


# Chapter 2

---


---

Instructions: Language of the  
Computer

# From logic gates to Computers




## COMPUTE SYSTEM




## RETI LOGICHE

## CIRCUITS




High-level language program (in C)

```
swap(int v[], int k)
{int temp;
  temp = v[k];
  v[k] = v[k+1];
  v[k+1] = temp;
}
```



Assembly language program (for RISC-V)

```
swap:
  slli x6, x11, 3
  add x6, x10, x6
  ld x5, 0(x6)
  ld x7, 8(x6)
  sd x7, 0(x6)
  sd x5, 8(x6)
  jalr x0, 0(x1)
```



Binary machine language program (for RISC-V)

```
00000000001101011001001100010011
00000000011001010000001100110011
0000000000000110011001010000011
0000000010000110011001110000011
0000000011100110011000000100011
00000000010100110011010000100011
0000000000000001000000001100111
```

# 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

# The RISC-V Instruction Set



- › Developed at UC Berkeley as open ISA starting in 2010
- › Now managed by the RISC-V Foundation ([riscv.org](http://riscv.org))
- › Typical of many modern ISAs
- › Similar ISAs have a large share of embedded core market
  - Applications in consumer electronics, network/storage equipment, cameras, printers, ...

# 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

# 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
```

```
sub f, t0, t1   // f = t0 - t1
```

# Register Operands



- › Arithmetic instructions use register operands
  
- › RISC-V has a  $32 \times 64$ -bit register file
  - Use for frequently accessed data
  - 64-bit data is called a “doubleword”
    - › 32 x 64-bit general purpose registers x0 to x31
  - 32-bit data is called a “word”
  
- › *Design Principle 2: Smaller is faster*
  - c.f. main memory: millions of locations

# 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



## RISC-V assembly language



Category	Instruction	Example	Meaning	Comments
Arithmetic	Add	add x5, x6, x7	$x5 = x6 + x7$	Three register operands; add
	Subtract	sub x5, x6, x7	$x5 = x6 - x7$	Three register operands; subtract
	Add immediate	addi x5, x6, 20	$x5 = x6 + 20$	Used to add constants
Data transfer	Load doubleword	ld x5, 40(x6)	$x5 = \text{Memory}[x6 + 40]$	Doubleword from memory to register
	Store doubleword	sd x5, 40(x6)	$\text{Memory}[x6 + 40] = x5$	Doubleword from register to memory
	Load word	lw x5, 40(x6)	$x5 = \text{Memory}[x6 + 40]$	Word from memory to register
	Load word, unsigned	lwu x5, 40(x6)	$x5 = \text{Memory}[x6 + 40]$	Unsigned word from memory to register
	Store word	sw x5, 40(x6)	$\text{Memory}[x6 + 40] = x5$	Word from register to memory
	Load halfword	lh x5, 40(x6)	$x5 = \text{Memory}[x6 + 40]$	Halfword from memory to register
	Load halfword, unsigned	lhu x5, 40(x6)	$x5 = \text{Memory}[x6 + 40]$	Unsigned halfword from memory to register
	Store halfword	sh x5, 40(x6)	$\text{Memory}[x6 + 40] = x5$	Halfword from register to memory
	Load byte	lb x5, 40(x6)	$x5 = \text{Memory}[x6 + 40]$	Byte from memory to register
	Load byte, unsigned	lbu x5, 40(x6)	$x5 = \text{Memory}[x6 + 40]$	Byte unsigned from memory to register
	Store byte	sb x5, 40(x6)	$\text{Memory}[x6 + 40] = x5$	Byte from register to memory
	Load reserved	lr.d x5, (x6)	$x5 = \text{Memory}[x6]$	Load; 1st half of atomic swap
	Store conditional	sc.d x7, x5, (x6)	$\text{Memory}[x6] = x5; x7 = 0/1$	Store; 2nd half of atomic swap
	Load upper immediate	lui x5, 0x12345	$x5 = 0x12345000$	Loads 20-bit constant shifted left 12 bits
Logical	And	and x5, x6, x7	$x5 = x6 \& x7$	Three reg. operands; bit-by-bit AND
	Inclusive or	or x5, x6, x8	$x5 = x6   x8$	Three reg. operands; bit-by-bit OR
	Exclusive or	xor x5, x6, x9	$x5 = x6 \wedge x9$	Three reg. operands; bit-by-bit XOR
	And immediate	andi x5, x6, 20	$x5 = x6 \& 20$	Bit-by-bit AND reg. with constant
	Inclusive or immediate	ori x5, x6, 20	$x5 = x6   20$	Bit-by-bit OR reg. with constant
	Exclusive or immediate	xori x5, x6, 20	$x5 = x6 \wedge 20$	Bit-by-bit XOR reg. with constant
Shift	Shift left logical	sll x5, x6, x7	$x5 = x6 \ll x7$	Shift left by register
	Shift right logical	srl x5, x6, x7	$x5 = x6 \gg x7$	Shift right by register
	Shift right arithmetic	sra x5, x6, x7	$x5 = x6 \gg x7$	Arithmetic shift right by register
	Shift left logical immediate	slli x5, x6, 3	$x5 = x6 \ll 3$	Shift left by immediate
	Shift right logical immediate	srl_i x5, x6, 3	$x5 = x6 \gg 3$	Shift right by immediate
	Shift right arithmetic immediate	srai x5, x6, 3	$x5 = x6 \gg 3$	Arithmetic shift right by immediate



Conditional branch	Branch if equal	beq x5, x6, 100	if (x5 == x6) go to PC+100	PC-relative branch if registers equal
	Branch if not equal	bne x5, x6, 100	if (x5 != x6) go to PC+100	PC-relative branch if registers not equal
	Branch if less than	blt x5, x6, 100	if (x5 < x6) go to PC+100	PC-relative branch if registers less
	Branch if greater or equal	bge x5, x6, 100	if (x5 >= x6) go to PC+100	PC-relative branch if registers greater or equal
	Branch if less, unsigned	bltu x5, x6, 100	if (x5 < x6) go to PC+100	PC-relative branch if registers less, unsigned
	Branch if greater or equal, unsigned	bgeu x5, x6, 100	if (x5 >= x6) go to PC+100	PC-relative branch if registers greater or equal, unsigned
Unconditional branch	Jump and link	jal x1, 100	x1 = PC+4; go to PC+100	PC-relative procedure call
	Jump and link register	jalr x1, 100(x5)	x1 = PC+4; go to x5+100	Procedure return; indirect call

# Register Operand Example



› C code:

```
f = (g + h) - (i + j);  
- f, ..., j in x19, x20, ..., x23
```

› Compiled RISC-V code:

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

# 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

# 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)
add     x9, x21, x9
sd     x9, 96(x22)
```

# 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!

# 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

# 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

$$\begin{aligned} & \blacksquare 0000\ 0000 \dots 0000\ 1011_2 \\ & \quad = 0 + \dots + 1 \times 2^3 + 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0 \\ & \quad = 0 + \dots + 8 + 0 + 2 + 1 = 11_{10} \end{aligned}$$

■ Using 64 bits: 0 to  
 $+18,446,774,073,709,551,615$



# 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

$$\begin{aligned} & \blacksquare 1111\ 1111\ \dots\ 1111\ 1100_2 \\ & \quad = -1 \times 2^{31} + 1 \times 2^{30} + \dots + 1 \times 2^2 + 0 \times 2^1 + 0 \times 2^0 \\ & \quad = -2,147,483,648 + 2,147,483,644 = -4_{10} \end{aligned}$$

■ Using 64 bits:  $-9,223,372,036,854,775,808$   
to  $9,223,372,036,854,775,807$

# 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

# Signed Negation



## › Complement and add 1

– Complement means  $1 \rightarrow 0, 0 \rightarrow 1$

$$x + \bar{x} = 1111\dots111_2 = -1$$

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

## ■ Example: negate +2

■  $+2 = 0000\ 0000 \dots 0010_{\text{two}}$

■  $-2 = 1111\ 1111 \dots 1101_{\text{two}} + 1$   
 $= 1111\ 1111 \dots 1110_{\text{two}}$

# 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
  - `lb`: sign-extend loaded byte
  - `lbu`: zero-extend loaded byte

