

Computer Organization & Assembly Language Programming (CSE 2312)

Lecture 5: Instructions, Memory, and Endianness

Taylor Johnson

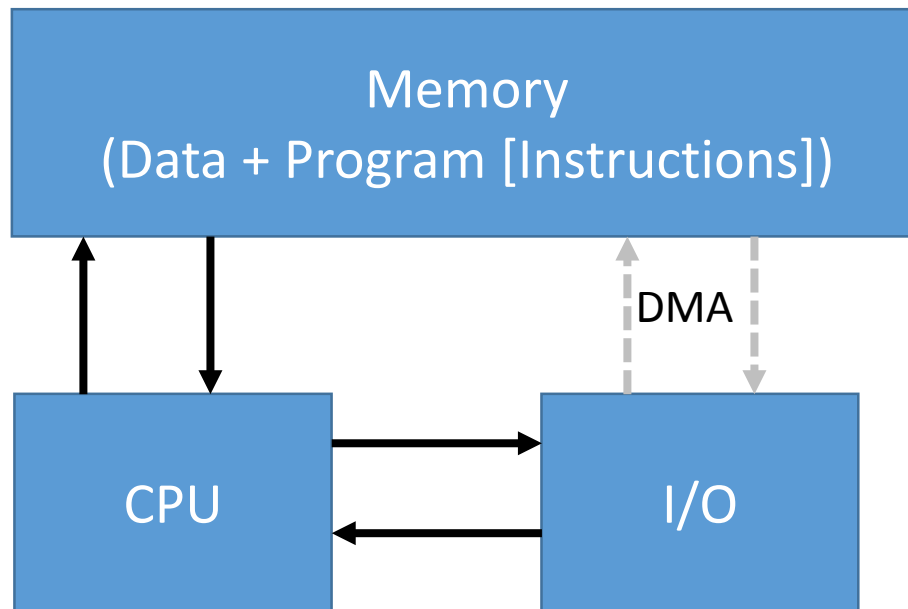
Important Concepts from Previous Lectures

- How do computers compute?
- Binary to decimal, decimal to binary, ASCII, signed numbers, hexadecimal
- Structured computers
- Performance metrics
 - Clock rates, cycle time/period, CPI, response time, throughput

