆ MEM: write the data into the Data Memory

Add an extra stage:

- ◆ WB: NOP

The Three Stages of beq

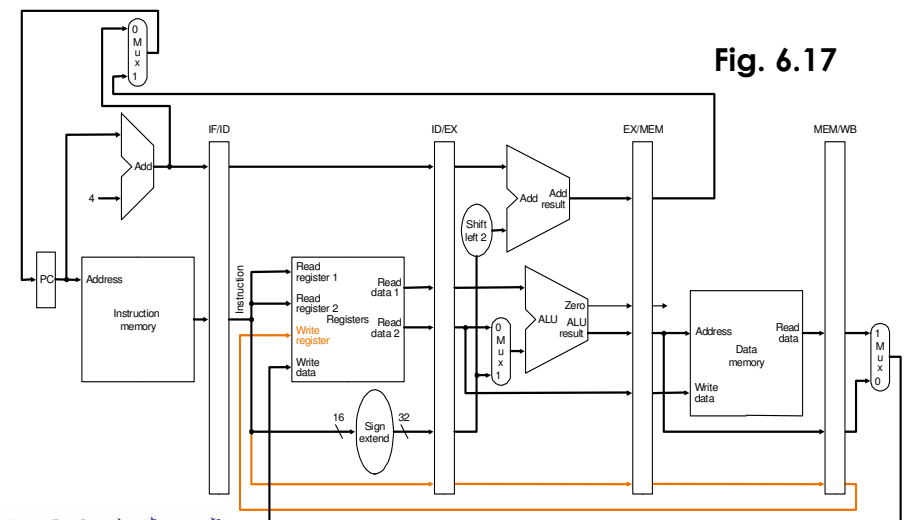


- ◆ IF: fetch the instruction from the Instruction Memory
- ◆ ID: registers fetch and instruction decode
- ◆ EX:
 - compares the two register operand
 - select correct branch target address
 - latch into PC

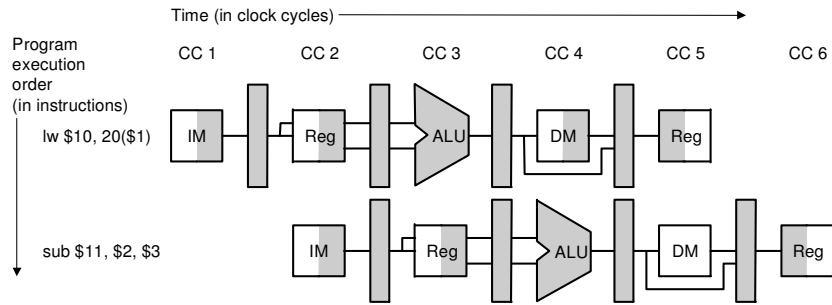
Add two extra stages:

- ◆ MEM: NOP
- ◆ WB: NOP

Pipelined Datapath



Graphically Representing Pipelines



- ◆ Can help with answering questions like:
 - How many cycles to execute this code?
 - What is the ALU doing during cycle 4?
 - Help understand datapaths

Example 1: Cycle 1

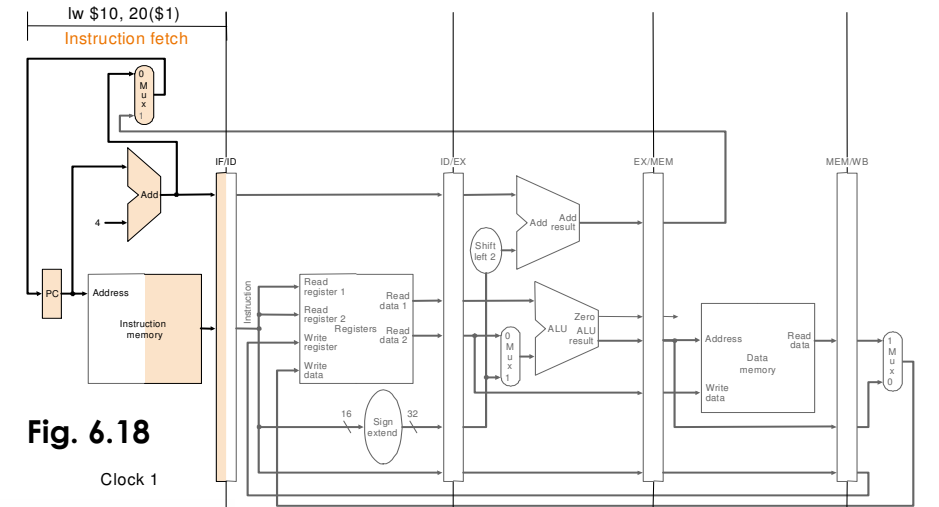


Fig. 6.18

Example 1: Cycle 2




Fig. 6.18

Example 1: Cycle 3



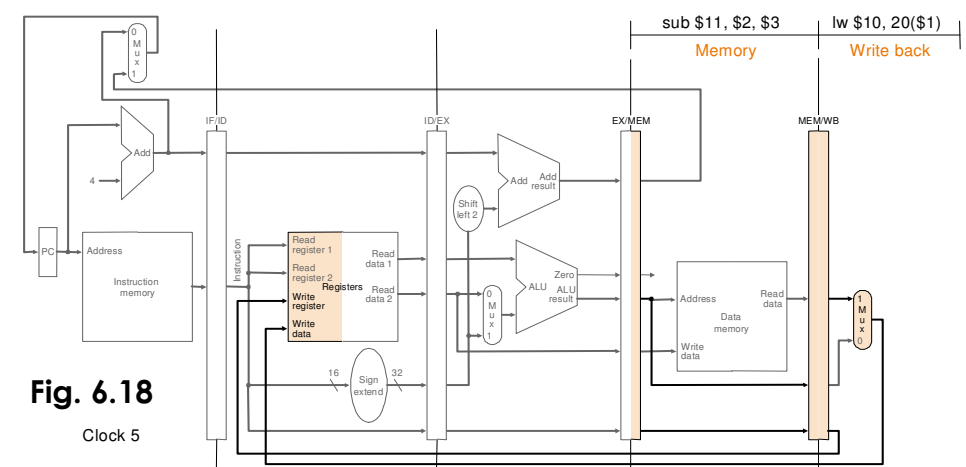


Fig. 6.18

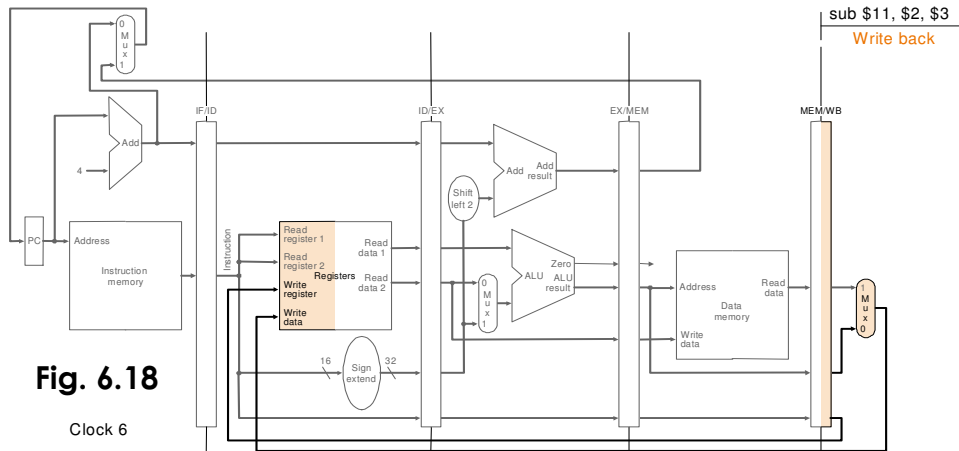
Example 1: Cycle 4



Example 1: Cycle 5



Example 1: Cycle 6



Outline

- ◆ An overview of pipelining
- ◆ A pipelined datapath
- ◆ **Pipelined control**
- ◆ Data hazards and forwarding
- ◆ Data hazards and stalls
- ◆ Branch hazards
- ◆ Exceptions
- ◆ Superscalar and dynamic pipelining

