

MK ROSEAN KAUFMAN COMPUTER ORGANIZATION AND DESIGN The Hardware/Software Interface RISC-V Edition

Chapter 2

Instructions: Language of the Computer

These slides are based on the slides by the authors. The slides doesn't include all the material covered in the lecture. The slides will be explained, modified, and sometime corrected in the lecture.

Instruction Set

- The repertoire of instructions of a computer
- Different computers have different instruction sets
 - But with many aspects in common
- Early computers had very simple instruction sets
 - Simplified implementation
- Many modern computers also have simple instruction sets

Chapter 2 — Instructions: Language of the Computer — 2

RISC vs. CISC Instruction set

- RISC
 - Reduced Instruction Set Computer
 - Fixed instruction size
 - Simple instructions (load/store)
- CISC
 - Complex Instruction Set Computer
 - Variable instruction length
 - Much more powerful instructions
- **Which is better?**

Chapter 2 — Instructions: Language of the Computer — 3

The RISC-V Instruction Set

- Used as the example throughout the book
- Developed at UC Berkeley as open ISA
- Now managed by the RISC-V Foundation (riscv.org)
- Typical of many modern ISAs
 - See RISC-V Reference Data tear-out card
- Similar ISAs have a large share of embedded core market
 - Applications in consumer electronics, network/storage equipment, cameras, printers, ...

MK MK Chapter 2 — Instructions: Language of the Computer — 4

Arithmetic Operations

- Add and subtract, three operands
 - Two sources and one destination
add a, b, c // a gets b + c
- All arithmetic operations have this form
- Design Principle 1:** Simplicity favours regularity
 - Regularity makes implementation simpler
 - Simplicity enables higher performance at lower cost

MK MK Chapter 2 — Instructions: Language of the Computer — 5

Arithmetic Example

- C code:


```
f = (g + h) - (i + j);
```
- Compiled RISC-V code:


```
add t0, g, h // temp t0 = g + h
add t1, i, j // temp t1 = i + j
add f, t0, t1 // f = t0 - t1
```

MK MK Chapter 2 — Instructions: Language of the Computer — 6

Register Operands

- Arithmetic instructions use register operands
- RISC-V has a 32×64 -bit register file
 - Use for frequently accessed data
 - 64-bit data is called a "doubleword"
 - 32×64 -bit general purpose registers $x0$ to $x30$
 - 32-bit data is called a "word"
- Design Principle 2: Smaller is faster*
 - c.f. main memory: millions of locations

Chapter 2 — Instructions: Language of the Computer — 7

RISC-V Registers

- $x0$: the constant value 0
- $x1$: return address
- $x2$: stack pointer
- $x3$: global pointer
- $x4$: thread pointer
- $x5 - x7, x28 - x31$: temporaries
- $x8$: frame pointer
- $x9, x18 - x27$: saved registers
- $x10 - x11$: function arguments/results
- $x12 - x17$: function arguments

Chapter 2 — Instructions: Language of the Computer — 8

Register Operand Example

- C code:


```
f = (g + h) - (i + j);
```

 - f, \dots, j in $x19, x20, \dots, x23$
- Compiled RISC-V code:


```
add x5, x20, x21
add x6, x22, x23
sub x19, x5, x6
```

Chapter 2 — Instructions: Language of the Computer — 9

Memory Operands

- Main memory used for composite data
 - Arrays, structures, dynamic data
- To apply arithmetic operations
 - Load values from memory into registers
 - Store result from register to memory
- Memory is byte addressed
 - Each address identifies an 8-bit byte
- RISC-V is Little Endian
 - Least-significant byte at least address of a word
 - c.f. Big Endian: most-significant byte at least address
- RISC-V does not require words to be aligned in memory
 - Unlike some other ISAs

Chapter 2 — Instructions: Language of the Computer — 10

Memory Operand Example

- C code:


```
A[12] = h + A[8];
```

 - h in x21, base address of A in x22
- Compiled RISC-V code:
 - Index 8 requires offset of 64
 - 8 bytes per doubleword

```
ld    x9, 64(x22) // load double word
add   x9, x21, x9
sd    x9, 96(x22)
```

Chapter 2 — Instructions: Language of the Computer — 11

Registers vs. Memory

- Registers are faster to access than memory
- Operating on memory data requires loads and stores
 - More instructions to be executed
- Compiler must use registers for variables as much as possible
 - Only spill to memory for less frequently used variables
 - Register optimization is important!