# 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!

# 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

# RISC-V R-format Instructions



## › 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)

# 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<sub>two</sub> =  
015A04B3<sub>16</sub>



# RISC-V I-format Instructions



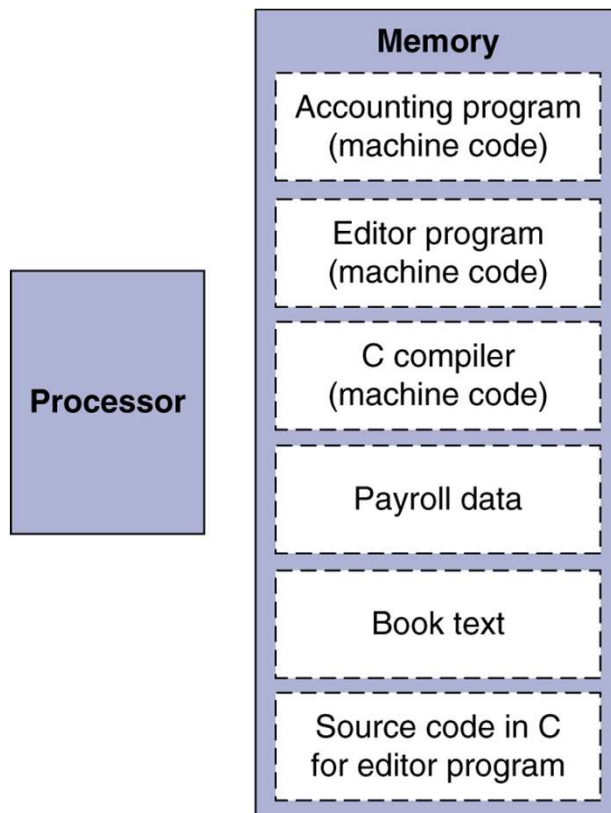
- › Immediate arithmetic and load instructions
  - rs1: source or base address register number
  - immediate: constant operand, or offset added to base address
    - › 2s-complement, sign extended
- › *Design Principle 3: Good design demands good compromises*
  - Different formats complicate decoding,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

# RISC-V S-format Instructions



- › 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

# Stored Program Computers



- › 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

# Logical Operations



› Instructions for bitwise manipulation

Operation	C	Java	RISC-V
Shift left	<code>&lt;&lt;</code>	<code>&lt;&lt;</code>	<code>slli</code>
Shift right	<code>&gt;&gt;</code>	<code>&gt;&gt;&gt;</code>	<code>srlr</code>
Bit-by-bit AND	<code>&amp;</code>	<code>&amp;</code>	<code>and, andi</code>
Bit-by-bit OR	<code> </code>	<code> </code>	<code>or, ori</code>
Bit-by-bit XOR	<code>^</code>	<code>^</code>	<code>xor, xori</code>
Bit-by-bit NOT	<code>~</code>	<code>~</code>	<code>xori FF..F</code>

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

# Shift Operations



- › I-format with just 6 bits for immediate
- › immed: how many positions to shift
- › Shift left logical
  - Shift left and fill with 0 bits
  - $sll\ i$  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)

# 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

# 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

# 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



# Conditional Operations

---



- › Branch to a labeled instruction if a condition is true
  - Otherwise, continue sequentially
  
- › `beq rs1, rs2, L1`
  - if (`rs1 == rs2`) branch to instruction labeled L1
  
- › `bne rs1, rs2, L1`
  - if (`rs1 != rs2`) branch to instruction labeled L1

# Compiling If Statements



> C code:


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

- f, g, ... in x19, x20, ...

> Compiled RISC-V code:

```
        bne x22, x23, Else  
        add x19, x20, x21  
        beq x0,x0,Exit // unconditional  
Else:   sub x19, x20, x21  
Exit:   ...
```

Assembler calculates addresses



# Compiling Loop Statements



› C code:

```
while (save[i] == k) i += 1;
```

– i in x22, k in x24, address of save in x25

› Compiled RISC-V code:


```
Loop: slli x10, x22, 3  
      add  x10, x10, x25  
      ld   x9, 0(x10)  
      bne  x9, x24, Exit  
      addi x22, x22, 1  
      beq  x0, x0, Loop
```

```
Exit: ...
```

# Basic Blocks



- > A basic block is a sequence of instructions with
  - No embedded branches (except at end)
  - No branch targets (except at beginning)



- A compiler identifies basic blocks for optimization
- An advanced processor can accelerate execution of basic blocks

# More Conditional Operations



- › `blt rs1, rs2, L1`
  - if ( $rs1 < rs2$ ) branch to instruction labeled L1
- › `bge rs1, rs2, L1`
  - if ( $rs1 \geq rs2$ ) branch to instruction labeled L1
- › Example
  - if ( $a > b$ ) `a += 1`;
  - a in x22, b in x23
  - `bge x23, x22, Exit`     // branch if  $b \geq a$
  - `addi x22, x22, 1`

Exit:

# Signed vs. Unsigned



- › Signed comparison: blt, bge
- › Unsigned comparison: bltu, bgeu
- › Example
  - x22 = 1111 1111 1111 1111 1111 1111 1111 1111
  - x23 = 0000 0000 0000 0000 0000 0000 0000 0001
  - x22 < x23 // signed
    - › -1 < +1
  - x22 > x23 // unsigned
    - › +4,294,967,295 > +1

# Procedure Calling

---



- › Steps required
  1. Place parameters in registers x10 to x17
  2. Transfer control to procedure
  3. Acquire storage for procedure
  4. Perform procedure's operations
  5. Place result in register for caller
  6. Return to place of call (address in x1)

# Procedure Call Instructions



- › Procedure call: jump and link

`jal x1, ProcedureLabel`

- Address of following instruction put in x1
- Jumps to target address (ProcedureLabel)

- › Procedure return: jump and link register

`jalr x0, 0(x1)`

- Like jal, but jumps to 0 + address in x1
- Use x0 as rd (x0 cannot be changed)
- Can also be used for computed jumps
  - › e.g., for case/switch statements

Unconditional branch	Jump and link	<code>jal x1, 100</code>	<code>x1 = PC+4; go to PC+100</code>	PC-relative procedure call
	Jump and link register	<code>jalr x1, 100(x5)</code>	<code>x1 = PC+4; go to x5+100</code>	Procedure return; indirect call



# Leaf Procedure Example



- › C code: **A procedure that doesn't call other procedures**

```
long long int leaf_example (  
    long long int g, long long int h,  
    long long int i, long long int j)  
{  
    long long int f;  
    f = (g + h) - (i + j);  
    return f;  
}
```

- Arguments g, ..., j in x10, ..., x13
- f in x20
- temporaries x5, x6
- Need to save x5, x6, x20 on stack (spill to mem)

## 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

# Leaf Procedure Example



RISC-V code:

```
addi sp, sp, -24
```

```
sd x5, 16(sp)
```

```
sd x6, 8(sp)
```

```
sd x20, 0(sp)
```

```
add x5, x10, x11
```

```
add x6, x12, x13
```

```
sub x20, x5, x6
```

```
addi x10, x20, 0
```

```
ld x20, 0(sp)
```

```
ld x6, 8(sp)
```

```
ld x5, 16(sp)
```

```
addi sp, sp, 24
```

```
jalr x0, 0(x1)
```