Announcements and Outline

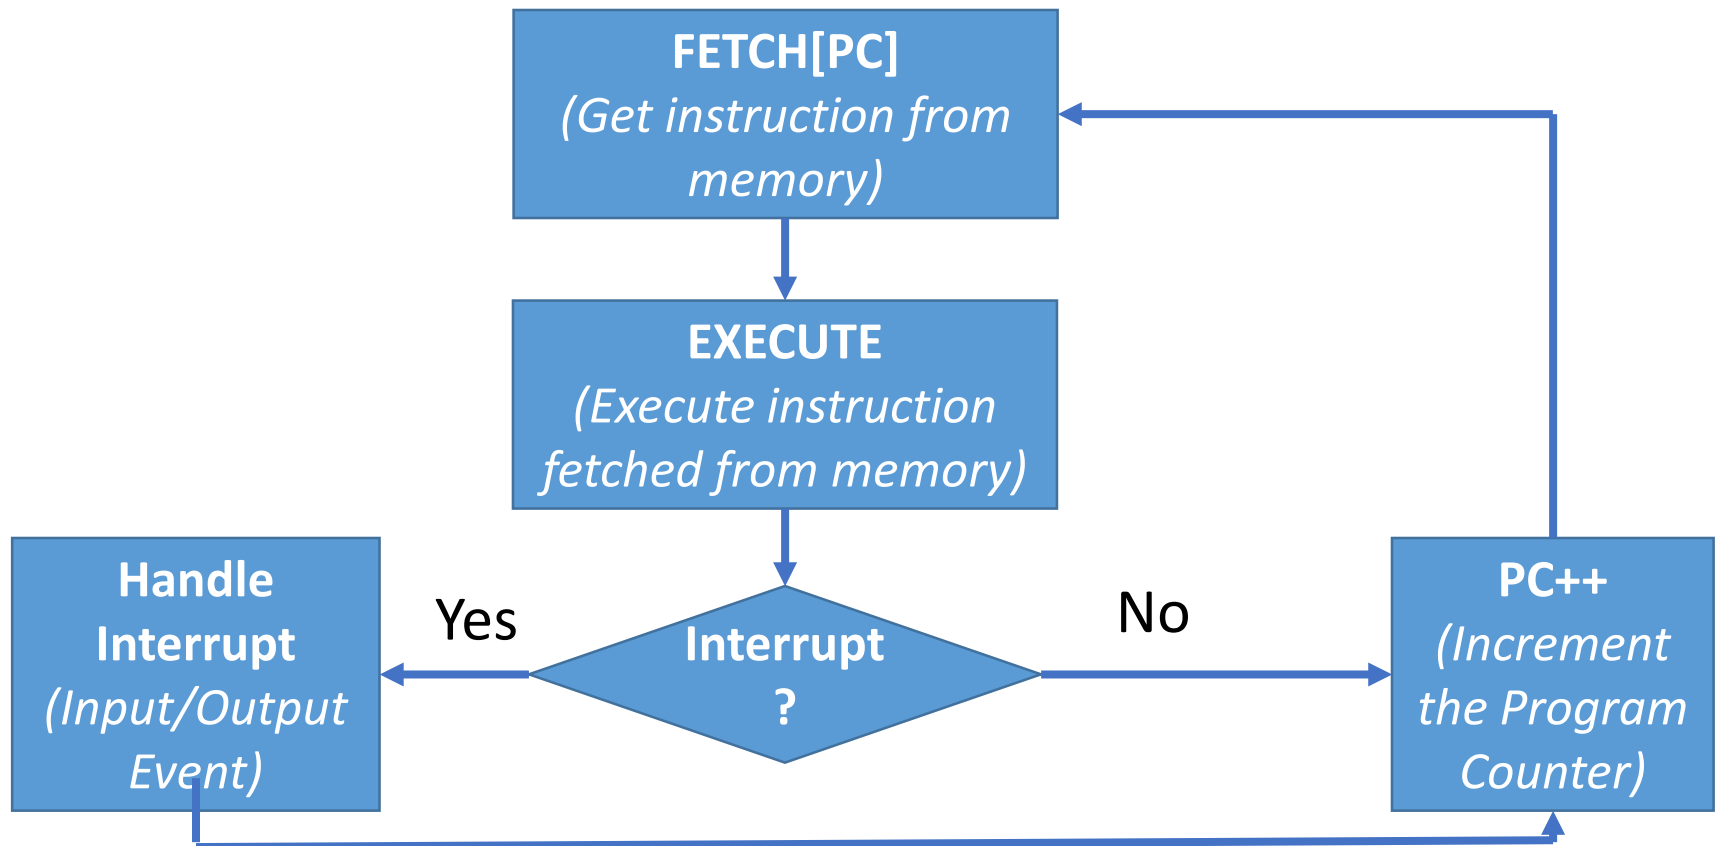
- Quiz 2 on Blackboard site (due 11:59PM Friday)
 - Review binary arithmetic, Boolean operations, and representing signed and unsigned numbers in binary
- Homework 1 due today
- Homework 2 assigned today
 - Reading chapter 2 (ARM version on Blackboard site)
- Review from last time
 - Signed vs. Unsigned Numbers (Two's Complement), Hex
 - Intro assembly
- Instructions: the Language of the Computer
- Memory
- Endianness

Review: Von Neumann Architecture



- Both data and program stored in memory
- Allows the computer to be “re-programmed”
- Input/output (I/O) goes through CPU
- I/O part is not representative of modern systems (direct memory access [DMA])
- Memory layout is representative of modern systems

Review: Abstract Processor Execution Cycle



Review: Two's Complement Signed Negation

- Complement and add 1
 - Complement means $1 \rightarrow 0, 0 \rightarrow 1$
 - Representation called one's complement

$$x + \bar{x} = 1111 \dots 111_2 = -1$$

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

- Example: negate +2
 - $+2 = 0000 \ 0000 \ \dots \ 0010_2$
 - $-2 = 1111 \ 1111 \ \dots \ 1101_2 + 1$
 $= 1111 \ 1111 \ \dots \ 1110_2$

Review: Hexadecimal

- Base 16
 - Compact representation of bit strings
 - 4 bits (also called a nibble or nybble) 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: 0xECA8 6420
 - 1110 1100 1010 1000 0110 0100 0010 0000