Pipeline Control: Control Signals

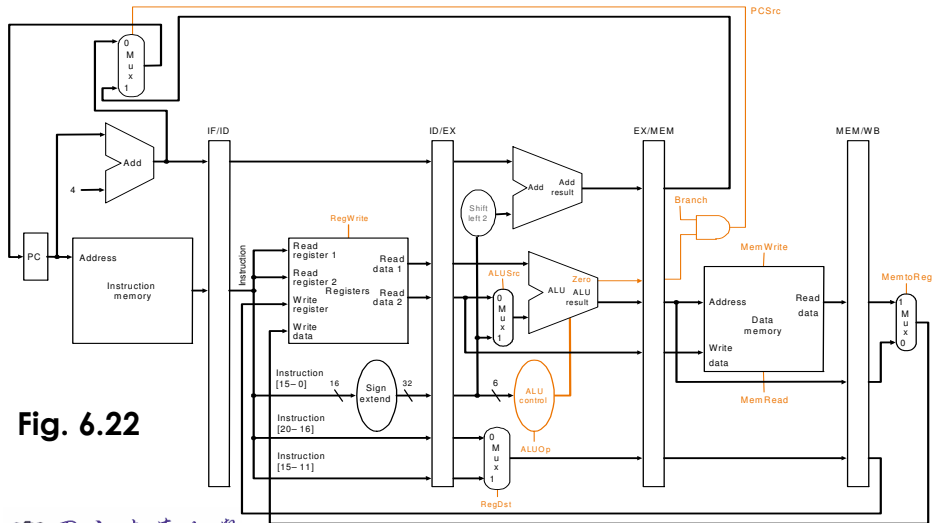


Fig. 6.22

Group Signals According to Stages

- Can use control signals of single-cycle CPU (Fig. 6.23, 6.24 \Leftrightarrow 5.12, 5.16)

Reg Dst	Execution/Address Calculation stage control lines			Memory access stage control lines			Write-back stage control lines	
	ALU Op1	ALU Op0	ALU Src	Branch	Mem Read	Mem Write	Reg write	Mem to Reg
1	1	0	0	0	0	0	1	0
0	0	0	1	0	1	0	1	1
X	0	0	1	0	0	1	0	X
X	0	1	0	1	0	0	0	X

Fig. 6.25

Data Stationary Control

- Pass control signals along just like the data
- Main control generates control signals during ID


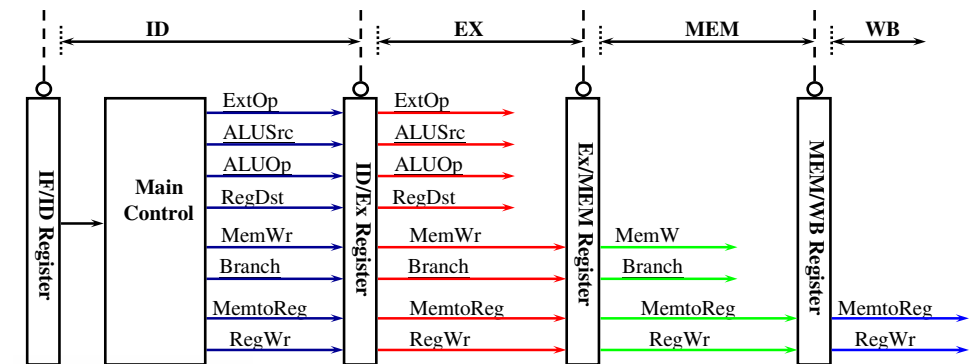


Fig. 6.26

Data Stationary Control (cont.)

- Signals for EX (ExtOp, ALUSrc, ...) are used 1 cycle later
- Signals for MEM (MemWr, Branch) are used 2 cycles later
- Signals for WB (MemtoReg, MemWr) are used 3 cycles later



Datapath with Control

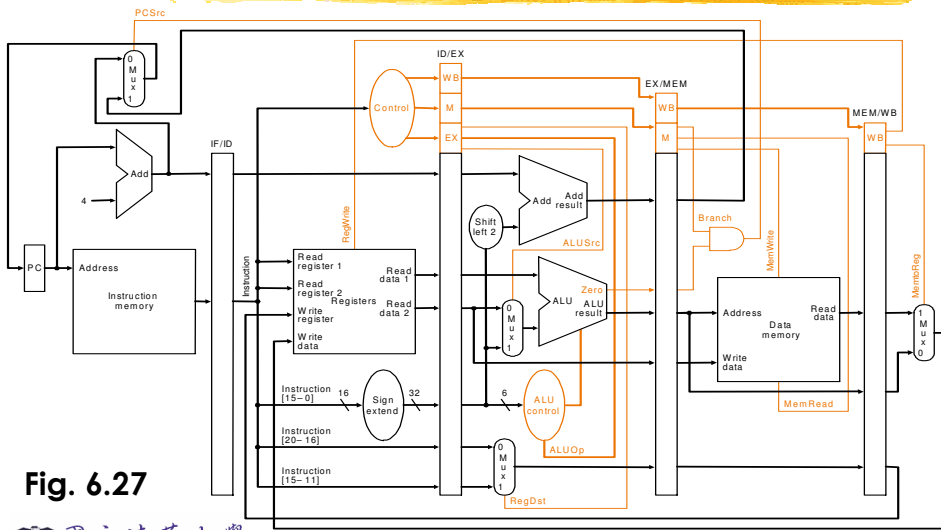
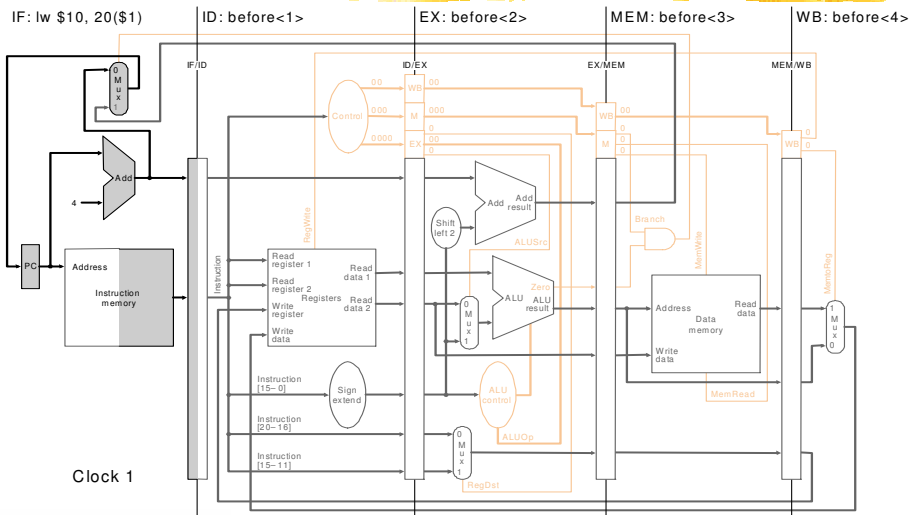