*Save x5, x6, x20 on stack*

*$x5 = g + h$*

*$x6 = i + j$*

*$f = x5 - x6$*

*copy f to return register*

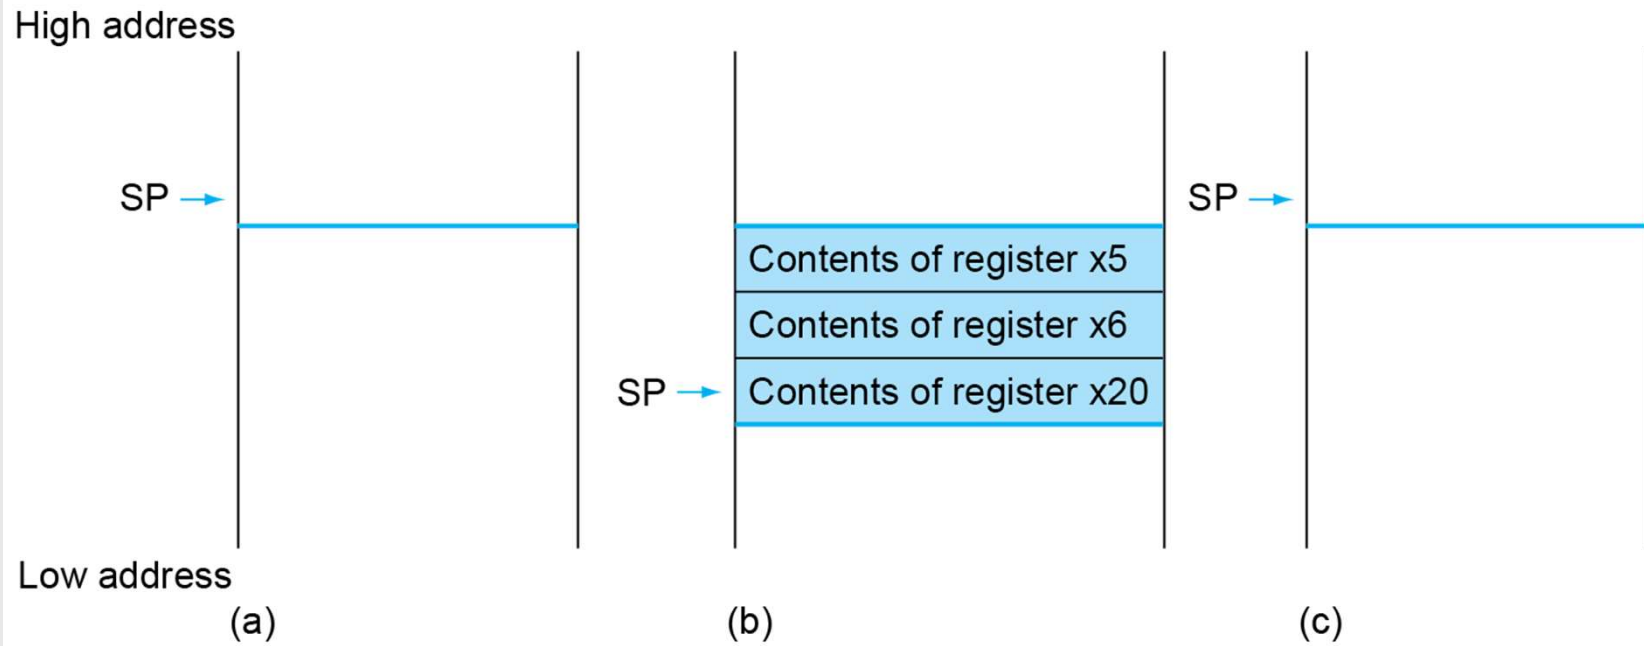
*Restore x5, x6, x20 from stack*

*Return to caller*

## 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

# Local Data on the Stack



# Register Usage – Calling convention



- › x5 – x7, x28 – x31: temporary registers
  - Not preserved by the **callee** (volatile across calls, must be saved by the **caller** if later used)
- › x8 – x9, x18 – x27: saved registers
  - Preserved across calls. If used, the **callee** saves and restores them
- › In previous example, the stores/loads on x5 and x6 can be dropped

## 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

# Register Usage – Calling convention



- › x5 – x7, x28 – x31: temporaries
- Not preserved across calls, must be saved

- › x8 – x17: saved registers
- Preserve the caller's state

- › In previous versions, x10 stores/loads can be dropped

## Chapter 18

### Calling Convention

This chapter describes the C compiler standards for RV32 and RV64 programs and two calling conventions: the convention for the base ISA plus standard general extensions (RV32G/RV64G), and the soft-float convention for implementations lacking floating-point units (e.g., RV32I/RV64I).

*Implementations with ISA extensions might require extended calling conventions.*

#### 18.1 C Datatypes and Alignment

Table 18.1 summarizes the datatypes natively supported by RISC-V C programs. In both RV32 and RV64 C compilers, the C type `int` is 32 bits wide. `long`s and pointers, on the other hand, are both as wide as an integer register, so in RV32, both are 32 bits wide, while in RV64, both are 64 bits wide. Equivalently, RV32 employs an ILP32 integer model, while RV64 is LP64. In both RV32 and RV64, the C type `long` is a 64-bit integer, `float` is a 32-bit IEEE 754-2008 floating-point number, `double` is a 64-bit IEEE 754-2008 floating-point number, and `long double` is a 128-bit IEEE floating-point number.

The C types `char` and `unsigned char` are 8-bit unsigned integers and are zero-extended when stored in a RISC-V integer register. `unsigned short` is a 16-bit unsigned integer and is zero-extended when stored in a RISC-V integer register. `signed char` is an 8-bit signed integer and is sign-extended when stored in a RISC-V integer register, i.e. bits (XLEN-1)..7 are all equal. `short` is a 16-bit signed integer and is sign-extended when stored in a register.

In RV64, 32-bit types, such as `int`, are stored in integer registers as proper sign extensions of their 32-bit values; that is, bits 63..31 are all equal. This restriction holds even for unsigned 32-bit types.

## Registers

- › x5 – x7, x28 – x31: temporaries
- › x8: frame pointer
- › x9, x18 – x27: saved registers
- › x10 – x11: function arguments/results
- › x12 – x17: function arguments

# Non-Leaf Procedures



- › Procedures that call other procedures
- › For nested call, caller needs to save on the stack:
  - Its return address
  - Any arguments and temporaries needed after the call
  - Not the saved registers
    - › *Those are handled by the **callee** (if used)*
- › Restore from the stack after the call

# Non-Leaf Procedure Example



> C code:

```
long long int fact (long long int n)
{
    if (n < 1) return 1;
    else return n * fact(n - 1);
}
```

- Argument n in x10
- Result in x10

# Leaf Procedure Example



RISC-V code:

fact:

```
addi sp,sp,-16
```

*Save return address and n on stack*

```
sd x1,8(sp)
```

```
sd x10,0(sp)
```

```
addi x5,x10,-1
```

*x5 = n - 1*

```
bge x5,x0,L1
```

*If n>0 (equiv. to n >= 1), go to L1*

```
addi x10,x0,1
```

*Else, set return value to 1*

```
addi sp,sp,16
```

*Pop stack, don't bother restoring values*

```
jalr x0,0(x1)
```

*Return*

```
L1: addi x10,x10,-1
```

*n = n - 1*

```
jal x1,fact
```

*Call fact(n-1)*

```
addi x6,x10,0
```

*Move result of fact(n - 1) to x6*

```
ld x10,0(sp)
```

*Restore caller's n*

```
ld x1,8(sp)
```

*Restore caller's return address*

```
addi sp,sp,16
```

*Pop stack*

```
mul x10,x10,x6
```

*Return n \* fact(n-1)*

```
jalr x0,0(x1)
```

*Return*

```
lli fact (lli n)
{
    if (n < 1)
        return 1;
    else
        return
            n*fact(n - 1);
}
```



# Across procedure call




Preserved	Not preserved
Saved registers: x8-x9, x18-x27	Temporary registers: x5-x7, x28-x31
Stack pointer register: x2(sp)	Argument/result registers: x10-x17
Frame pointer: x8(fp)	
Return address: x1(ra)	
Stack above the stack pointer	Stack below the stack pointer