Review: Arithmetic Operations

- Add and subtract, three **operands**
 - **Operand**: *quantity on which an operation is performed*
 - Two sources and one destination

add a, b, c # a updated to 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

Review: Arithmetic Example

- C code:

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

- Compiled MIPS code:

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

- Compiled ARM code:

```
add r0, g, h    # temp r0 = g + h
add r1, i, j    # temp r1 = i + j
sub f, r0, r1   # f = t0 - t1
```

- Notice: registers “=” variables


Review: Some Processor Components

CPU

Register File

- Program Counter (PC)
- Instruction Register (IR)
- General Purpose Registers
 - Word size
 - Typically 16-32 of these
 - PC sometimes one of these
- Floating Point Registers

Arithmetic logic unit (ALU)



Floating Point Unit (FPU)

Review: ARM: Load/Store Architecture

- ARM is a load/store architecture
- This means that memory can only be accessed by load and store instructions
- All arguments for arithmetic and logical instructions must either:
 - Come from registers
 - Be constants specified within the instruction
 - (more examples of that later)
- This may not seem like a big deal to you, as you have not experienced the alternative
 - However, it makes life much easier
 - This is one reason why we chose ARM 7 for this course

ARM 7 Registers

- 16 32-bit general purpose registers
- 32 32-bit floating-point registers (not available on every device)

Register	Alt. name	Function
R0–R3	A1–A4	Holds parameters to the procedure being called
R4–R11	V1–V8	Holds local variables for the current procedure
R12	IP	Intraprocedure call register (for 32-bit calls)
R13	SP	Stack pointer
R14	LR	Link register (return address for current function)
R15	PC	Program counter

Version 7 ARM's general registers.

ARM 7 Registers

- The Vx registers hold data needed by procedures (functions)
- They should be stored in memory when calling another procedure
- They should be restored from memory when returning from another procedure

Register	Alt. name	Function
R0–R3	A1–A4	Holds parameters to the procedure being called
R4–R11	V1–V8	Holds local variables for the current procedure
R12	IP	Intraprocedure call register (for 32-bit calls)
R13	SP	Stack pointer
R14	LR	Link register (return address for current function)
R15	PC	Program counter

Version 7 ARM's general registers.

ARM 7 Registers

- The Ax registers are used for passing parameters to procedures
- Four dedicated registers have special roles: IP, SP, LR, PC.
 - We will see more details on these registers are later.
- Who ensures that these registers are used as specified here?

Register	Alt. name	Function
R0–R3	A1–A4	Holds parameters to the procedure being called
R4–R11	V1–V8	Holds local variables for the current procedure
R12	IP	Intraprocedure call register (for 32-bit calls)
R13	SP	Stack pointer
R14	LR	Link register (return address for current function)
R15	PC	Program counter

Version 7 ARM's general registers.

ARM 7 Registers

- The Ax registers are used for passing parameters to procedures
- Four dedicated registers have special roles: IP, SP, LR, PC.
 - We will see more details on these registers are later
- Who ensures that these registers are used as specified here?
 - **You!!! (The programmer)**

Register	Alt. name	Function
R0–R3	A1–A4	Holds parameters to the procedure being called
R4–R11	V1–V8	Holds local variables for the current procedure
R12	IP	Intraprocedure call register (for 32-bit calls)
R13	SP	Stack pointer
R14	LR	Link register (return address for current function)
R15	PC	Program counter

Version 7 ARM's general registers.

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
- Words are aligned in memory
 - Address must be a multiple of 4
- MIPS/ARM are Big Endian
 - Most-significant byte at least address of a word
 - *c.f.* Little Endian: least-significant byte at least address

Register Operand Example

- C code:

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

- f, ..., j in:

- \$s0, ..., \$s4 (MIPS)
 - r0, ..., r4 (ARM)

- Compiled MIPS code:

```
add $t0, $s1, $s2
add $t1, $s3, $s4
sub $s0, $t0, $t1
```

- Compiled ARM code:

```
add r0, r1, r2
add r1, r3, r4 // overwrite r1
sub r0, r0, r1
```