Fig. 6.27

Let's Try it Out


```
lw $t0, 20($t1)
sub $t1, $t2, $t3
and $t2, $t4, $t5
or $t3, $t6, $t7
add $t4, $t8, $t9
```

Example 2: Cycle 1




Clock 1

Example 2: Cycle 2




Clock 2

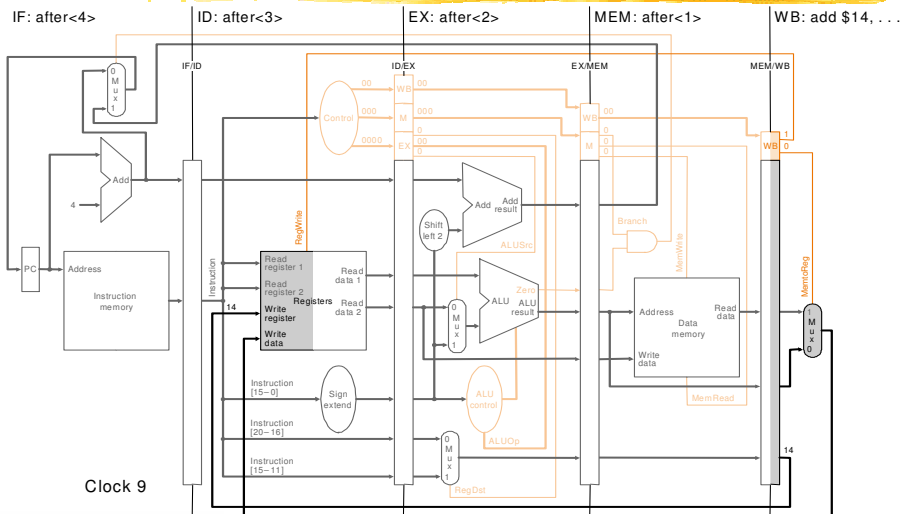
Example 2: Cycle 7



Example 2: Cycle 8



Example 2: Cycle 9



Summary of Pipeline Basics

- ◆ Pipelining is a fundamental concept
 - Multiple steps using distinct resources
 - Utilize capabilities of datapath by pipelined instruction processing
 - Start next instruction while working on the current one
 - Limited by length of longest stage (plus fill/flush)
 - Need to detect and resolve hazards
- ◆ What makes it easy in MIPS?
 - All instructions are of the same length
 - Just a few instruction formats
 - Memory operands only in loads and stores
- ◆ What makes it hard? **hazards**

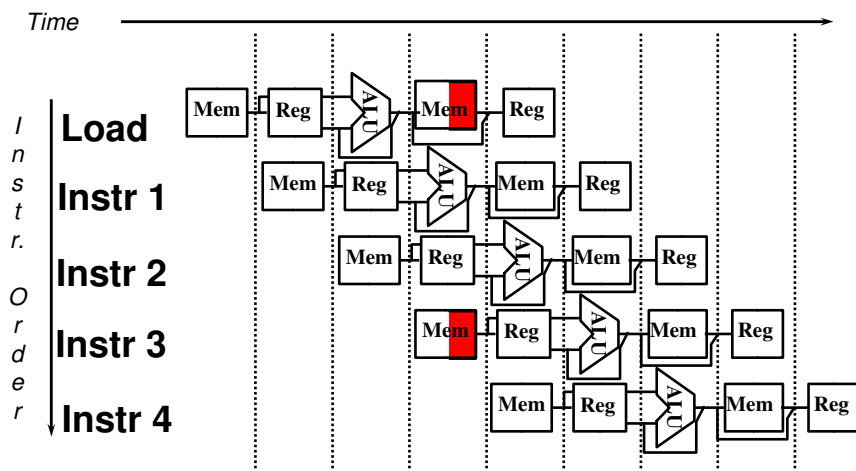
Outline

- ◆ An overview of pipelining
- ◆ A pipelined datapath
- ◆ Pipelined control
- ◆ Data hazards and forwarding
- ◆ Data hazards and stalls
- ◆ Branch hazards
- ◆ Exceptions
- ◆ Superscalar and dynamic pipelining

Pipeline Hazards

- ◆ Pipeline Hazards:
 - **Structural hazards:** attempt to use the same resource in two different ways at the same time
 - Ex.: combined washer/dryer or folder busy doing something else (watching TV)
 - **Data hazards:** attempt to use item before ready
 - Instruction depends on result of prior instruction still in the pipeline
 - **Control hazards:** attempt to make decision before condition is evaluated
 - Ex.: wash football uniforms and need to see result of previous load to get proper detergent level
 - Branch instructions
- ◆ Can always resolve hazards by **waiting**
 - pipeline control must detect the hazard
 - take action (or delay action) to resolve hazards