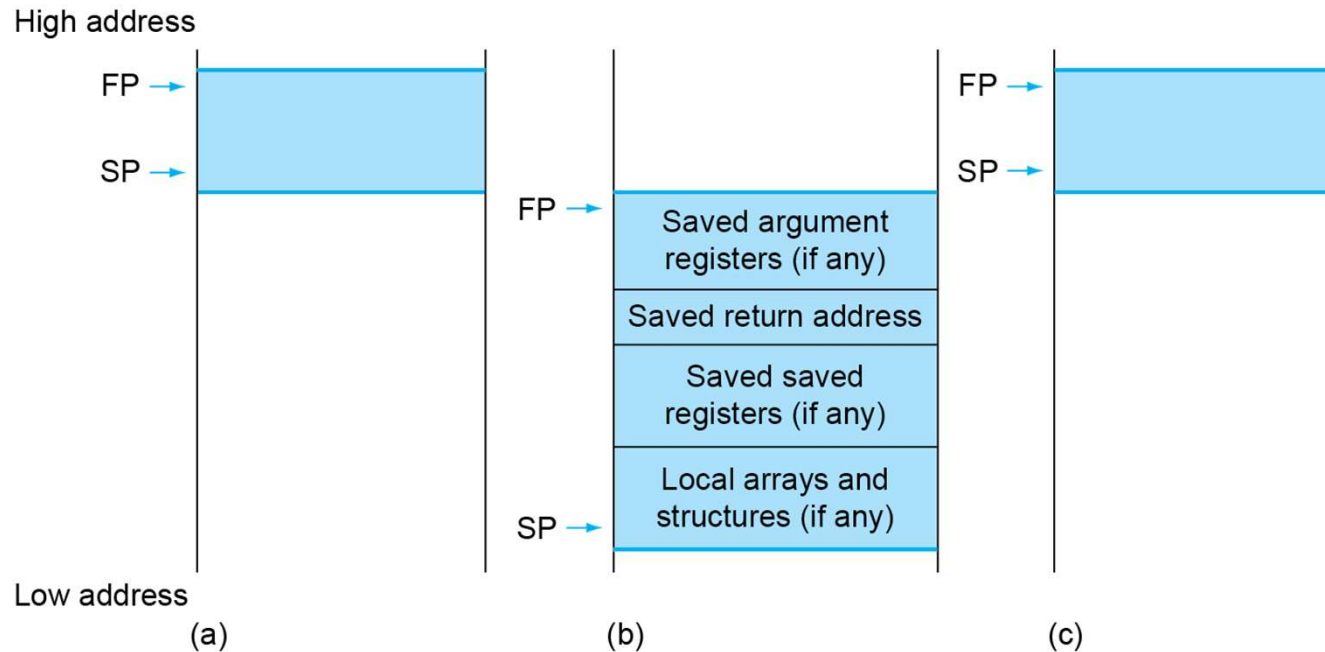
# Memory Layout



- › Text: program code
- › Static data: global variables
  - e.g., static variables in C, constant arrays and strings
  - x3 (global pointer) initialized to address allowing  $\pm$ offsets into this segment
- › Dynamic data: heap
  - E.g., malloc in C, new in Java
- › Stack: automatic storage



# Local Data on the Stack



- › Local data allocated by callee
  - e.g., C automatic variables
- › Procedure frame (activation record)
  - Used by some compilers to manage stack storage

# Character Data



- › Byte-encoded character sets
  - ASCII: 128 characters
    - › 95 graphic, 33 control
  - Latin-1: 256 characters
    - › ASCII, +96 more graphic characters
- › Unicode: 32-bit character set
  - Used in Java, C++ wide characters, ...
  - Most of the world's alphabets, plus symbols
  - UTF-8, UTF-16: variable-length encodings

# Byte/Halfword/Word Operations



- › RISC-V byte/halfword/word load/store
  - Load byte/halfword/word: Sign extend to 64 bits in rd
    - › `lb rd, offset(rs1)`
    - › `lh rd, offset(rs1)`
    - › `lw rd, offset(rs1)`
  - Load byte/halfword/word unsigned: Zero extend to 64 bits in rd
    - › `lbu rd, offset(rs1)`
    - › `lhu rd, offset(rs1)`
    - › `lwu rd, offset(rs1)`
  - Store byte/halfword/word: Store rightmost 8/16/32 bits
    - › `sb rs2, offset(rs1)`
    - › `sh rs2, offset(rs1)`
    - › `sw rs2, offset(rs1)`

# String Copy Example



› C code:

– Null-terminated string

```
void strcpy (char x[], char y[])
{ size_t i;
  i = 0;
  while ((x[i]=y[i])!='\0')
    i += 1;
}
```

# String Copy Example



> RISC-V code: `x10 = x; x11 = y; x19 = i;`

```
strcpy:
    addi sp,sp,-8           // adjust stack for 1 dw
    sd   x19,0(sp)         // push x19
    add  x19,x0,x0          // i=0
L1:   add  x5,x19,x11       // x5 = addr of y[i]
    lbu  x6,0(x5)           // x6 = y[i]
    add  x7,x19,x10        // x7 = addr of x[i]
    sb   x6,0(x7)           // x[i] = y[i]
    beq  x6,x0,L2          // if y[i] == 0 then exit
    addi x19,x19,1         // i = i + 1
    jal  x0,L1             // next iteration of loop
L2:   ld   x19,0(sp)        // restore saved x19
    addi sp,sp,8           // pop 1 dw from stack
    jalr x0,0(x1)          // and return
```

Could have used a temporary register (e.g., x28-x31)  
instead of x19 to avoid needing to save it

# 32-bit Constants



- › Most constants are small
  - 12-bit immediate is sufficient
- › For the occasional 32-bit constant

## `lui rd, constant`

- Copies 20-bit constant to bits [31:12] of rd
- Extends bit 31 to bits [63:32]
- Clears bits [11:0] of rd to 0

`lui x19, 976 // 0x003D0`

0000 0000 0000 0000	0000 0000 0000 0000	0000 0000 0011 1101 0000	0000 0000 0000
---------------------	---------------------	--------------------------	----------------

`addi x19,x19,128 // 0x500`

0000 0000 0000 0000	0000 0000 0000 0000	0000 0000 0011 1101 0000	0101 0000 0000
---------------------	---------------------	--------------------------	----------------



# Branch Addressing



- › Branch instructions specify
  - Opcode, two registers, target address
- › Most branch targets are near branch
  - Forward or backward
- › SB format:




- PC-relative addressing
  - Target address = PC + immediate × 2  
(Addressing instructions down to halfword)

# Jump Addressing



- › Jump and link (jal) target uses 20-bit immediate for larger range
- › UJ format:




- For long jumps, eg, to 32-bit absolute address
  - lui: load address[31:12] to temp register
  - jalr: add address[11:0] and jump to target


# RISC-V Addressing Summary





## 1. Immediate addressing





## 2. Register addressing



## 3. Base addressing



## 4. PC-relative addressing



# RISC-V Encoding Summary



Name (Field Size)	Field						Comments
	7 bits	5 bits	5 bits	3 bits	5 bits	7 bits	
R-type	funct7	rs2	rs1	funct3	rd	opcode	Arithmetic instruction format
I-type	immediate[11:0]		rs1	funct3	rd	opcode	Loads & immediate arithmetic
S-type	immed[11:5]	rs2	rs1	funct3	immed[4:0]	opcode	Stores
SB-type	immed[12,10:5]	rs2	rs1	funct3	immed[4:1,11]	opcode	Conditional branch format
UJ-type	immediate[20,10:1,11,19:12]				rd	opcode	Unconditional jump format
U-type	immediate[31:12]				rd	opcode	Upper immediate format

# Synchronization



- › Two processors sharing an area of memory
  - P1 writes, then P2 reads
  - Data race if P1 and P2 don't synchronize
    - › Result depends of order of accesses
- › Hardware support required
  - Atomic read/write memory operation
  - No other access to the location allowed between the read and write
- › Could be a single instruction
  - E.g., atomic swap of register  $\leftrightarrow$  memory
  - Or an atomic pair of instructions

# Synchronization in RISC-V



- › Load reserved: `l r . d rd, (rs1)`
  - Load from address in rs1 to rd
  - Place reservation on memory address
  
- › Store conditional: `sc . d rd, rs2, (rs1)`
  - Store from rs2 to address in rs1
  - Succeeds if location not changed since the `l r . d`
    - › Returns 0 in rd
  - Fails if location is changed
    - › Returns non-zero value in rd

# Synchronization in RISC-V



## › Example 1: Lock / Unlock

### Lock

```
        addi x12,x0,1      // copy locked value
again:  lr.d  x10,(x20)    // read lock
        bne  x10,x0,again // check if it is 0 yet
        sc.d x11,x12,(x20) // attempt to store
        bne  x11,x0,again // branch if fails
```

0 = LOCK LIBERO  
1 = LOCK OCCUPATO

### Unlock:

```
sd  x0,0(x20) // free lock
```

### Thread 1

```
local_s1= 0
for i = 0, n/2-1
    local_s1 = local_s1 + sqr(A[i])
    lock(lk);
    s = s + local_s1
    unlock(lk); (*)
```

