



## The Essentials of Computer Organization and Architecture

Linda Null and Julia Lobur  
Jones and Bartlett Publishers, 2003

# Chapter 5 Instructor's Manual

---

## Chapter Objectives

Chapter 5, A Closer Look at Instruction Set Architectures, provides a closer look at instruction set architectures, including instruction formats, instruction types, endianness, and addressing modes. Instruction-level pipelining is introduced as well. Real-world ISAs (including Intel, MIPS, and Java) are discussed to reinforce the concepts presented in the chapter.

This chapter should be covered after Chapter 4.

Lectures should focus on the following points:

- **Instruction formats.** Important issues here include instruction length, little versus big endian, register usage (and the use of stacks), and expanding opcodes.
- **Addressing.** Although addressing is an instruction design issue, there are many aspects to consider, the most important of which is addressing modes.
- **Instruction-level pipelining.** The fetch-decode-execute cycle can be overlapped, resulting in faster execution time. Resource conflicts, conditional branching, and data dependencies can slow this process down.
- **Real-world examples of ISAs.** Intel, MIPS, and the Java virtual machine are all presented to reinforce the concepts of the chapter.

## Required Lecture Time

The important concepts in Chapter 5 can typically be covered in 3 lecture hours. However, if a teacher wants the students to have a mastery of all topics in Chapter 5, 8 lecture hours are more reasonable. If lecture time is limited, we suggest that the focus be on instruction formats (including big versus little endian, use of registers, and instruction length) and addressing (with considerable attention paid to the various modes).

## Lecture Tips

It is important to motivate the little versus big endian debate. Many students comprehend the concept, but don't see the significance. The various software applications listed in this chapter will help the instructor make this point.

Students often have difficulty understanding stack machines. It is important to emphasize that not ALL instructions on stack machine have zero operands, but rather, the instructions that allow for operands are limited.

The concept of expanding opcodes is sometimes difficult for students as well. We suggest going over a small example in detail.

When covering the addressing modes, we suggest that instructors include many examples, as this is one concept that tends to be easier to understand through examples.

Students often confuse instruction-level pipelining with other types of pipelining. It is important to stress that there are many types of pipelining, but in this chapter only pipelining the fetch-decode-execute cycle is addressed.

Although the real-life examples can be left for students to read, these case studies provide the instructor with a means to tie the concepts from this chapter together, as well as a method for motivating study of the concepts from this chapter.

## Answers to Exercises

1. Assume you have a machine that uses 32-bit integers and you are storing the hex value 1234 at address 0.
  - a. Show how this is stored on a big endian machine.
  - b. Show how this is stored on a little endian machine.
  - c. If you wanted to increase the hex value to 123456, which byte assignment would be more efficient, big or little endian? Explain your answer.

Ans.

a and b.

Address →	00	01	10	11
Big Endian	00	00	12	34
Little Endian	34	12	00	00

- c. Little endian is more efficient because the additional information simply needs to be appended. With big endian, the "12" and "34" would need to shift to maintain the correct byte ordering.

- 
2. Show how the following values would be stored by machines with 32-bit words, using little endian and then big endian format. Assume each value starts at address  $10_{16}$ . Draw a diagram of memory for each, placing the appropriate values in the correct (and labeled) memory locations.

- a. 456789A<sub>16</sub>
- b. 0000058A<sub>16</sub>
- c. 14148888<sub>16</sub>

Ans.

a.

Address →	10 <sub>16</sub>	11 <sub>16</sub>	12 <sub>16</sub>	13 <sub>16</sub>
Big Endian	45	67	89	A1
Little Endian	A1	89	67	45

b.

Address →	10 <sub>16</sub>	11 <sub>16</sub>	12 <sub>16</sub>	13 <sub>16</sub>
Big Endian	00	00	05	8A
Little Endian	8A	05	00	00

c.