Structural Hazard: Single Memory



Use 2 memory: data memory and instruction memory

Data Hazards

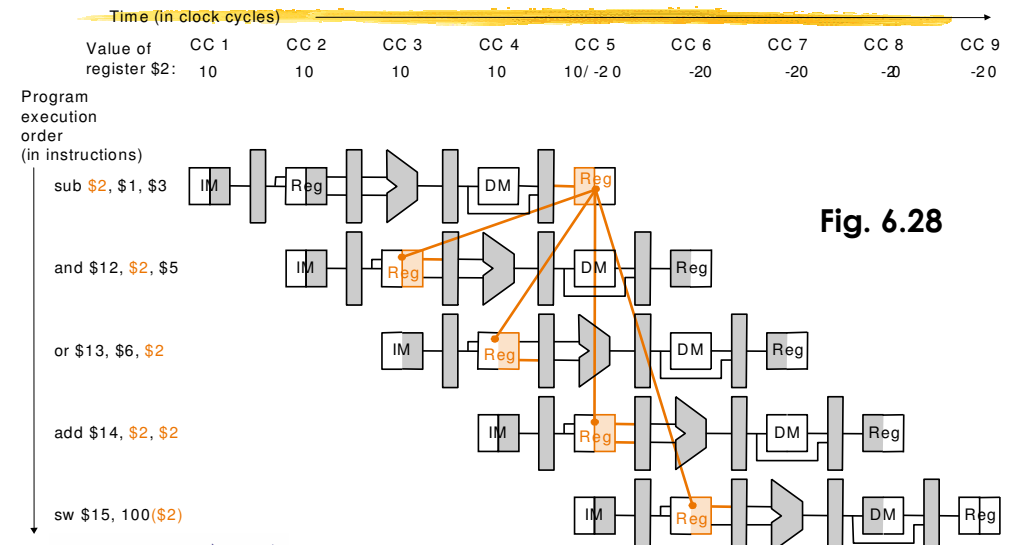


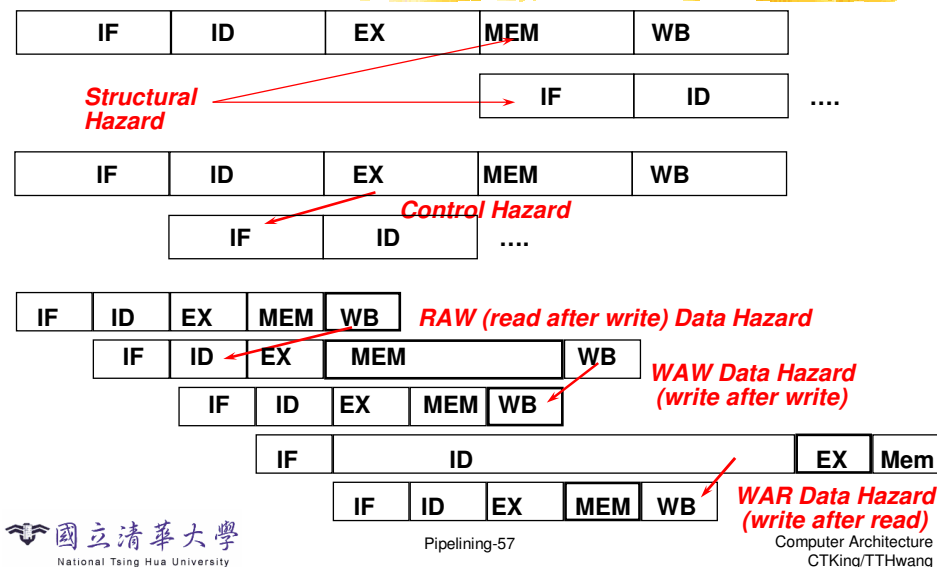
Fig. 6.28

Types of Data Hazards

Three types: (inst. i1 followed by inst. i2)

- ◆ **RAW (read after write):**
i2 tries to read operand before i1 writes it
- ◆ **WAR (write after read):**
i2 tries to write operand before i1 reads it
 - Gets wrong operand, e.g., autoincrement addr.
 - Can't happen in MIPS 5-stage pipeline because:
 - All instructions take 5 stages, and reads are always in stage 2, and writes are always in stage 5
- ◆ **WAW (write after write):**
i2 tries to write operand before i1 writes it
 - Leaves wrong result (i1's not i2's); occur only in pipelines that write in more than one stage
 - Can't happen in MIPS 5-stage pipeline because:
 - All instructions take 5 stages, and writes are always in stage 5

Pipeline Hazards Illustrated



Handling Data Hazards

- ◆ Use simple, fixed designs
 - Eliminate WAR by always fetching operands early (ID) in pipeline
 - Eliminate WAW by doing all write backs in order (last stage, static)
 - These features have a lot to do with ISA design
- ◆ Internal forwarding in register file:
 - Write in first half of clock and read in second half
 - Read delivers what is written, resolve hazard between sub and add
- ◆ **Detect and resolve** remaining ones
 - Compiler inserts NOP
 - Stall
 - Forward

Software Solution

- ◆ Have compiler guarantee no hazards
- ◆ Where do we insert the NOPs?

```

sub   $2, $1, $3
and   $12, $2, $5
or    $13, $6, $2
add   $14, $2, $2
sw    $15, 100($2)
    
```

- ◆ Problem: this really slows us down!

Data Hazards

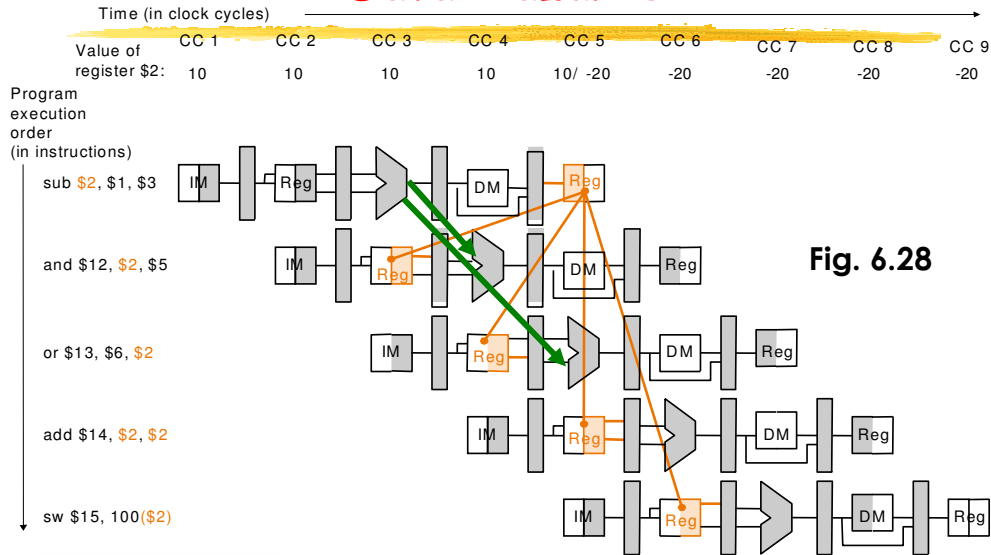


Fig. 6.28

Detecting Data Hazards

- ◆ Hazard conditions:
 - 1a. EX/MEM.RegWrite = ID/EX.RegWrite
 - 1b. EX/MEM.RegWrite = ID/EX.RegWrite
 - 2a. MEM/WB.RegWrite = ID/EX.RegWrite
 - 2b. MEM/WB.RegWrite = ID/EX.RegWrite
- ◆ Two optimizations:
 - Don't forward if instruction does not write register => check if RegWrite is asserted
 - Don't forward if destination register is \$0 => check if RegisterRd = 0

Detecting Data Hazards (cont.)

- ◆ Hazard conditions using control signals:
 - At EX stage:
EX/MEM.RegWrite and (EX/MEM.RegRd≠0)
and (EX/MEM.RegRd=ID/EX.RegRs)
 - At MEM stage:
MEM/WB.RegWrite and (MEM/WB.RegRd≠0)
and (MEM/WB.RegRd=ID/EX.RegRs)
 - (replace ID/EX.RegRt for ID/EX.RegRs for the other two conditions)

Resolving Hazards: Forwarding

- ◆ Use temporary results, e.g., those in pipeline registers, don't wait for them to be written

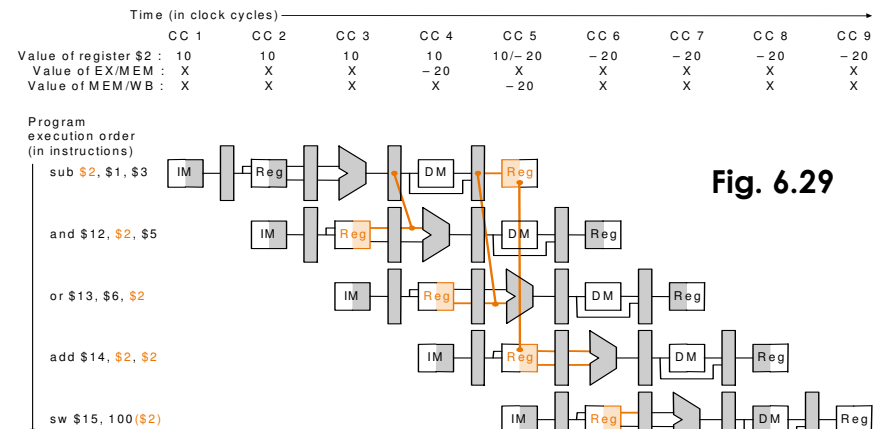
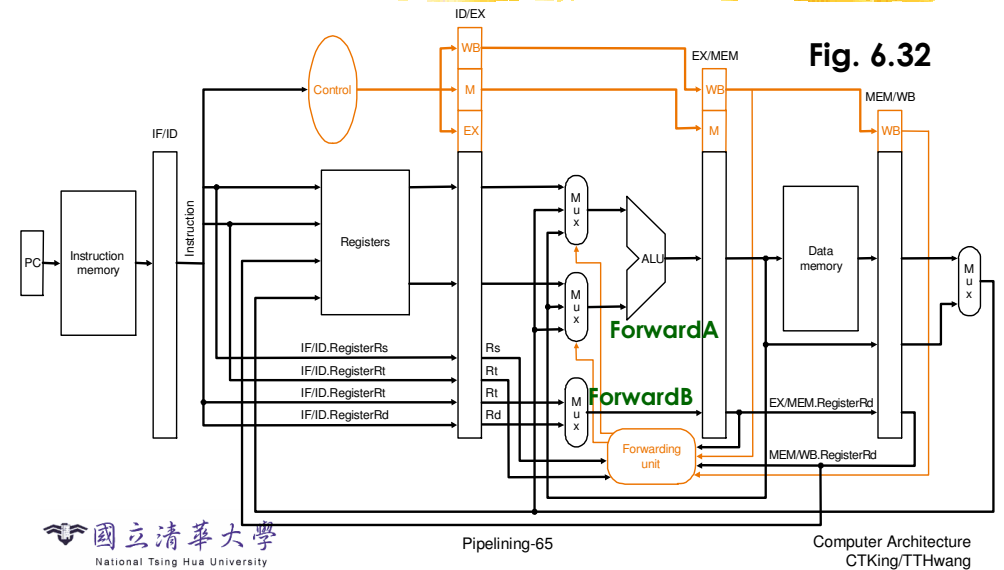