(\*) stralcio di codice estratto da *a) Approfondimento\_Sincronizzazione.pdf*

# Synchronization in RISC-V



- › Example 2: **Atomic swap** (to test/set lock variable)

```
again:  lr.d x10, (x20)
        sc.d x11, x23, (x20) // x11 = status
        bne x11, x0, again // branch if failed
        addi x23, x10, 0 // x23 = loaded value
```

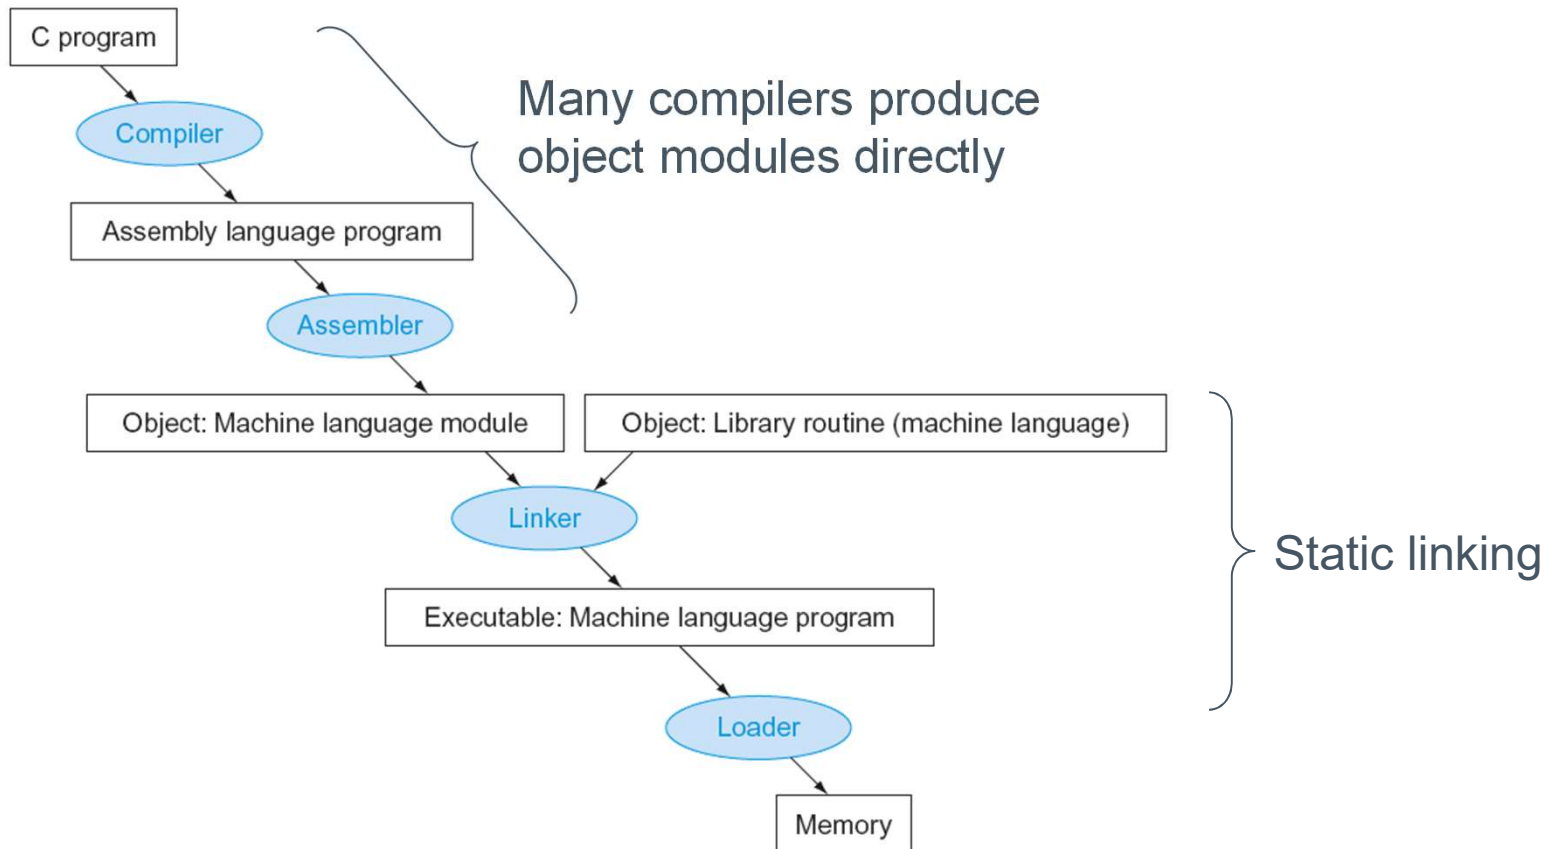
**SCAMBIA ATOMICAMENTE IL VALORE DELLA LOCAZIONE DI MEMORIA IL CUI INDIRIZZO È SPECIFICATO IN x20 COL VALORE SPECIFICATO IN x23. EQUIVALE A:**

```
long long int swap (long long int * p, long long int val)
{
    LOCK
        long long int temp = *p; // ld x10, (x20)
        *p = val; // sd x23, (x20)
    UNLOCK

    return temp; // addi x23, x10, 0
}
```



# Translation and Startup




# Producing an Object Module



- › Assembler (or compiler) translates program into machine instructions
- › Provides information for building a complete program from the pieces
  - **Header**: described contents of object module
  - **Text** segment: translated instructions
  - **Static** data segment: data allocated for the life of the program
  - **Relocation info**: for contents that depend on absolute location of loaded program
  - **Symbol table**: global definitions and external refs
  - **Debug info**: for associating with source code

SP → 0000 003f ffff fff0<sub>hex</sub>

PC → 0000 0000 0040 0000<sub>hex</sub>




# Linking Object Modules



- › Produces an executable image
  1. Merges segments
  2. Resolve labels (determine their addresses)
  3. Patch location-dependent and external refs
- › Could leave location dependencies for fixing by a relocating loader
  - But with virtual memory, no need to do this
  - Program can be loaded into absolute location in virtual memory space

SP → 0000 003f ffff fff0<sub>hex</sub>

PC → 0000 0000 0040 0000<sub>hex</sub>



0

# Loading a Program




## › Load from image file on disk into memory

1. Read header to determine segment sizes
2. Create virtual address space
3. Copy text and initialized data into memory
  - › Or set page table entries so they can be faulted in
4. Set up arguments on stack
5. Initialize registers (including sp, fp, gp)
6. Jump to startup routine

- › Copies arguments to x10, ... and calls main
- › When main returns, do **exit** syscall

SP → 0000 003f ffff fff0<sub>hex</sub>

PC → 0000 0000 0040 0000<sub>hex</sub>



# Dynamic Linking



- › Only link/load library procedure when it is called
  - Requires procedure code to be **relocatable**
  - Avoids image bloat caused by static linking of all (transitively) referenced libraries
  - Automatically picks up new library versions