Address →	10 <sub>16</sub>	11 <sub>16</sub>	12 <sub>16</sub>	13 <sub>16</sub>
Big Endian	14	14	88	88
Little Endian	88	88	14	14

- ◆ 3. The first two bytes of a 2M x 16 main memory have the following hex values:

Byte 0 is FE

Byte 1 is 01

If these bytes hold a 16-bit two's complement integer, what is its actual decimal value if:

- ◆ a. memory is big endian?
- ◆ b. memory is little endian?

Ans.

a.  $FE01_{16} = 1111\ 1110\ 0000\ 0001_2 = -511_{10}$

b.  $01FE_{16} = 0000\ 0001\ 1111\ 1110_2 = 510_{10}$

4. What kinds of problems do you think endian-ness can cause if you wished to transfer data from a big endian machine to a little endian machine? Explain.

Ans.

If the machines receiving the data uses different endian-ness than the machine sending the data, the values can be misinterpreted. For example, the value from Problem 3, sent as the value -511 on a big endian machine, would be read as the value 510 on a little endian machine.

- ◆ 5. The Population Studies Institute monitors the population of the United States. In 2000, this institute wrote a program to create files of the numbers representing the various states, as well as the total population of the U.S. This program, which runs on a Motorola

processor, projects the population based on various rules, such as the average number of births and deaths per year. The Institute runs the program and then ships the output files to state agencies so the data values can be used as input into various applications. However, one Pennsylvania agency, running all Intel machines, encountered difficulties, as indicated by the following problem. When the 32-bit unsigned integer  $1D2F37E8_{16}$  (representing the overall U.S. population predication for 2003) is used as input, and the agency's program simply outputs this input value, the U.S. population forecast for 2003 is far too large. Can you help this Pennsylvania agency by explaining what might be going wrong?

Ans.

Hint: They are run on different processors.

Intel and Motorola use different endian-ness. If the program on the Intel machine does not adjust for this, then the integer is interpreted incorrectly.

6. There are reasons for machine designers to want all instructions to be the same length. Why is this not a good idea on a stack machine?


Ans.

The only instructions on a stack machine that need to address memory are push and pop. So an operand field is required, which implies the instruction field must be divided into an opcode and an operand. However, the other instructions need not access memory and can thus consist of only the opcode. To make them all the same length, these instructions would need to be "artificially lengthened".

- ◆7. A computer has 32-bit instructions and 12-bit addresses. Suppose there are 250 2-address instructions. How many 1-address instructions can be formulated? Explain your answer.

Ans.

There are 250 2-address instructions. There are only a total of 256 2-address instructions allowed if we have 32-bit instructions (two addresses take up 24 bits, leaving only 8 bits for the opcode). Looking at the 8 bit opcode, assume bit patterns 00000000 (0) through 11111001 (249) are used for the 250 two-address instructions. Then there are 6 bit patterns left for one address instructions. However, each one of these can use the remaining 12 bits gained from having only one operand, so we have  $6 * 2^{12}$ .



8. Convert the following expressions from infix to reverse Polish (postfix) notation.

- ◆ a.  $X * Y + W * Z + V * U$
- b.  $W * X + W * (U * V + Z)$
- c.  $(W * (X + Y * (U * V))) / (U * (X + Y))$

Ans.

- a.  $X Y * W Z * V U * + +$

- b. W X \* W U V \* Z + \* +
- c. W X Y U V \* \* + \* U X Y + \* /

9. Convert the following expressions from reverse Polish notation to infix notation.

- a. W X Y Z - + \*
- b. U V W X Y Z + \* + \* +
- c. X Y Z + V W - \* Z + +

Ans.

- a. W \* (X + Y - Z)
- b. U + (V \* (W + (X \* (Y + Z))))
- c. X + ((Y + Z) \* (V - W) + Z)

10. a. Write the following expression in postfix (Reverse Polish) notation. Remember the rules of precedence for arithmetic operators!

$$X = \frac{A - B + C * (D * E - F)}{G + H * K}$$

- b. Write a program to evaluate the above arithmetic statement using a stack organized computer with zero-address instructions (so only `pop` and `push` can access memory).

