

MK HORGAN KAUFMANN

COMPUTER ORGANIZATION AND DESIGN
The Hardware/Software Interface

5th Edition

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

The diagram illustrates the CPU overview. It shows a PC (Program Counter) that provides an address to Instruction memory. Instruction memory outputs an instruction. The instruction is then used to access the Register file, which provides register numbers to the ALU. The ALU performs operations on data from the registers. The ALU result is used to calculate the next PC value (PC + 4) and to access Data memory. Data memory outputs data back to the Register file.

Chapter 4 — The Processor — 4

Multiplexers

The diagram shows the CPU overview with red circles highlighting the need for multiplexers. It indicates that simply joining wires together is not sufficient for routing data from multiple sources (like the ALU and registers) to a single destination (like the next PC value or the Register file).

- Can't just join wires together
- Use multiplexers

Chapter 4 — The Processor — 5

Control

The diagram shows the CPU overview with control signals. It includes a Control unit that generates signals for Branch, MemWrite, and MemRead. These signals are used to control the ALU operation, the Register file (for MemWrite), and the Data memory (for MemRead).

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
 - $Y = S ? I_1 : I_0$
- Arithmetic/Logic Unit
 - $Y = F(A, B)$

Chapter 4 — The Processor — 8

Multiplexors(1)

A two-input multiplexor has two data inputs (*A* and *B*) labeled 0 and 1, one selector input (*S*), and an output *C*.

Chapter 4 — The Processor — 9

Multiplexors(2)

- A multiplexor is arrayed 32 times to perform a selection between two 32-bit inputs.
- One data selection signal used for all 32 1-bit multiplexors.

Chapter 4 — The Processor — 10

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 — 11

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 — 12

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 — 13

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 — 14

Instruction Fetch

A portion of the datapath used for fetching instructions and incrementing the program counter. The fetched instruction is used by other parts of the datapath.

Chapter 4 — The Processor — 15

R-Format Instructions

- Read two register operands
- Perform arithmetic/logical operation
- Write register result

Chapter 4 — The Processor — 16

Load/Store Instructions


- Read register operands
- Calculate address using 16-bit offset
 - Use ALU, but sign-extend offset
- Load: Read memory and update register
- Store: Write register value to memory

Chapter 4 — The Processor — 17



Branch Instructions

- Read register operands
- Compare operands
 - Use ALU, subtract and check Zero output
- Calculate target address
 - Sign-extend displacement
 - Shift left 2 places (word displacement)
 - Add to PC + 4
 - Already calculated by instruction fetch

Chapter 4 — The Processor — 18



- ### Composing the Elements
- Simple data path does one instruction in one clock cycle
 - Each datapath element can only do one function at a time
 - Hence, we need separate instruction and data memories
 - Use multiplexers where alternate data sources are used for different instructions
- Chapter 4 — The Processor — 20



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 — 23

ALU Control (2)

- 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 — 24


The Main Control Unit

Control signals derived from instruction

R-type	0	rs	rt	rd	shamt	funct
	31:26	25:21	20:16	15:11	10:6	5:0
Load/Store	35 or 43	rs	rt	address		
	31:26	25:21	20:16	15:0		
Branch	4	rs	rt	address		
	31:26	25:21	20:16	15:0		

opcode
 always read
 read, except for load
 write for R-type and load
 sign-extend and add


Chapter 4 — The Processor — 25



The Main Control Unit (3)

Signal name	Effect when deasserted	Effect when asserted
RegDst	The register destination number for the Write register comes from the rt field (bits 20:16).	The register destination number for the Write register comes from the rd field (bits 15:11).
RegWrite	None.	The register on the Write register input is written with the value on the Write data input.
ALUSrc	The second ALU operand comes from the second register file output (Read data 2).	The second ALU operand is the sign-extended, lower 16 bits of the instruction.
PCSrc	The PC is replaced by the output of the adder that computes the branch target.	The PC is replaced by the output of the adder that computes the value of PC + 4.
MemRead	None.	Data memory contents designated by the address input are put on the Read data output.
MemWrite	None.	Data memory contents designated by the address input are replaced by the value on the Write data input.
MemtoReg	The value fed to the register Write data input comes from the ALU.	The value fed to the register Write data input comes from the data memory.

Chapter 4 — The Processor — 27




Datapath With Control (2)

Instruction	RegDst	ALUSrc	MemtoReg	RegWrite	MemRead	MemWrite	Branch	ALUOp1	ALUOp0
R-format	1	0	0	1	0	0	0	1	0
lw	0	1	1	1	1	0	0	0	0
sw	X	1	X	0	0	1	0	0	0
beq	X	0	X	0	0	0	1	0	1

- The setting of the control lines is completely determined by the opcode fields of the instruction.
- The first row of the table corresponds to the R-format instructions (add, sub, AND, OR, and slt). For all these instructions, the source register fields are rs and rt, and the destination register field is rd; this defines how the signals ALUSrc and RegDst are set. R-type instruction writes a register (Reg Write = 1), but neither reads nor writes data memory.
- When the Branch control signal is 0, the PC is unconditionally replaced with PC + 4; otherwise, the PC is replaced by the branch target if the Zero output of the ALU is also high.
- The ALUOp field for R-type instructions is set to 10 to indicate that the ALU control should be generated from the funct field.

Chapter 4 — The Processor — 29




R-Type Instruction (2)

For example, `add $t1,$t2,$t3`

- Four steps to execute the instruction in one clock cycle
1. The instruction is fetched, and the PC is incremented
 2. Registers \$t2 and \$t3 are read from the register file. Also, the main control unit computes the setting of the control lines during this step
 3. The ALU operates on the data read from the register file, using the function code (bits 5:0, funct field) to generate the ALU function
 4. The result from the ALU is written into the register file using bits 15:11 of the instruction to select the destination register (\$t1)



Load Instruction




Load Instruction (2)

For example, `lw $t1, offset($t2)`

- Five steps to execute the instruction in one clock cycle
1. Instruction is fetched from the instruction memory and PC is incremented
 2. Register (\$t2) value is read from the register file
 3. ALU computes the sum of the value read from register file and sign-extended offset
 4. Sum from ALU is used as the address for the data memory
 5. Data from the memory unit is written into register file. Register destination is given by bits 20:16 of the instruction (\$t1)



Branch-on-Equal Instruction




Branch-on-Equal Instruction (2)

For example, `beq $t1, $t2, offset`

- Four steps to execute the instruction in one clock cycle
1. An instruction is fetched from instruction memory and PC is incremented
 2. Register (\$t2) value is read from the register file
 3. ALU performs a subtract on the data values read from the register file. PC + 4 is added to sign-extended offset shifted left by two; result is branch target address
 4. Zero result from ALU is used to decide which adder result to store into PC




Implementing Jumps




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





- ### 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




- ### 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

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



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

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

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

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

Data Hazards

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

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

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 — 49

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 — 50

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

Chapter 4 — The Processor — 51

Stall on Branch

- Wait until branch outcome determined before fetching next instruction

Chapter 4 — The Processor — 52

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

Chapter 4 — The Processor — 53

MIPS with Predict Not Taken

Chapter 4 — The Processor — 54

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



Chapter 4 — The Processor — 55

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



Chapter 4 — The Processor — 56

Acknowledgement

The slides are adopted from Computer Organization and Design, 5th Edition by David A. Patterson and John L. Hennessy 2014, published by MK (Elsevier)



Chapter 4 — The Processor — 57