## ESTENSIONI STANDARD DELLE LIBRERIE DINAMICHE

*.so	nei sistemi LINUX
*.dll	nei sistemi WINDOWS

# Lazy Linkage

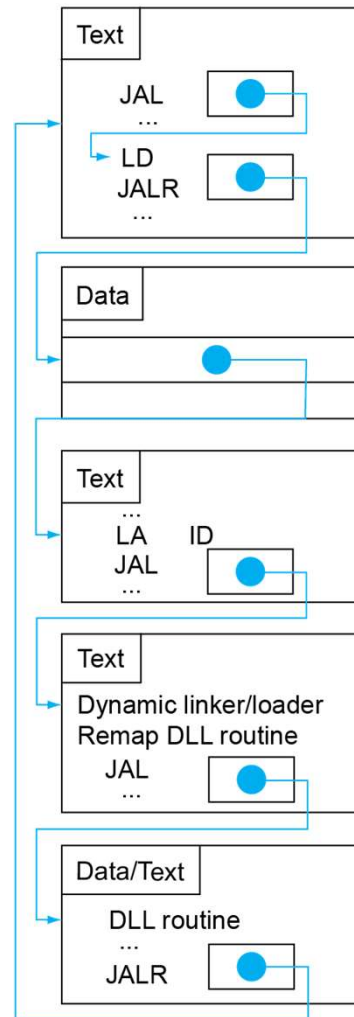


Indirection table


Stub: Loads routine ID,  
Jump to linker/loader

Linker/loader code

Dynamically  
mapped code

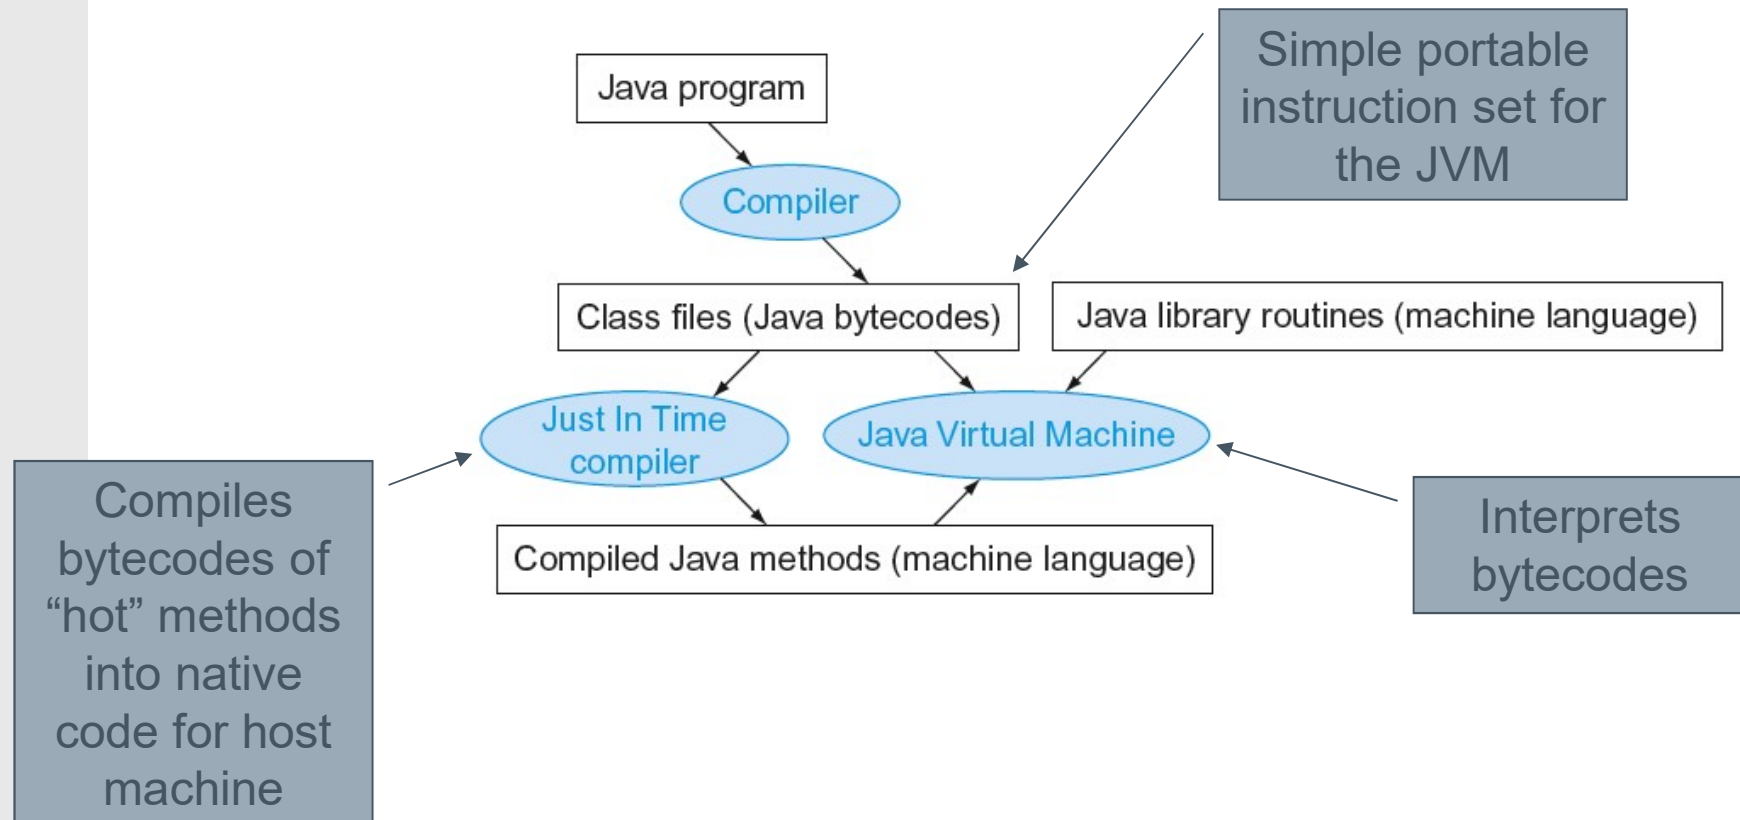


(a) First call to DLL routine



(b) Subsequent calls to DLL routine

# Starting Java Applications



# C Sort Example



- › Illustrates use of assembly instructions for a C sort function
- › Swap procedure (leaf)

```
void swap (long long int v[],
           long long int k)
{
    long long int temp;
    temp = v[k];
    v[k] = v[k+1];
    v[k+1] = temp;
}
```

– v in x10, k in x11, temp in x5



# The Procedure Swap



swap:

```
slli x6,x11,3    // reg x6 = k * 8
add  x6,x10,x6   // reg x6 = v + (k * 8)
ld   x5,0(x6)   // reg x5 (temp) = v[k]
ld   x7,8(x6)   // reg x7 = v[k + 1]
sd   x7,0(x6)   // v[k] = reg x7
sd   x5,8(x6)   // v[k+1] = reg x5 (temp)
jalr x0,0(x1)   // return to calling routine
```

# The Sort Procedure in C



## › Non-leaf (calls swap)

```
void sort (long long int v[], size_t n)
{
    size_t i, j;
    for (i = 0; i < n; i += 1) {
        for (j = i - 1;
             j >= 0 && v[j] > v[j + 1];
             j -= 1)
        {
            swap(v, j);
        }
    }
}
```

– v in x10, n in x11, i in x19, j in x20

# The Outer Loop



› Skeleton of outer loop:

```
- for (i = 0; i < n; i += 1) {
```

```
    li    x19, 0           // i = 0
```

```
for1tst:
```

```
    bge  x19, x11, exit1  // goto exit1 if x19 ≥ x11 (i ≥ n)
```

```
    (body of outer for-loop)
```

```
    addi x19, x19, 1      // i += 1
```

```
    j    for1tst         // branch to test outer loop
```

```
exit1:
```

# The Inner Loop



› Skeleton of inner loop:

– for (j = i - 1; j >= 0 && v[j] > v[j + 1]; j - = 1) {

```
        addi x20,x19,-1    // j = i -1
for2tst:
    blt  x20,x0,exit2    // go to exit2 if x20 < 0 (j < 0)
    slli x5,x20,3        // reg x5 = j * 8
    add  x5,x10,x5        // reg x5 = v + (j * 8)
    ld   x6,0(x5)        // reg x6 = v[j]
    ld   x7,8(x5)        // reg x7 = v[j + 1]
    ble  x6,x7,exit2    // go to exit2 if x6 ≤ x7
    mv   x21, x10        // copy parameter x10 into x21
    mv   x22, x11        // copy parameter x11 into x22
    mv   x10, x21        // first swap parameter is v
    mv   x11, x20        // second swap parameter is j
    jal  x1,swap         // call swap
    addi x20,x20,-1      // j -= 1
    j    for2tst         // branch to test of inner loop
exit2:
```

# Preserving Registers



## > Preserve saved registers:

```
addi sp,sp,-40 // make room on stack for 5 regs
sd   x1,32(sp) // save x1 on stack
sd   x22,24(sp) // save x22 on stack
sd   x21,16(sp) // save x21 on stack
sd   x20,8(sp)  // save x20 on stack
sd   x19,0(sp)  // save x19 on stack
```