- Note: syntax and semantics (meaning) differences

Representing Instructions

- Instructions are encoded in binary
 - Called machine code
- ARM (and MIPS) instructions
 - Encoded as 32-bit instruction words
 - Small number of formats encoding operation code (opcode), register numbers, ...
 - Regularity!
- Register numbers
 - r0 referenced (addressed) by `b0000`
 - r1 referenced by `b0001`
 - ...
 - r15 referenced by `b1111`

ARM Arithmetic Instructions in Machine Language

Cond	F	I	Opcode	S	Rn	Rd	Operand2
4 bits	2 bits	1 bit	4 bits	1 bit	4 bits	4 bits	12 bits

- Opcode: Basic operation of the instruction
- Rd: The register destination operand. It gets the result of the operation
- Rn: The first register source operand
- Operand2: The second source operand
- I: Immediate. If I is 0, the second source operand is a register. If I is 1, the second source operand is a 12-bit immediate
- S: Set Condition Code. This field is related to conditional branch instructions
- Cond: Condition. Related to conditional branch instructions
- F: Instruction Format. This field allows ARM to different instruction formats when needed

ARM Arithmetic Instructions in Machine Language

Cond	F	I	Opcode	S	Rn	Rd	Operand2
1110	00	0	0100	0	0001	0101	0000 0000 0010

4 bits 2 bits 1 bit 4 bits 1 bit 4 bits 4 bits 12 bits

- Example: `add r5, r1, r2`
- C equivalent: `r5 = r1 + r2`
- Machine language encoding above
- Opcode: 0100 means add (dependent on digital logic, some encoding)
- Rd: register destination operand. It gets the result of the operation
- Rn: first register source operand
- Operand2: second source operand
- I: Immediate. If I is 0, the second source operand is a register. If I is 1, the second source operand is a 12-bit immediate
- S: Set Condition Code
- Cond: Condition. Related to conditional branch instructions
- F: Instruction Format

Machine Code: MIPS R-format Instructions



- Instruction fields

- op: operation code (opcode)
- rs: first source register number
- rt: second source register number
- rd: destination register number
- shamt: shift amount (00000 for now)
- funct: function code (extends opcode)

Machine Code: R-format Example

op	rs	rt	rd	shamt	funct
6 bits	5 bits	5 bits	5 bits	5 bits	6 bits

add \$t0, \$s1, \$s2

special	\$s1	\$s2	\$t0	0	add
0	17	18	8	0	32
000000	10001	10010	01000	00000	100000

$$00000010001100100100000000100000_2 = 02324020_{16}$$

Machine Code: MIPS I-format Instructions



- Immediate arithmetic and load/store instructions
 - rt: destination or source register number
 - Constant: -2^{15} to $+2^{15} - 1$
 - Address: offset added to base address in rs
- *Design Principle 4: Good design demands good compromises*
 - Different formats complicate decoding, but allow 32-bit instructions uniformly
 - Keep formats as similar as possible

Memory Cells and Addresses

- **Memory cell:** a piece of memory that contains a specific number of bits
 - How many bits depends on the architecture
 - In modern architectures, it is almost universal that a cell contains *8 bits (1 byte)*, and that will be also our convention in this course
- **Memory address:** a number specifying a location of a memory cell containing data
 - Essentially, a number specifying the location of a byte of memory