Fig. 6.29


Forwarding Logic

- ◆ Forwarding: input to ALU from any pipe reg.
 - Add multiplexors to ALU input
 - Control forwarding in EX => carry Rs in ID/EX
 - ◆ Control signals for forwarding:
 - If both WB and MEM forward, e.g., add \$1, \$1, \$2; add \$1, \$1, \$3; add \$1, \$1, \$4; => let MEM forward
 - EX hazard:
 - if (EX/MEM.RegWrite and (EX/MEM.RegRd≠0) and (EX/MEM.RegRd=ID/EX.RegRs)) ForwardA=10
 - MEM hazard:
 - if (MEM/WB.RegWrite and (MEM/WB.RegRd≠0) and (EX/MEM.RegRd ≠ ID/EX.Reg.Rs) and (MEM/WB.RegRd=ID/EX.RegRs)) ForwardA=01
- (ID/EX.RegRt<->ID/EX.RegRs, ForwardB<-> ForwardA)


Pipeline with Forwarding




Example 3: Cycle 3



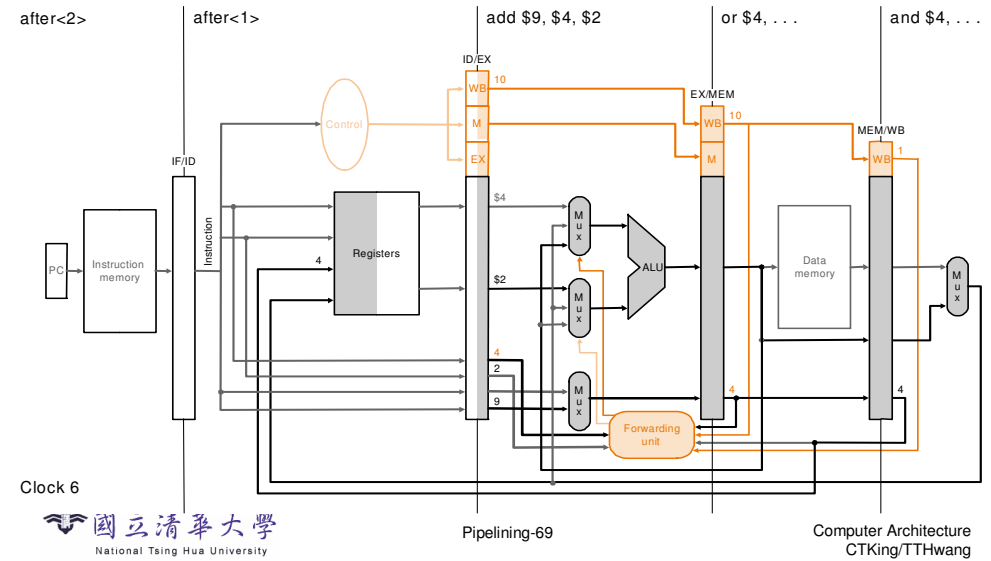
Example 3: Cycle 4



Example 3: Cycle 5

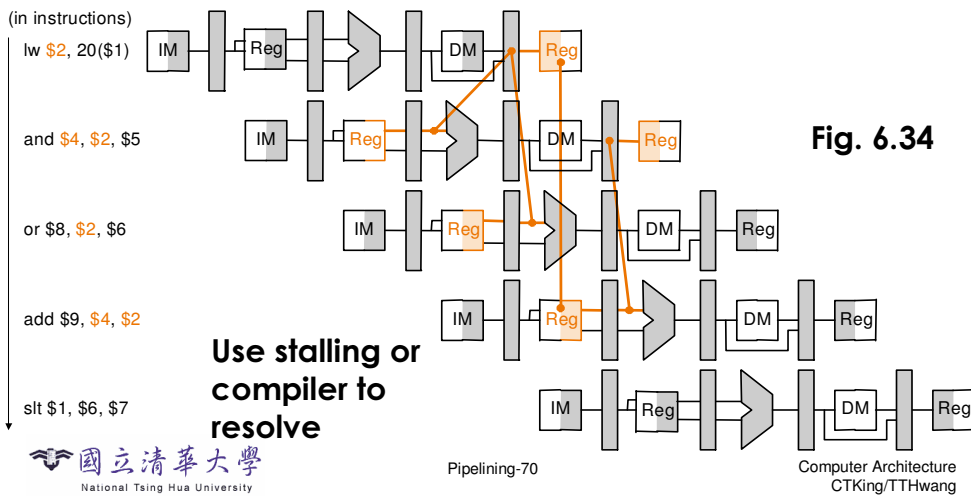


Example 3: Cycle 6



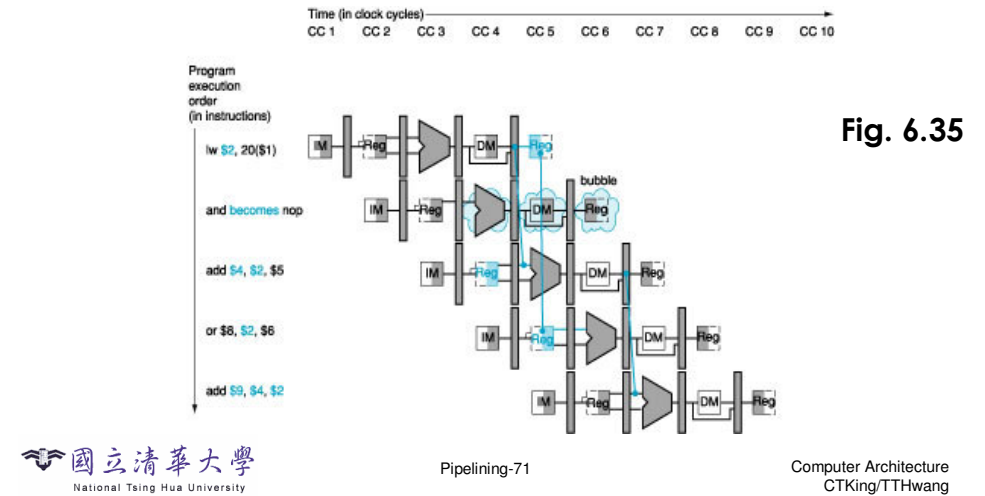
Can't Always Forward

- ◆ **lw can still cause a hazard:**
 - if it is followed by an instruction to read the loaded reg.