Chapter 2 — Instructions: Language of the Computer — 12

Immediate Operands

- Constant data specified in an instruction
addi x22, x22, 4
- Make the common case fast**
 - Small constants are common
 - Immediate operand avoids a load instruction

MK MK Chapter 2 — Instructions: Language of the Computer — 13

Unsigned Binary Integers

- Given an n-bit number

$$x = x_{n-1}2^{n-1} + x_{n-2}2^{n-2} + \dots + x_12^1 + x_02^0$$
- Range: 0 to $2^n - 1$
- Example
 - 0000 0000 ... 0000 1011₂
 $= 0 + \dots + 1 \times 2^3 + 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0$
 $= 0 + \dots + 8 + 0 + 2 + 1 = 11_{10}$
- Using 64 bits: 0 to +18,446,774,073,709,551,615

MK MK Chapter 2 — Instructions: Language of the Computer — 14

2s-Complement Signed Integers

- Given an n-bit number

$$x = -x_{n-1}2^{n-1} + x_{n-2}2^{n-2} + \dots + x_12^1 + x_02^0$$
- Range: -2^{n-1} to $2^{n-1} - 1$
- Example
 - 1111 1111 ... 1111 1100₂
 $= -1 \times 2^{31} + 1 \times 2^{30} + \dots + 1 \times 2^2 + 0 \times 2^1 + 0 \times 2^0$
 $= -2,147,483,648 + 2,147,483,644 = -4_{10}$
- Using 64 bits: -9,223,372,036,854,775,808 to 9,223,372,036,854,775,807

MK MK Chapter 2 — Instructions: Language of the Computer — 15

2s-Complement Signed Integers

- Bit 63 is sign bit
 - 1 for negative numbers
 - 0 for non-negative numbers
- $-(-2^n - 1)$ can't be represented
- Non-negative numbers have the same unsigned and 2s-complement representation
- Some specific numbers
 - 0: 0000 0000 ... 0000
 - 1: 1111 1111 ... 1111
 - Most-negative: 1000 0000 ... 0000
 - Most-positive: 0111 1111 ... 1111

Chapter 2 — Instructions: Language of the Computer — 16

Signed Negation

- Complement and add 1
 - Complement means $1 \rightarrow 0, 0 \rightarrow 1$

$$\overline{x + \overline{x}} = 1111 \dots 1111_2 = -1$$

$$\overline{x + 1} = -x$$

- Example: negate +2
 - +2 = 0000 0000 ... 0010_{two}
 - 2 = 1111 1111 ... 1101_{two} + 1
 - = 1111 1111 ... 1110_{two}

Chapter 2 — Instructions: Language of the Computer — 17

Sign Extension

- Representing a number using more bits
 - Preserve the numeric value
- Replicate the sign bit to the left
 - c.f. unsigned values: extend with 0s
- Examples: 8-bit to 16-bit
 - +2: 0000 0010 => 0000 0000 0000 0010
 - 2: 1111 1110 => 1111 1111 1111 1110
- In RISC-V instruction set
 - l b: sign-extend loaded byte
 - l bu: zero-extend loaded byte

Chapter 2 — Instructions: Language of the Computer — 18

Representing Instructions

- Instructions are encoded in binary
 - Called machine code
- RISC-V instructions
 - Encoded as 32-bit instruction words
 - Small number of formats encoding operation code (opcode), register numbers, ...
 - Regularity!

Chapter 2 — Instructions: Language of the Computer — 19

Hexadecimal

- Base 16
 - Compact representation of bit strings
 - 4 bits per hex digit

0	0000	4	0100	8	1000	c	1100
1	0001	5	0101	9	1001	d	1101
2	0010	6	0110	a	1010	e	1110
3	0011	7	0111	b	1011	f	1111

- Example: eca8 6420
 - 1110 1100 1010 1000 0110 0100 0010 0000

Chapter 2 — Instructions: Language of the Computer — 20

RISC-V R-format Instructions

funct7	rs2	rs1	funct3	rd	opcode
7 bits	5 bits	5 bits	3 bits	5 bits	7 bits

- Instruction fields
 - opcode: operation code
 - rd: destination register number
 - funct3: 3-bit function code (additional opcode)
 - rs1: the first source register number
 - rs2: the second source register number
 - funct7: 7-bit function code (additional opcode)

Chapter 2 — Instructions: Language of the Computer — 21

R-format Example

funct7	rs2	rs1	funct3	rd	opcode
7 bits	5 bits	5 bits	3 bits	5 bits	7 bits

add x9, x20, x21

0	21	20	0	9	51
---	----	----	---	---	----

0000000	10101	10100	000	01001	0110011
---------	-------	-------	-----	-------	---------

0000 0001 0101 1010 0000 0100 1011 0011_{two} =
015A04B3₁₆

Chapter 2 — Instructions: Language of the Computer — 22

RISC-V I-format Instructions

immediate	rs1	funct3	rd	opcode
12 bits	5 bits	3 bits	5 bits	7 bits

- Immediate arithmetic and load instructions
 - rs1: source or base address register number
 - immediate: constant operand, or offset added to base address
 - 2s-complement, sign extended
- Design Principle 3: Good design demands good compromises
 - Different formats complicate decoding, but allow 32-bit instructions uniformly
 - Keep formats as similar as possible