Memory Cells and Addresses

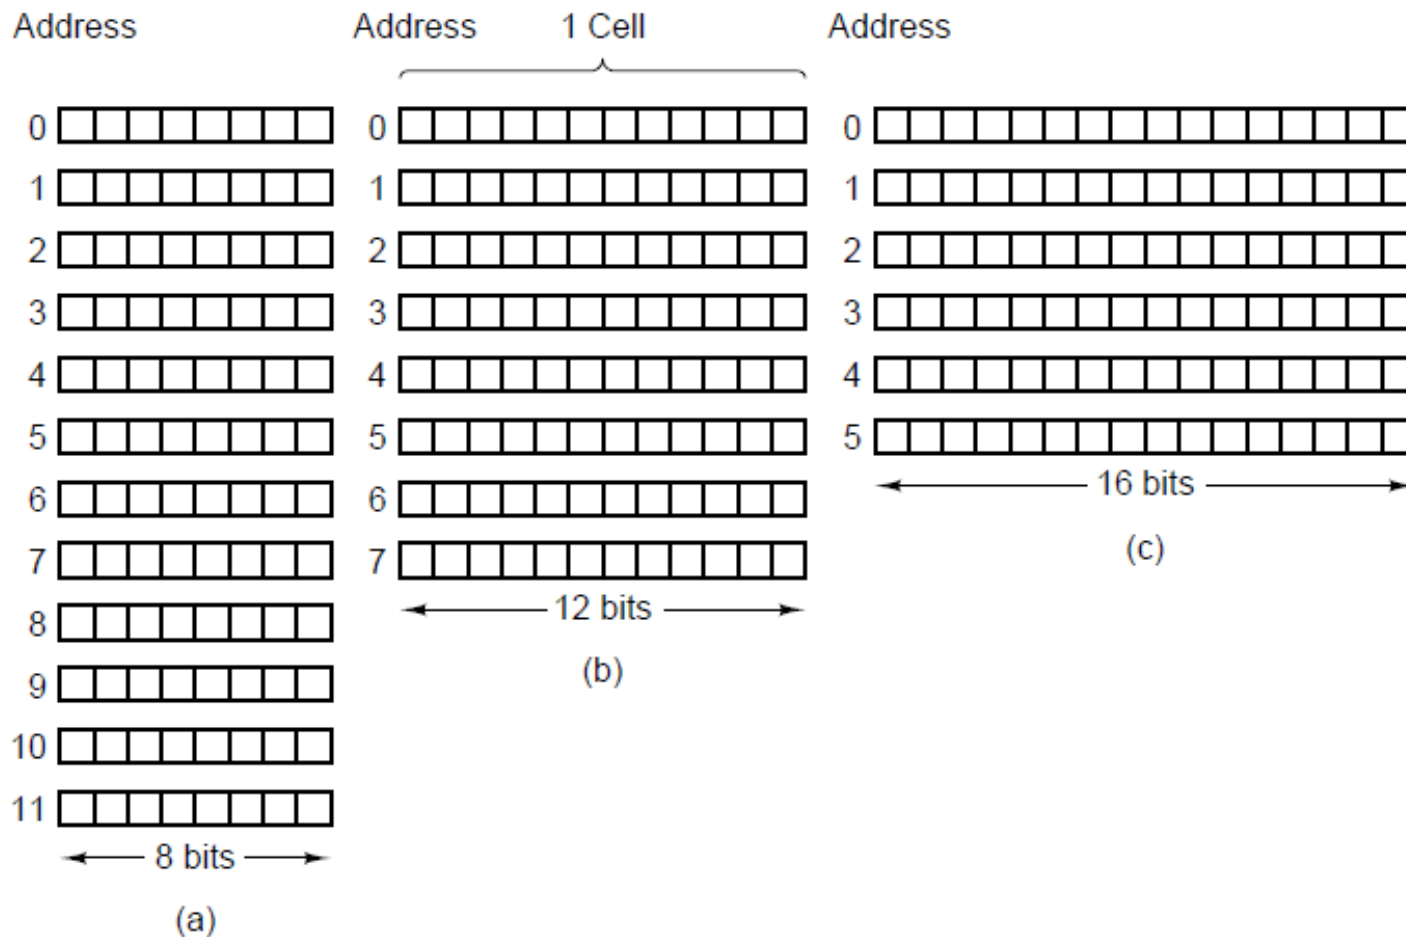
- The number of unique memory addresses depends on the size of the memory and the size of each cell
- For example, suppose we have a 96-bit memory.
- If each cell is 8 bits, we have ??? addresses?
- If each cell is 12 bits, we have ??? addresses?
- If each cell is 16 bits, we have ??? addresses?

Memory Cells and Addresses

- The number of unique memory addresses depends on the size of the memory and the size of each cell
- For example, suppose we have a 96-bit memory.
- If each cell is 8 bits, we have 12 addresses?
- If each cell is 12 bits, we have 8 addresses?
- If each cell is 16 bits, we have 6 addresses?

- Convention used almost everywhere, and in this course: if a memory has n cells, the addresses of these cells will be from **0** to **$n-1$** .