Stalling

- ◆ **Stall pipeline by keeping instructions in same stage and inserting a NOP instead**



Handling Stalls

- ◆ Hazard detection unit in ID to insert stall between a load instruction and its use:

```

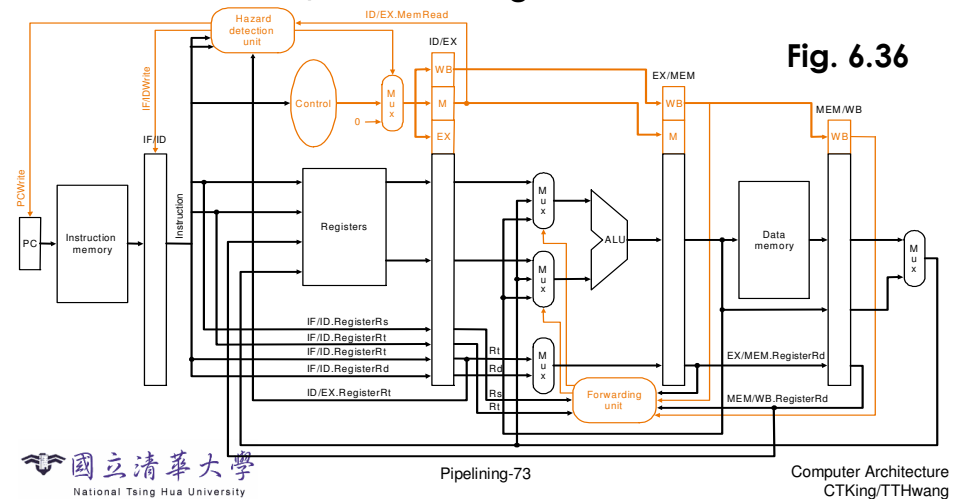
if (ID/EX.MemRead and
    ((ID/EX.RegisterRt = IF/ID.RegisterRs) or
     (ID/EX.RegisterRt = IF/ID.registerRt))
    stall the pipeline for one cycle
(ID/EX.MemRead=1 indicates a load instruction)
    
```

- ◆ How to stall?


- Stall instruction in IF and ID: not change PC and IF/ID => the stages re-execute the instructions
- What to move into EX: insert an NOP by changing EX, MEM, WB control fields of ID/EX pipeline register to 0
 - as control signals propagate, all control signals to EX, MEM, WB are deasserted and no registers or memories are written

Pipeline with Stalling Unit

- ◆ Forwarding controls ALU inputs, hazard detection controls PC, IF/ID, control signals




Example 4: Cycle 2




Clock 2

Example 4: Cycle 3




Clock 3

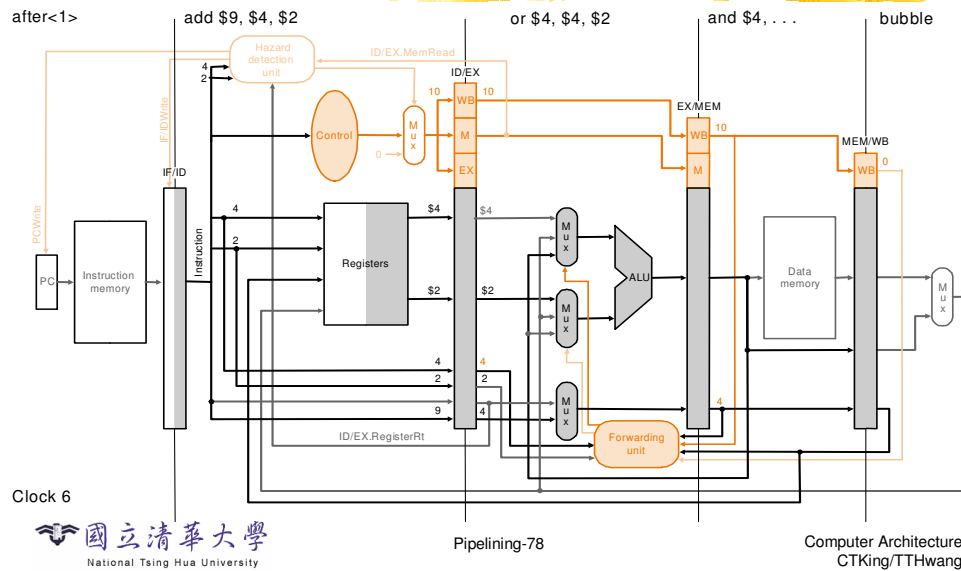
Example 4: Cycle 4




Example 4: Cycle 5



Example 4: Cycle 6



Example 4: Cycle 7



Outline

- ◆ An overview of pipelining
- ◆ A pipelined datapath
- ◆ Pipelined control
- ◆ Data hazards and forwarding
- ◆ Data hazards and stalls
- ◆ **Branch hazards**
- ◆ Exceptions
- ◆ Superscalar and dynamic pipelining

Pipeline Datapath with Control Signals

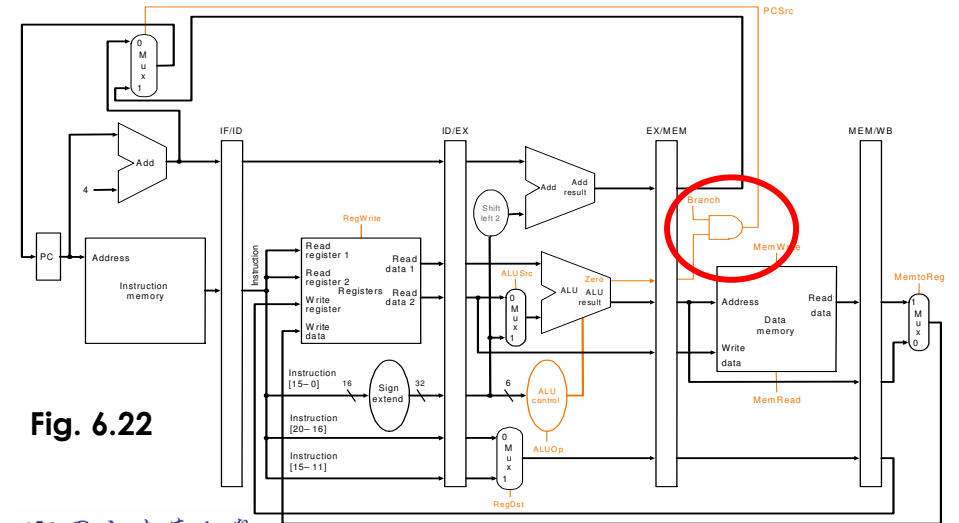


Fig. 6.22

Branch Hazards

- ◆ When decide to branch, other inst. are in pipeline!