## > Restore saved registers:

exit1:

```
ld   x19,0(sp) // restore x19 from stack
ld   x20,8(sp) // restore x20 from stack
ld   x21,16(sp) // restore x21 from stack
ld   x22,24(sp) // restore x22 from stack
ld   x1,32(sp) // restore x1 from stack
addi sp,sp, 40 // restore stack pointer
jalr x0,0(x1)
```

### Register Usage – Calling convention

- > x5 – x7, x28 – x31: temporary registers
  - Not preserved by the **callee** (volatile across calls, must be saved by the **caller** if later used)
- > x8 – x9, x18 – x27: saved registers
  - Preserved across calls. If used, the **callee** saves and restores them
- > In previous example, the stores/loads on x5 and x6 can be dropped

#### 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
- > x12 – x17: function arguments

**calling convention**

# The full sort procedure



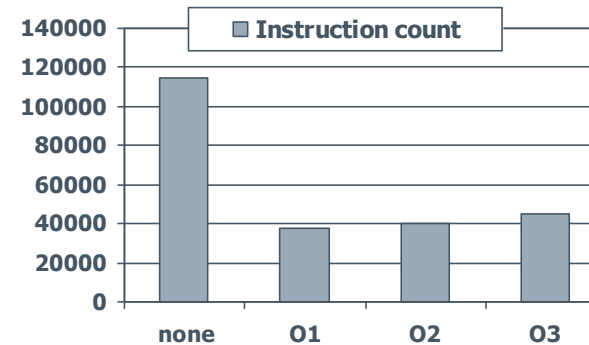



Saving registers		
	sort:	<pre> addi sp, sp, -40    # make room on stack for 5 registers sd x1, 32(sp)      # save return address on stack sd x22, 24(sp)     # save x22 on stack sd x21, 16(sp)     # save x21 on stack sd x20, 8(sp)      # save x20 on stack sd x19, 0(sp)      # save x19 on stack </pre>
Procedure body		
Move parameters		<pre> mv x21, x10        # copy parameter x10 into x21 mv x22, x11        # copy parameter x11 into x22 </pre>
Outer loop		<pre> li x19, 0          # i = 0 for1tst:bge x19, x22, exit1 # go to exit1 if i &gt;= n </pre>
Inner loop		<pre> addi x20, x19, -1 # j = i - 1 for2tst:blt x20, x0, exit2 # go to exit2 if j &lt; 0 slli x5, x20, 3   # x5 = j * 8 add x5, x21, x5   # x5 = v + (j * 8) ld x6, 0(x5)      # x6 = v[j] ld x7, 8(x5)      # x7 = v[j + 1] ble x6, x7, exit2 # go to exit2 if x6 &lt; x7 </pre>
Pass parameters and call		<pre> mv x10, x21        # first swap parameter is v mv x11, x20        # second swap parameter is j jal x1, swap       # call swap </pre>
Inner loop		<pre> addi x20, x20, -1 # j for2tst j for2tst         # go to for2tst </pre>
Outer loop	exit2:	<pre> addi x19, x19, 1  # i += 1 j for1tst        # go to for1tst </pre>
Restoring registers		
	exit1:	<pre> ld x19, 0(sp)     # restore x19 from stack ld x20, 8(sp)     # restore x20 from stack ld x21, 16(sp)    # restore x21 from stack ld x22, 24(sp)    # restore x22 from stack ld x1, 32(sp)     # restore return address from stack addi sp, sp, 40   # restore stack pointer </pre>
Procedure return		
		<pre> jalr x0, 0(x1)    # return to calling routine </pre>

# Effect of Compiler Optimization






Compiled with gcc for Pentium 4 under Linux



Bubble sort on 100,000 32-bit words initialized to random values.

Pentium 4 @3.06 GHz, 533 MHz system bus with 2 GB of PC2100 DDR SDRAM. Linux 2.4.20.

# Effect of Language and Algorithm





# Lessons Learnt



- › Instruction count and CPI are not good performance indicators in isolation
- › Compiler optimizations are sensitive to the algorithm
- › Java/JIT compiled code is significantly faster than JVM interpreted
  - Comparable to optimized C in some cases
- › Nothing can fix a dumb algorithm!

# Arrays vs. Pointers

---



- > Array indexing involves
  - Multiplying index by element size
  - Adding to array base address
- > Pointers correspond directly to memory addresses
  - Can avoid indexing complexity

# Example: Clearing an Array



```
clear1(int array[], int size) {
    int i;
    for (i = 0; i < size; i += 1)
        array[i] = 0;
}
```

```
clear2(int *array, int size) {
    int *p;
    for (p = &array[0]; p < &array[size];
        p = p + 1)
        *p = 0;
}
```

```
li    x5,0          // i = 0
loop1:
slli  x6,x5,3       // x6 = i * 8
add   x7,x10,x6     // x7 = address
                        // of array[i]
sd    x0,0(x7)     // array[i] = 0
addi  x5,x5,1      // i = i + 1
blt   x5,x11,loop1 // if (i<size)
                        // go to loop1
```

```
mv    x5,x10       // p = address
                        // of array[0]
slli  x6,x11,3     // x6 = size * 8
add   x7,x10,x6    // x7 = address
                        // of array[size]
loop2:
sd    x0,0(x5)     // Memory[p] = 0
addi  x5,x5,8      // p = p + 8
bltu  x5,x7,loop2
                        // if (p<&array[size])
                        // go to loop2
```

# Comparison of Array vs. Ptr

---



- › Multiply “strength reduced” to shift
- › Array version requires shift to be inside loop
  - Part of index calculation for incremented  $i$
  - c.f. incrementing pointer
- › Compiler can achieve same effect as manual use of pointers
  - Induction variable elimination
  - Better to make program clearer and safer

# MIPS Instructions



- › MIPS: commercial predecessor to RISC-V
- › Similar basic set of instructions
  - 32-bit instructions
  - 32 general purpose registers, register 0 is always 0
  - 32 floating-point registers
  - Memory accessed only by load/store instructions
    - › Consistent use of addressing modes for all data sizes
- › Different conditional branches
  - For <, <=, >, >=
  - RISC-V: blt, bge, bltu, bgeu
  - MIPS: slt, sltu (set less than, result is 0 or 1)
    - › Then use beq, bne to complete the branch

# Instruction Encoding



## Register-register

	31	25 24	20 19	15 14	12 11	7 6	0
RISC-V	funct7(7)		rs2(5)	rs1(5)	funct3(3)	rd(5)	opcode(7)
	31	26 25	21 20	16 15	11 10	6 5	0
MIPS	Op(6)		Rs1(5)	Rs2(5)	Rd(5)	Const(5)	Opx(6)

## Load

	31	20 19	15 14	12 11	7 6	0	
RISC-V	immediate(12)			rs1(5)	funct3(3)	rd(5)	opcode(7)
	31	26 25	21 20	16 15		0	
MIPS	Op(6)		Rs1(5)	Rs2(5)	Const(16)		

## Store

	31	25 24	20 19	15 14	12 11	7 6	0
RISC-V	immediate(7)		rs2(5)	rs1(5)	funct3(3)	immediate(5)	opcode(7)
	31	26 25	21 20	16 15		0	
MIPS	Op(6)		Rs1(5)	Rs2(5)	Const(16)		

## Branch

	31	25 24	20 19	15 14	12 11	7 6	0
RISC-V	immediate(7)		rs2(5)	rs1(5)	funct3(3)	immediate(5)	opcode(7)
	31	26 25	21 20	16 15		0	
MIPS	Op(6)		Rs1(5)	Opx/Rs2(5)	Const(16)		

# The Intel x86 ISA



- › Evolution with backward compatibility
  - 8080 (1974): 8-bit microprocessor
    - › Accumulator, plus 3 index-register pairs
  - 8086 (1978): 16-bit extension to 8080
    - › Complex instruction set (CISC)
  - 8087 (1980): floating-point coprocessor
    - › Adds FP instructions and register stack
  - 80286 (1982): 24-bit addresses, MMU
    - › Segmented memory mapping and protection
  - 80386 (1985): 32-bit extension (now IA-32)
    - › Additional addressing modes and operations
    - › Paged memory mapping as well as segments