Cells and Addresses



Three ways of organizing a 96-bit memory

Words

- These days, memory cells are 8-bit
- That is why we have a special term for 8 bits, and we call them a **byte**
 - The term **octet** is also used instead of a byte
- The next memory unit is a **word**
- The size of a word is not universally fixed, but it depends on the architecture
- A 32-bit architecture has 4-byte words
- A 64-bit architecture (standard on PCs these days) has 8-byte words

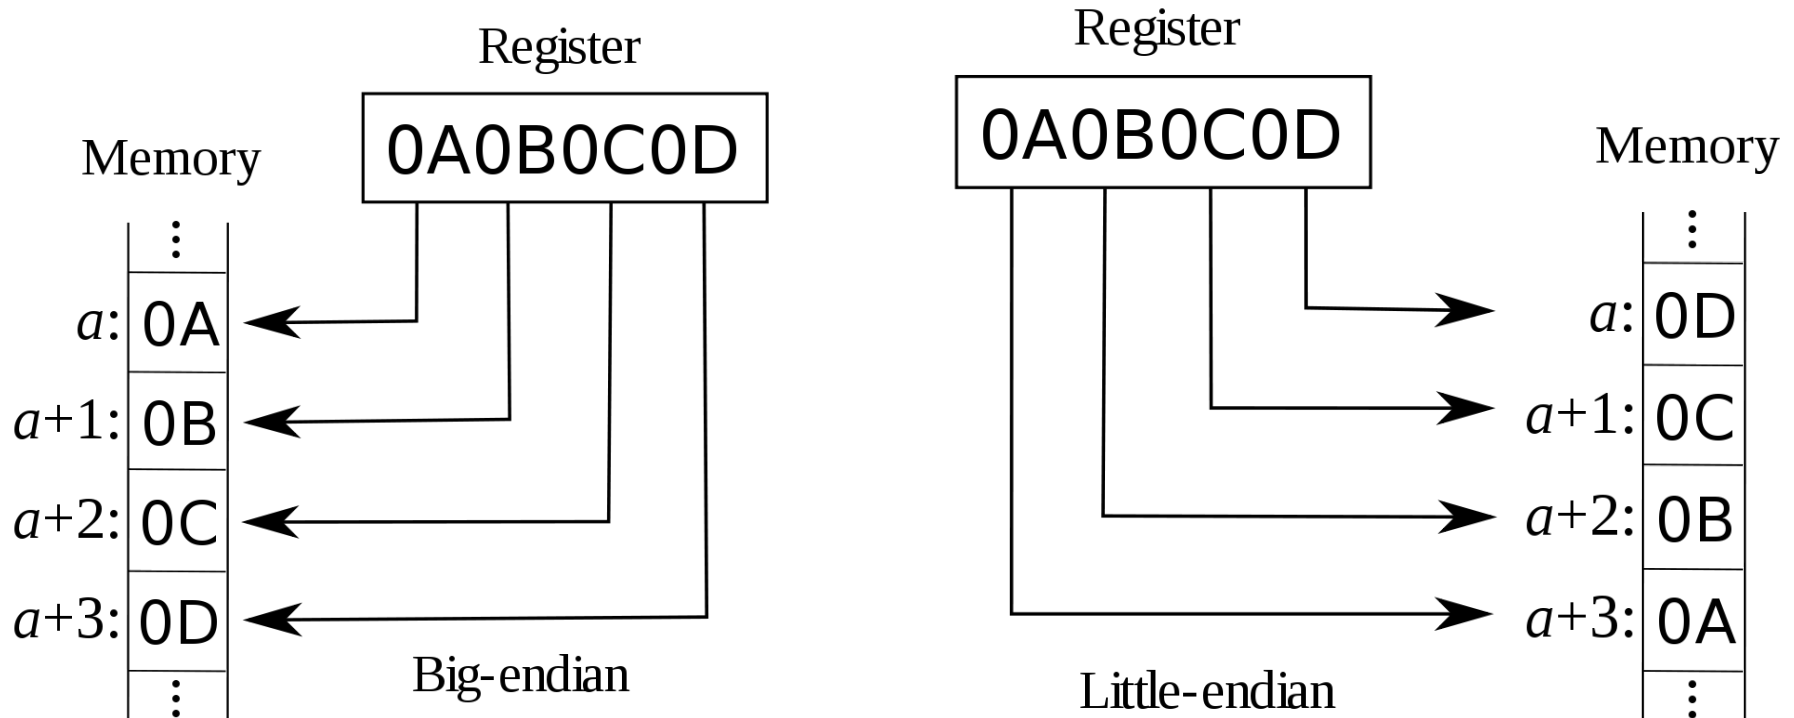
Byte Ordering - Endianness

- How do we store an integer in memory?
- Simple answer: in binary
- Actual answer: yes, in binary, but this does not fully specify how we store the number
- Unfortunately, we have two choices
- Common architectures may follow either choice, and mess ensues, unless we are aware of this issue and we deal with it explicitly
- This is the problem of **endianness**