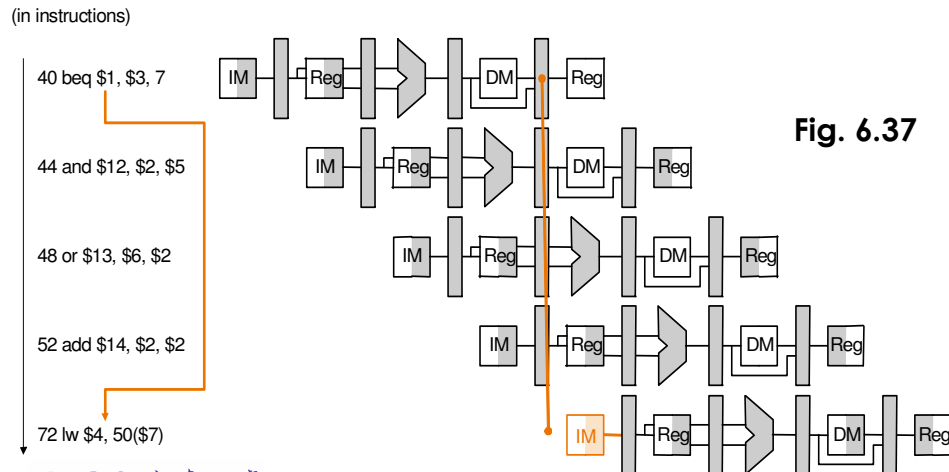


Fig. 6.37

Handling Branch Hazard

- ◆ Predict branch always not taken
 - Need to add hardware for flushing inst. if wrong
 - Branch decision made at MEM => need to flush instruction in IF, ID, EX by changing control values to 0
- ◆ Reduce delay of taken branch by moving branch execution earlier in the pipeline
 - Move up branch address calculation to ID
 - Check branch equality at ID (using XOR) by comparing the two registers read during ID
 - Branch decision made at EX => one instruction to flush
 - Add a control signal, IF.Flush, to zero instruction field of IF/ID => making the instruction a NOP
- ◆ Dynamic branch prediction
- ◆ Compiler rescheduling, delay branch

Pipeline with Flushing

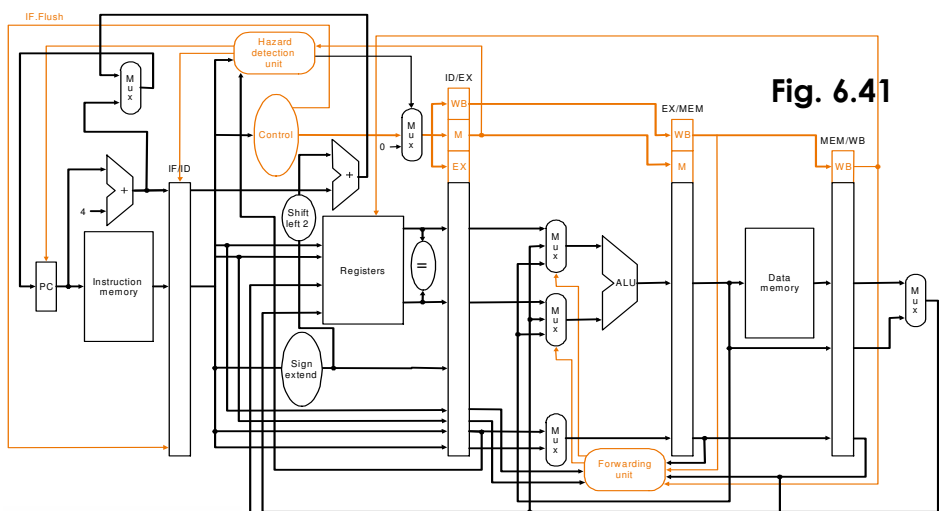


Fig. 6.41

Example 5: Cycle 3




Fig. 6.38

Clock 3

Example 5: Cycle 4


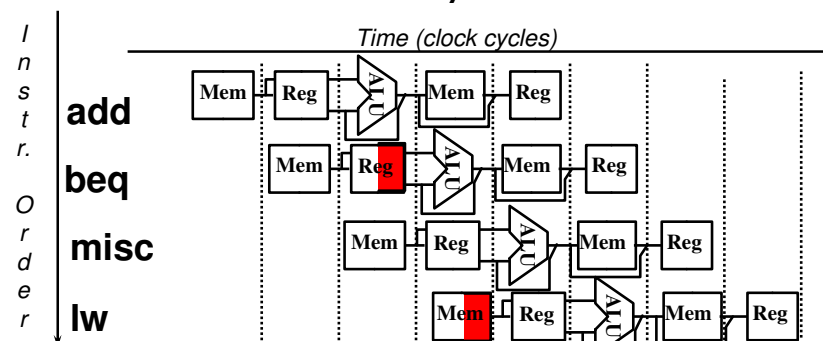


Fig. 6.38

Clock 4

Delayed Branch

- ◆ Predict-not-taken + branch decision at ID
=> the following instruction is always executed
=> branches take effect 1 cycle later



- 0 clock cycle penalty per branch instruction if can find instruction to put in slot (≈50% of time)

Outline

- ◆ An overview of pipelining
- ◆ A pipelined datapath
- ◆ Pipelined control
- ◆ Data hazards and forwarding
- ◆ Data hazards and stalls
- ◆ Branch hazards
- ◆ **Exceptions**
- ◆ Superscalar and dynamic pipelining

What about Exceptions?

- ◆ 5 instructions executing in 5 stage pipeline
 - How to stop the pipeline? restart?
 - Who caused the interrupt?
 - Who to serve first, if multiple interrupts at the same time?
- ◆ Need to know in which stage an exception can occur => help determine cause

Stage	Problem interrupts occurring
IF	Page fault; misaligned memory access; memory-protection violation
ID	Undefined or illegal opcode
EX	Arithmetic exception
MEM	Page fault; misaligned memory access; memory error; mem-protection violation;

Handling Exceptions

- ◆ Suppose overflow occur at add \$1, \$2, \$1
 - Disable writes of instructions till trap hits WB, e.g., flush following instructions using IF.Flush, ID.Flush, EX.Flush to cause multiplexers to zero control signals (overflow exception detected at EX => flush offending instruction)
 - Force trap instruction into IF, e.g., fetch from 4000 0040hex by adding 4000 0040hex to PC input MUX
 - Save address of offending instruction in EPC
- ◆ Multiple interrupts: use priority hardware to choose the earliest instruction to interrupt
- ◆ External interrupts: flexible in when to interrupt

Pipeline with Exception

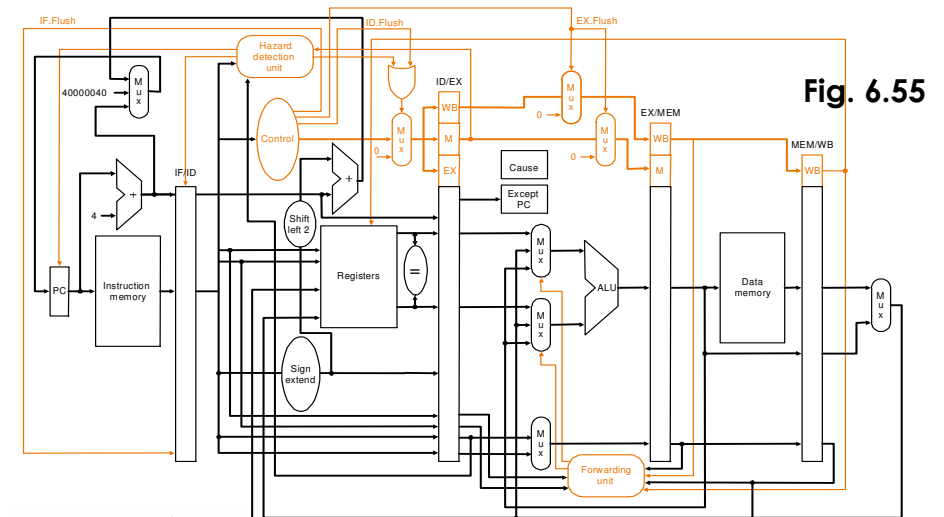



Fig. 6.55

Outline

- ◆ An overview of pipelining
- ◆ A pipelined datapath
- ◆ Pipelined control
- ◆ Data hazards and forwarding
- ◆ Data hazards and stalls
- ◆ Branch hazards
- ◆ Exceptions
- ◆ **Superscalar and dynamic pipelining**