Ans.

- a. A B - C D E \* F - \* + G H K \* + /
- b. The program would be:

```

Push A
Push B
Subtract
Push C
Push D
Push E
Mult
Push F
Subtract
Mult
Add
Push G
Push H
Push K
Mult
Add
Div
Pop X

```

11. a. In a computer instruction format, the instruction length is 11 bits and the size of an address field is 4 bits. Is it possible to have:

- 5      2-address instructions
- 45     1-address instructions

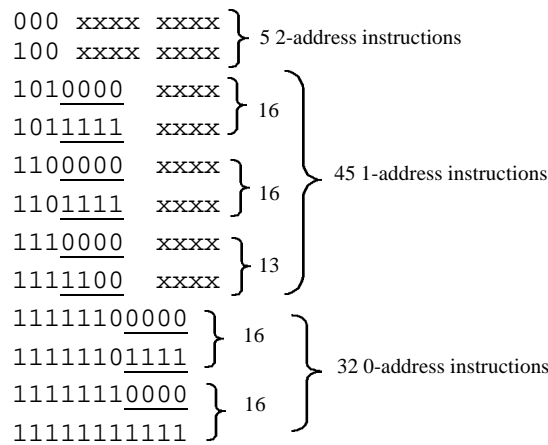
32 0-address instructions

using the format? Justify your answer.

- b. Assume that a computer architect has already designed 6 two-address and 24 zero-address instructions using the instruction format above. What is the maximum number of one-address instructions that can be added to the instruction set?

Ans.

- a. Yes. The 2-address instructions could be represented 000xxxxxxx through 100xxxxxxx (using 000 through 100 for opcodes). The 1-address instructions could use 1010000 xxxxx through 1011111 (16), 1100000 through 1101111 (16), and 1110000 through 1111100 (13 more, for a total of 45). The 0-address instructions could use 1111100000 through 1111101111 (16), and 1111110000 through 1111111111 (16). So we have:



- b. Assume the two-address instructions use bit patterns 000 xxxx xxxx through 101 xxxx xxxx. Assume also that the zero-address instructions are of the format 11111100000 through 11111101111 (8), 11111100000 through 11111101111 (8), and 11111110000 through 11111111111 (8) (These constitute the last 16 binary numbers possible with 11 bits). Then all instructions beginning with 110 (1100000 xxxx through 1101111 xxxx) could be one address instructions (16). In addition, 1110000 xxxx through 1111101 xxxx could be one address instructions, giving us 14 more, for a total of 30 1-address instructions.

12. What is the difference between using direct and indirect addressing? Give an example.

Ans.

Direct addressing provides the actual memory address of the operand in the instruction, whereas indirect addressing provides, as part of the instruction, a pointer to a memory location. For example, the instruction Load X interpreted using direct addressing would go to memory location X and load the value found there. Using indirect addressing, memory location X would be used as the effective address of what should actually be loaded. So if a value of 200 were found at location X, the value *located* at address 200 would be loaded.

- ◆ 13. Suppose we have the instruction Load 1000. Given memory and register R1 contain the values below:

Memory		
1000	1400	R1 <div style="border: 1px solid black; display: inline-block; padding: 2px 5px;">200</div>
...		
1100	400	
...		
1200	1000	
...		
1300	1100	
...		
1400	1300	

Assuming R1 is implied in the indexed addressing mode, determine the actual value loaded into the accumulator and fill in the table below:

Mode	Value Loaded into AC
Immediate	
Direct	
Indirect	
Indexed	

Ans.