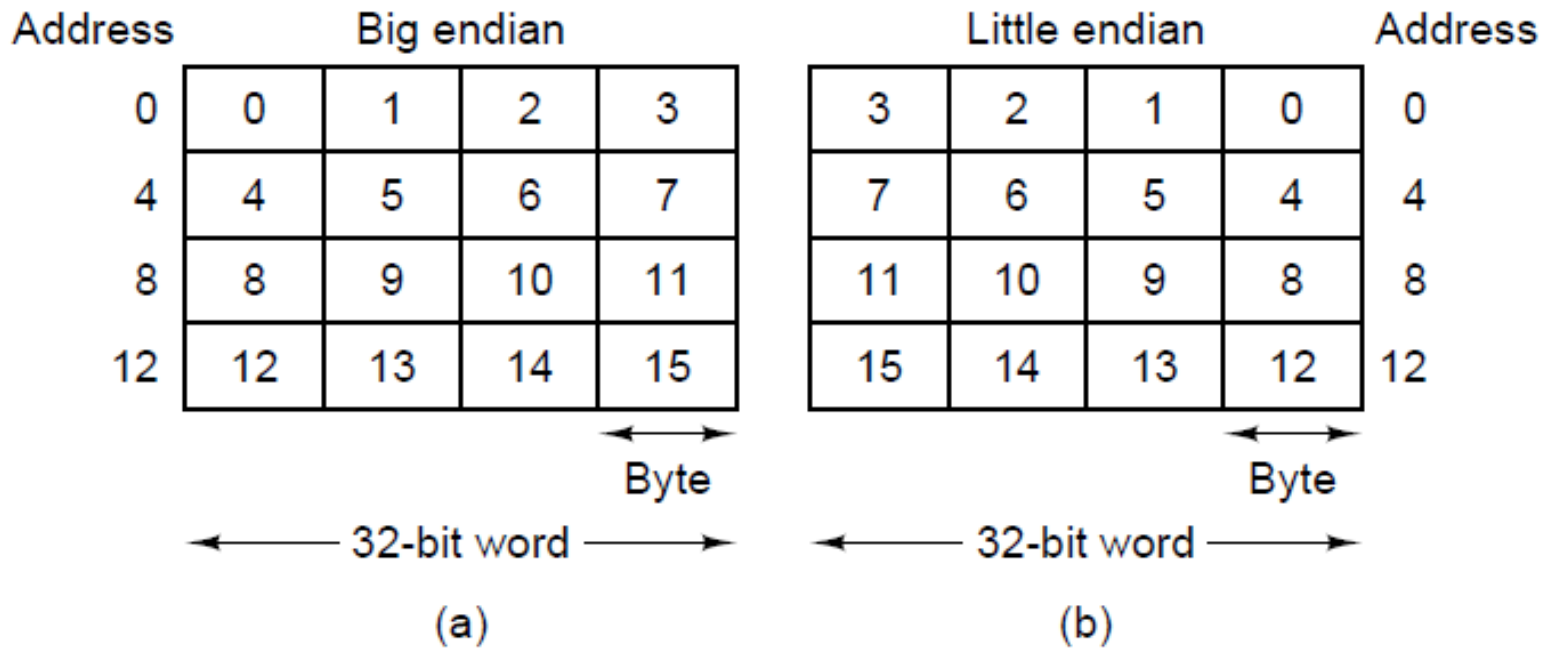
Endianness

- Little-endian: increasing numeric significance with increasing memory addresses
- Big-endian: decreasing numeric significance with increasing memory addresses
- Little-Endian Examples
 - x86, x86-64, 8051, DEC Alpha, Atmel AVR
- Big-Endian Examples
 - Motorola 6800 and 68k series, Xilinx Microblaze, IBM POWER, and System/360, MIPS
- Bi-Endianness
 - Ability for computer to operate using either
 - SPARC
 - ARM architecture: little-endian before version 3, now bi-endian

Endianness Example



Byte Ordering Visualization



(a) Big endian memory. (b) Little endian memory.