# The Intel x86 ISA



- › Further evolution...
  - i486 (1989): pipelined, on-chip caches and FPU
    - › Compatible competitors: AMD, Cyrix, ...
  - Pentium (1993): superscalar, 64-bit datapath
    - › Later versions added MMX (Multi-Media eXtension) instructions
    - › The infamous FDIV bug
  - Pentium Pro (1995), Pentium II (1997)
    - › New microarchitecture (see Colwell, *The Pentium Chronicles*)
  - Pentium III (1999)
    - › Added SSE (Streaming SIMD Extensions) and associated registers
  - Pentium 4 (2001)
    - › New microarchitecture
    - › Added SSE2 instructions



# The Intel x86 ISA



- › And further...
  - AMD64 (2003): extended architecture to 64 bits
  - EM64T – Extended Memory 64 Technology (2004)
    - › AMD64 adopted by Intel (with refinements)
    - › Added SSE3 instructions
  - Intel Core (2006)
    - › Added SSE4 instructions, virtual machine support
  - AMD64 (announced 2007): SSE5 instructions
    - › Intel declined to follow, instead...
  - Advanced Vector Extension (announced 2008)
    - › Longer SSE registers, more instructions
- › If Intel didn't extend with compatibility, its competitors would!
  - Technical elegance ≠ market success

# Basic x86 Registers



Name	31	0	Use
EAX	[31-bit register]		GPR 0
ECX	[31-bit register]		GPR 1
EDX	[31-bit register]		GPR 2
EBX	[31-bit register]		GPR 3
ESP	[31-bit register]		GPR 4
EBP	[31-bit register]		GPR 5
ESI	[31-bit register]		GPR 6
EDI	[31-bit register]		GPR 7
	CS	[31-bit register]	Code segment pointer
	SS	[31-bit register]	Stack segment pointer (top of stack)
	DS	[31-bit register]	Data segment pointer 0
	ES	[31-bit register]	Data segment pointer 1
	FS	[31-bit register]	Data segment pointer 2
	GS	[31-bit register]	Data segment pointer 3
EIP	[31-bit register]		Instruction pointer (PC)
EFLAGS	[31-bit register]		Condition codes

# Intel 80386



- › 16 registers
  - Only 8 x 32-bits General-Purpose Registers (GPR)!
  - 6 segment pointer registers @16-bits
  - EIP, EFLAGS 32-bit
- › Arithmetic/logical instructions must have one operand act as both a source and a destination
  - Not all GPR can be used for all operations
- › One of the operands can be in memory
  - Only one: no memory-to-memory
- › immediates may be 8, 16 or 32 bits
- › Memory addressing with 8 or 32 bits displacement

# Intel 80386



- › Data size can be 8, 16 or 32 bits
  - Default specified as a bit in the CS
  - Override with an 8-bit prefix in the instruction
  - AMD64 adds 64 bits addresses and data
- › Conditional branches check the result of a previous operation (as in MIPS)
- › Byte-aligned PC-relative branch
  - No alignment restriction for instructions
- › Instruction length varies between 1 to 15 bytes!
  - Postfix bytes specify addressing mode
  - Prefix bytes modify operation
    - › Operand length, repetition, locking, address size

# Basic x86 Addressing Modes



- › Only two operands per instruction

Source/dest operand	Second source operand
Register	Register
Register	Immediate
Register	Memory
Memory	Register
Memory	Immediate

- › Memory addressing:

Mode	Description	Register restrictions	RISC-V equivalent
Register indirect	Address is in a register.	Not ESP or EBP	<code>ld x10, 0(x11)</code>
Based mode with 8- or 32-bit displacement	Address is contents of base register plus displacement.	Not ESP	<code>ld x10, 40(x11)</code>
Base plus scaled index	The address is $\text{Base} + (2^{\text{Scale}} \times \text{Index})$ where Scale has the value 0, 1, 2, or 3.	Base: any GPR Index: not ESP	<code>slli x12, x12, 3</code> <code>add x11, x11, x12</code> <code>ld x10, 0(x11)</code>
Base plus scaled index with 8- or 32-bit displacement	The address is $\text{Base} + (2^{\text{Scale}} \times \text{Index}) + \text{Displacement}$ where Scale has the value 0, 1, 2, or 3.	Base: any GPR Index: not ESP	<code>slli x12, x12, 3</code> <code>add x11, x11, x12</code> <code>ld x10, 40(x11)</code>

# X86 instructions




Instruction	Meaning
<b>Control</b>	<b>Conditional and unconditional branches</b>
jnz, jz	Jump if condition to EIP + 8-bit offset; JNE (for JNZ), JE (for JZ) are alternative names
jmp	Unconditional jump—8-bit or 16-bit offset
call	Subroutine call—16-bit offset; return address pushed onto stack
ret	Pops return address from stack and jumps to it
loop	Loop branch—decrement ECX; jump to EIP + 8-bit displacement if ECX ≠ 0
<b>Data transfer</b>	<b>Move data between registers or between register and memory</b>
move	Move between two registers or between register and memory
push, pop	Push source operand on stack; pop operand from stack top to a register
les	Load ES and one of the GPRs from memory
<b>Arithmetic, logical</b>	<b>Arithmetic and logical operations using the data registers and memory</b>
add, sub	Add source to destination; subtract source from destination; register-memory format
cmp	Compare source and destination; register-memory format
shl, shr, rcr	Shift left; shift logical right; rotate right with carry condition code as fill
cbw	Convert byte in eight rightmost bits of EAX to 16-bit word in right of EAX
test	Logical AND of source and destination sets condition codes
inc, dec	Increment destination, decrement destination
or, xor	Logical OR; exclusive OR; register-memory format
<b>String</b>	<b>Move between string operands; length given by a repeat prefix</b>
movs	Copies from string source to destination by incrementing ESI and EDI; may be repeated
lods	Loads a byte, word, or doubleword of a string into the EAX register

# x86 Instruction Encoding



a. JE EIP + displacement




> Variable length encoding


b. CALL




c. MOV EBX, [EDI + 45]



d. PUSH ESI



e. ADD EAX, #6765



f. TEST EDX, #42



# Implementing IA-32

---



- › Complex instruction set makes implementation difficult
  - Hardware translates instructions to simpler microoperations
    - › Simple instructions: 1–1
    - › Complex instructions: 1–many
  - Microengine similar to RISC
  - Market share makes this economically viable
- › Comparable performance to RISC
  - Compilers avoid complex instructions



# Other RISC-V Instructions



## › Base integer instructions (RV64I)

Those previously described, plus:

- `auipc rd, imm` //  $rd = (imm \ll 12) + pc$ 
  - › follow by `jalr` (adds 12-bit imm) for long jump
- `slt`, `sltu`, `slti`, `sltui`: set less than (like MIPS)
- `addw`, `subw`, `addiw`: 32-bit add/sub
- `sllw`, `srlw`, `slliw`, `srliw`, `sraiw`: 32-bit shift

## › 32-bit variant: RV32I

- registers are 32-bits wide, 32-bit operations

# Instruction Set Extensions

---



- › M: integer multiply, divide, remainder
- › A: atomic memory operations
- › F: single-precision floating point
- › D: double-precision floating point
- › C: compressed instructions
  - 16-bit encoding for frequently used instructions

# Fallacies




- › Powerful instruction  $\Rightarrow$  higher performance
  - Fewer instructions required
  - But complex instructions are hard to implement
    - › May slow down all instructions, including simple ones
  - Compilers are good at making fast code from simple instructions
  
- › Use assembly code for high performance
  - But modern compilers are better at dealing with modern processors
  - More lines of code  $\Rightarrow$  more errors and less productivity

# Fallacies



- › Backward compatibility  $\Rightarrow$  instruction set doesn't change
  - But they do accrete more instructions



x86 instruction set

# Pitfalls



- › Sequential words are not at sequential addresses
  - Increment by 4, not by 1!
- › Keeping a pointer to an automatic variable after procedure returns
  - e.g., passing pointer back via an argument
  - Pointer becomes invalid when stack popped

# Concluding Remarks

---



- › Design principles
  1. Simplicity favors regularity
  2. Smaller is faster
  3. Good design demands good compromises
- › Make the common case fast
- › Layers of software/hardware
  - Compiler, assembler, hardware
- › RISC-V: typical of RISC ISAs
  - c.f. x86