Mode	Value
Immediate	1000
Direct	1400
Indirect	1300
Indexed	1000

14. Suppose we have the instruction Load 500. Given memory and register R1 contain the values below:

100	600	R1 <div style="border: 1px solid black; display: inline-block; padding: 2px 5px;">200</div>
...		
400	300	
...		
500	100	
...		
600	500	
...		
700	800	

Assuming R1 is implied in the indexed addressing mode, determine the actual value loaded into the accumulator and fill in the table below:

Mode	Value Loaded into AC
Immediate	
Direct	
Indirect	
Indexed	

Ans.

Mode	Value
Immediate	500
Direct	100
Indirect	600
Indexed	800

15. A nonpipelined system takes 200ns to process a task. The same task can be processed in a 5-segment pipeline with a clock cycle of 40ns. Determine the speedup ratio of the pipeline for 200 tasks. What is the maximum speedup that could be achieved with the pipeline unit over the nonpipelined unit?

Ans.

$$\text{SpeedUp} = (200\text{ns} \times 200) / ((5+200-1)(40\text{ns})) = 40000 / 8160 = 4.91$$

$$\text{Max SpeedUp} = 5$$

16. A nonpipelined system takes 100ns to process a task. The same task can be processed in a 5-segment pipeline with a clock cycle of 20ns. Determine the speedup ratio of the pipeline for 100 tasks. What is the theoretical speedup that could be achieved with the pipeline system over a nonpipelined system?

Ans.

$$\text{SpeedUp} = (100\text{ns} \times 100) / ((5+100-1)(20\text{ns})) = 10000 / 2080 = 4.8$$

$$\text{Max SpeedUp} = 5$$

17. Write code to implement the expression: **A = (B + C) \* (D + E)** on 3-, 2-, 1- and 0-address machines. In accordance with programming language practice, computing the expression should not change the values of its operands.

Ans.

Add R1, B, C
Add R2, D, E
Mult A, R1, R2

3-address machine

Load R1, B
Add R1, C
Load R2, E
Add R2, E
Mult R2, R1
Store A, R2

2-address machine

Load B
Add C
Store Temp
Load D
Add E
Mult Temp
Store A

1-address machine

Push B
Push C
Add
Push D
Push E
Add
Mult
Store A

0-address machine

◆ 18. A digital computer has a memory unit with 24 bits per word. The instruction set consists of 150 different operations. All instructions have an operation code part (opcode) and an



address part (allowing for only one address). Each instruction is stored in one word of memory.

- ◆ a. How many bits are needed for the opcode?
- ◆ b. How many bits are left for the address part of the instruction?
- ◆ c. What is the maximum allowable size for memory?
- ◆ d. What is the largest unsigned binary number that can be accommodated in one word of memory?

Ans.

- a. 150 instructions implies  $2^8$  ( $2^7$  will only give us 128 instructions), or 8 bits for the opcode.
  - b.  $24 - 8 = 16$ .
  - c.  $2^{16}$ , or 32M.
  - d. 24 1's, or  $2^{24} - 1$ .
- 

19. The memory unit of a computer has 256K words of 32 bits each. The computer has an instruction format with 4 fields: an opcode field; a mode field to specify 1 of 7 addressing modes; a register address field to specify one of 60 registers; and a memory address field. Assume an instruction is 32 bits long. Answer the following:

- a. How large must the mode field be?
- b. How large must the register field be?
- c. How large must the address field be?
- d. How large is the opcode field?

Ans.

- a. We need to identify 1 of 7 items, so there must be 3 bits ( $2^3 = 8$ )
  - b. 60 registers implies 6 bits ( $2^6 = 64$ )
  - c.  $256K = 2^8 2^{10} = 2^{18}$ , or 18 bits
  - d.  $32 - (3 + 6 + 18) = 5$  bits
- 

20. Suppose an instruction takes four cycles to execute in a nonpipelined CPU: one cycle to fetch the instruction, one cycle to decode the instruction, one cycle to perform the ALU operation, and one cycle to store the result. In a CPU with a 4-stage pipeline, that instruction still takes four cycles to execute, so how can we say the pipeline speeds up the execution of the program?