Main difference: ordering of bytes in a word

- Left-to-right in big endian.
- Right-to-left in little-endian.

Memory: Words and Alignment

- Bytes are grouped into **words**
- Depending on the machine, a word can be:
 - 32 bits (4 bytes) , or
 - 64 bits (8 bytes), or ... (16-bits, 128 bits, etc.)
- Oftentimes it is required that words are aligned
- This means that:
 - 4-byte words can only begin at memory addresses that are multiples of 4: 0, 4, 8, 12, 16...
 - 8-byte words can only begin at memory addresses that are multiples of 8: 0, 8, 16, 24, 32, ...

Memory Models



An 8-byte word in a little-endian memory.
 (a) Aligned. (b) Not aligned. Some machines require that words in memory be aligned.

Memory as an Array

- Think of memory and addressing like you think of arrays

```
MEM[ ADDR-1 ] 0x05
```

```
MEM[ ADDR ] 0xAB
```

```
MEM[ ADDR+1 ] 0xF1
```

Suppose ADDR = 0x1000

```
MEM[ 0x0FFF ] 0x05
```

```
MEM[ 0x1000 ] 0xAB
```

```
MEM[ 0x1001 ] 0xF1
```

```
MEM[ ... ] ...
```

How large is this memory?

Address Spaces For Instructions and Data

- Typically memory can be accessed using a single address space
 - For example, if we have 4 GB of memory, each byte has an address from 0 to $2^{32} - 1$
 - Each memory location may store instructions at some point and data at some other point
- An alternative is to have separate address spaces for instructions and data
 - In that case, a memory location is permanently dedicated to either storing instructions or to storing data
 - Instead of a single **load** instruction, we have **load_instructions** and **load_data**

Effects of Separate Address Spaces

- If A is a valid memory address, `load_instructions A` and `load_data A` access different memory locations.
 - `load_instructions A` accesses address A in the instructions space.
 - `load_data A` accesses address A in the data space.
- This makes it harder for malware to cause trouble. Why?

Effects of Separate Address Spaces


- If A is a valid memory address, `load_instructions A` and `load_data A` access different memory locations.
 - `load_instructions A` accesses address A in the instructions space.
 - `load_data A` accesses address A in the data space.
- This makes it harder for malware to cause trouble. Why?
- A common way for malware to attack is to:
 - Run as regular program.
 - Modify memory locations that store instructions, thus modifying other programs (such as the operating system).
- If instruction memory is accessed with different instructions, such behavior can easily be prevented.

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!

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



Operands Types

- Register operand: operand comes from the binary valued stored in a particular register in the CPU
 - Example: `add r0, r1, r2`
 - C code: `r0 = r1 + r2;`
- Immediate operand: operand value comes from instruction itself
 - Example: `add r0, r1, #1`
 - C code: `r0 = r1 + 1;`
- Memory operand: operand refers to memory
 - Example: `str r0, [r1]`
 - C code (roughly): `MEM[r1] = r0;`
 - Only for load / store instructions!
 - Several addressing modes (more on this later)

Memory Operand Example 1

- C code:

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

- g in r1, h in r2, base address of A in r3


- Compiled ARM code:

- Index 8 requires offset of 8 words
 - 4 bytes per word

```
@ load word
```

```
ldr r0, [r3, #32] @ r0 = MEM[r3 + 32]
```

```
add r1, r2, r0
```



Memory Operand Example 2

- C code:

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

- h in r2, base address of A in r3

- Compiled ARM code:

- Index 8 requires offset of 32 (8 bytes, 4 bytes per word)

```
@ load word
```

```
ldr r0, [r3, #32] @ r0 = MEM[r3 + 32]
```

```
add r0, r2, r0
```

```
@ store word
```

```
str r0, [r3, #48] @ MEM[r3 + 48] = r0
```

Immediate Operands

- Constant data specified in an instruction
`add r3, r3, #4`
- *Design Principle 3: Make the common case fast*
 - Small constants are common
 - Immediate operand avoids a load instruction

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

Summary

- Operations (instructions), operands
- Machine encoding of assembly
- Assembly to C and C to assembly (for basic examples, no functions yet)
- Memory
- Endianness

Questions?