Chapter 2 — Instructions: Language of the Computer — 23

RISC_V I-format Example

Imm[11:15]	rs2	rs1	funct3	Imm[4:1 11]	opcode
7 bits	5 bits	5 bits	3 bits	5 bits	7 bits

ld x9, 64(x3)

040	3	3	9	3
-----	---	---	---	---

00001000000	00011	011	01001	0000011
-------------	-------	-----	-------	---------

0000 0001 0101 1010 0000 0100 1011 0011_{two} =
015A04B3₁₆

Chapter 2 — Instructions: Language of the Computer — 24

RISC-V S-format Instructions

imm[11:5]	rs2	rs1	funct3	imm[4:0]	opcode
7 bits	5 bits	5 bits	3 bits	5 bits	7 bits

- Different immediate format for store instructions
 - rs1: base address register number
 - rs2: source operand register number
 - immediate: offset added to base address
 - Split so that rs1 and rs2 fields always in the same place

Chapter 2 — Instructions: Language of the Computer — 25

RISC-V S-format Instructions

imm[11:5]	rs2	rs1	funct3	imm[4:0]	opcode
7 bits	5 bits	5 bits	3 bits	5 bits	7 bits

sd x10, 80(x2)

2	10	2	3	9	51
---	----	---	---	---	----

0000010	01010	00010	011	10000	0100011
---------	-------	-------	-----	-------	---------

0000010	10000	=0x050	=80
---------	-------	--------	-----

Chapter 2 — Instructions: Language of the Computer — 26

RISC-V S-Format Instructions

imm[11:5]	rs2	rs1	funct3	imm[4:0]	opcode
7 bits	5 bits	5 bits	3 bits	5 bits	7 bits

sd x10, 80(x2)

0000010	01010				
7 bits	5 bits	5 bits	3 bits	5 bits	7 bits

	rs2	rs1	funct3		opcode
7 bits	5 bits	5 bits	3 bits	5 bits	7 bits

Chapter 2 — Instructions: Language of the Computer — 27

Stored Program Computers

The BIG Picture

- Instructions represented in binary, just like data
- Instructions and data stored in memory
- Programs can operate on programs
 - e.g., compilers, linkers, ...
- Binary compatibility allows compiled programs to work on different computers
 - Standardized ISAs

MK MK Chapter 2 — Instructions: Language of the Computer — 28

Logical Operations

- Instructions for bitwise manipulation

Operation	C	Java	RISC-V
Shift left	<<	<<	sl l i
Shift right	>>	>>>	srl i
Bit-by-bit AND	&	&	and, andi
Bit-by-bit OR			or, ori
Bit-by-bit XOR	^	^	xor, xori
Bit-by-bit NOT	~	~	

- Useful for extracting and inserting groups of bits in a word

MK MK Chapter 2 — Instructions: Language of the Computer — 29

Shift Operations

funct6	immed	rs1	funct3	rd	opcode
6 bits	6 bits	5 bits	3 bits	5 bits	7 bits

- immed: how many positions to shift
- Shift left logical
 - Shift left and fill with 0 bits
 - sl l i by i bits multiplies by 2^i
- Shift right logical
 - Shift right and fill with 0 bits
 - srl i by i bits divides by 2^i (unsigned only)

MK MK Chapter 2 — Instructions: Language of the Computer — 30

AND Operations

- Useful to mask bits in a word
 - Select some bits, clear others to 0

and x9, x10, x11

x10	00000000 00000000 00000000 00000000 00000000 00000000 00001101 11000000
x11	00000000 00000000 00000000 00000000 00000000 00000000 00111100 00000000
x9	00000000 00000000 00000000 00000000 00000000 00000000 00001100 00000000

MK MK Chapter 2 — Instructions: Language of the Computer — 31

OR Operations

- Useful to include bits in a word
 - Set some bits to 1, leave others unchanged

or x9, x10, x11

x10	00000000 00000000 00000000 00000000 00000000 00000000 00001101 11000000
x11	00000000 00000000 00000000 00000000 00000000 00000000 00111100 00000000
x9	00000000 00000000 00000000 00000000 00000000 00000000 00111101 11000000

MK MK Chapter 2 — Instructions: Language of the Computer — 32

XOR Operations

- Differencing operation
 - Set some bits to 1, leave others unchanged

xor x9, x10, x12 // NOT operation

x10	00000000 00000000 00000000 00000000 00000000 00000000 00001101 11000000
x12	11111111 11111111 11111111 11111111 11111111 11111111 11111111 11111111
x9	11111111 11111111 11111111 11111111 11111111 11111111 11110010 00111111

MK MK Chapter 2 — Instructions: Language of the Computer — 33