Ans.

For one instruction, there is no speedup. The speedup comes with the *parallel* execution of multiple instructions. While the first instruction is decoding, the second can be fetched; while the first instruction is performing the ALU instruction, the second can be decoding, and the third can be fetched, etc.

---

21. Pick an architecture (other than those covered in this chapter). Do research to find out how your architecture approaches the concepts introduced in this chapter, as was done for Intel, MIPS, and Java.

Ans.

No answer given.

---

## Sample Exam Questions

1. Suppose we have the instruction LDA 800. Given memory as follows:

Memory		
800	900	What would be loaded into the AC if the addressing mode for the operand is:  a. immediate _____  b. direct     _____  c. indirect    _____
...		
900	1000	
...		
1000	500	
...		
1100	600	
...		
1200	800	

Ans.

- a. 800
- b. 900
- c. 1000

2. Write the following expression in postfix (reverse Polish notation). Remember the rules of precedence for arithmetic operators.

$$X = A - B + C * (D * E - F) / (G + H * K)$$

Ans.

$$AB-CDE*F-*+GHK**+/$$

3. A nonpipelined system takes 300ns to process a task. The same task can be processed in a 5-segment pipeline with a clock cycle of 60ns. Determine the speedup ratio of the pipeline for 100 tasks. What is the maximum speedup that could be achieved with the pipeline unit over the nonpipelined unit?

Ans.

$$\text{SpeedUp} = (300\text{ns} \times 100) / ((5+100-1)(60\text{ns})) = 30000 / 6240 = 4.8$$

$$\text{Max SpeedUp} = 5$$

4. A digital computer has a memory unit with 32 bits per word. The instruction set consists of 128 different operations. All instructions have an operation code part (opcode) and an address part (allowing for only one address). Each instruction is stored in one word of memory.

- a. How many bits are needed for the opcode?
- b. How many bits are left for the address part of the instruction?
- c. What is the maximum allowable size for memory?

Ans.

- a. 128 instructions =  $2^7$ , or 7 bits for the opcode

- b.  $32 - 7 = 25$ .
- c.  $2^{25}$ , or 32G.

5. Show how the following values would be stored by machines with 32-bit words, using little endian and big endian format. Assume each value starts at address  $0_{16}$ . Draw a diagram of memory for each, placing the appropriate values in the correct (and labeled) memory locations.

- a.  $01234568_{16}$
- b.  $00001122_{16}$

Ans.

a.

Address $\longrightarrow$	$00_{16}$	$01_{16}$	$02_{16}$	$03_{16}$
Big Endian	01	23	45	68
Little Endian	68	45	23	01

b.

Address $\longrightarrow$	$00_{16}$	$01_{16}$	$02_{16}$	$03_{16}$
Big Endian	00	00	11	22
Little Endian	22	11	00	00

### TRUE or FALSE

- \_\_\_\_\_ 1. The term endian refers to the byte ordering, or the way a computer stores the bytes of a multiple-byte data element.
- \_\_\_\_\_ 2. Accumulator architectures use sets of general purpose registers to store operands.
- \_\_\_\_\_ 3. Most architectures today are accumulator based.
- \_\_\_\_\_ 4. Indexed or based addressing techniques add the value in the index or base register to the operand to produce the effective address.
- \_\_\_\_\_ 5. The core elements of an ISA include the memory model, registers, data types, instruction formats, addressing, and instruction types.
- \_\_\_\_\_ 6. Alternative addressing modes were prompted by the advances in memory technologies which resulted in larger memories.
- \_\_\_\_\_ 7. A fixed length instruction must have a fixed length opcode.
- \_\_\_\_\_ 8. The best architecture for evaluating postfix notation is the stack-based architecture.

Ans.

- 1. T                      5. T
- 2. F                      6. T
- 3. F                      7. F
- 4. T                      8. T