Different Pipelined Designs

□ Pipelining




Limitation

Issue rate, FU stalls, FU depth

□ Super-pipeline


- Issue one instruction per (fast) cycle
- ALU takes multiple cycles



Clock skew, FU stalls, FU depth

□ Super-scalar


- Issue multiple scalar instructions per cycle



Hazard resolution

□ VLIW (EPIC)


- Each instruction specifies multiple scalar operations
- Compiler determines parallelism



Packing

□ Vector operations

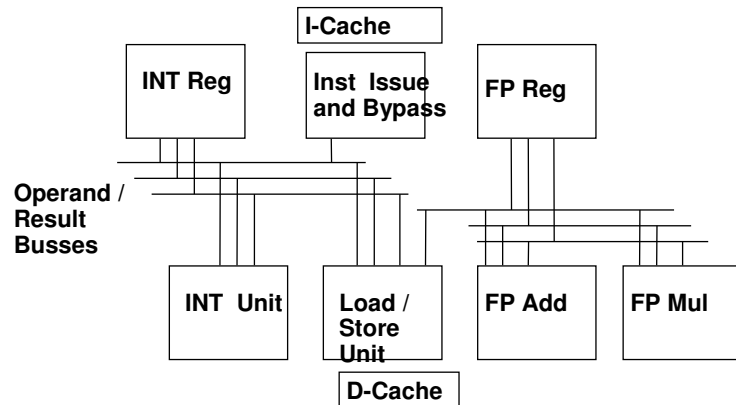
- Each instruction specifies series of identical operations



Applicability

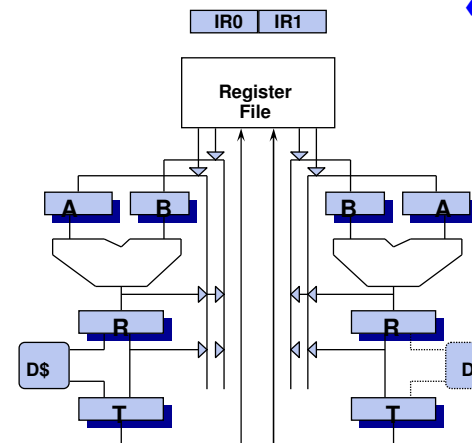
Alternative Simple Superscalar

- ◆ Independent INT and FP issue to separate pipelines



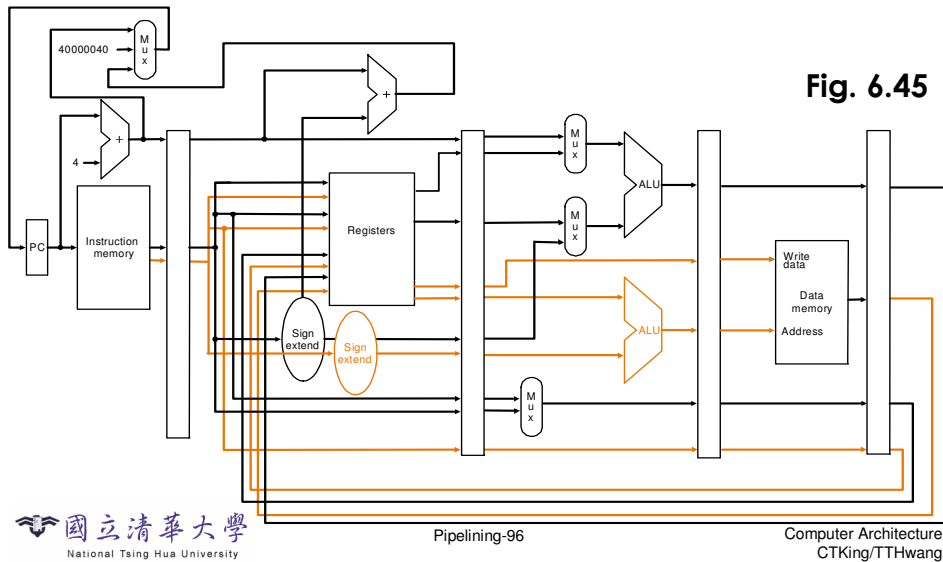
Multiple Pipes (Harder Superscalar)

- ◆ Issues:



- Register file ports
- Detecting data dependencies
- Bypassing
- RAW Hazard
- WAR Hazard
- Multiple load/store ops?
- Branches

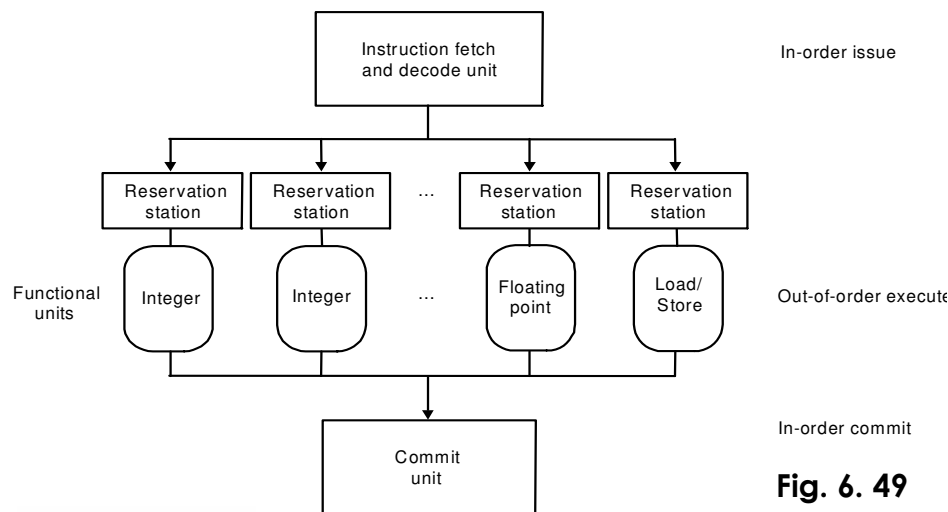
Simple MIPS Superscalar



Dynamic Scheduling

- ◆ The hardware performs the scheduling?
 - hardware tries to find instructions to execute
 - out of order execution is possible
 - speculative execution and dynamic branch prediction
- ◆ All modern processors are very complicated
 - DEC Alpha 21264: 9 stage pipeline, 6 instruction issue
 - PowerPC and Pentium: branch history table
 - Compiler technology important

Three Primary Units



Summary

- ◆ Pipelines pass control information down the pipe just as data moves down pipe
- ◆ Forwarding/stalls handled by local control
- ◆ Exceptions stop the pipeline
- ◆ MIPS instruction set architecture made pipeline visible (delayed branch, delayed load)
- ◆ More performance from deeper pipelines, parallelism