

COMPUTER ORGANIZATION AND DESIGN The Hardware/Software Interface



Chapter 1

Computer Abstractions and Technology

The Computer Revolution

- Progress in computer technology
 - Underpinned by Moore's Law
- Makes novel applications feasible
 - Computers in automobiles
 - Cell phones
 - Human genome project
 - World Wide Web
 - Search Engines
- Computers are pervasive

Classes of Computers

- Desktop computers
 - General purpose, variety of software
 - Subject to cost/performance tradeoff
- Server computers
 - Network based
 - High capacity, performance, reliability
 - Range from small servers to building sized
 - Embedded computers
 - Hidden as components of systems
 - Stringent power/performance/cost constraints

The Processor Market



Chapter 1 — Computer Abstractions and Technology — 4

What You Will Learn

- How programs are translated into the machine language
 - And how the hardware executes them
- The hardware/software interface
- What determines program performance
 - And how it can be improved
- How hardware designers improve performance
- What is parallel processing

Understanding Performance

- Algorithm
 - Determines number of operations executed
- Programming language, compiler, architecture
 - Determine number of machine instructions executed per operation
- Processor and memory system
 - Determine how fast instructions are executed
- I/O system (including OS)
 - Determines how fast I/O operations are executed

Below Your Program

- Application software
 - Written in high-level language
- System software



- Compiler: translates HLL code to machine code
- Operating System: service code
 - Handling input/output
 - Managing memory and storage
 - Scheduling tasks & sharing resources

Hardware

Processor, memory, I/O controllers



Levels of Program Code

High-level language

- Level of abstraction closer to problem domain
- Provides for productivity and portability
- Assembly language
 - Textual representation of instructions
- Hardware representation
 - Binary digits (bits)
 - Encoded instructions and data



Components of a Computer





- Same components for all kinds of computer
 - Desktop, server, embedded

Input/output includes

- User-interface devices
 - Display, keyboard, mouse
- Storage devices
 - Hard disk, CD/DVD, flash
- Network adapters
 - For communicating with other computers

Anatomy of a Computer



Anatomy of a Mouse

- Optical mouse
 - LED illuminates desktop
 - Small low-res camera
 - Basic image processor
 - Looks for x, y movement
 - Buttons & wheel
- Supersedes roller-ball mechanical mouse







Chapter 1 — Computer Abstractions and Technology — 12

Opening the Box





Inside the Processor (CPU)

- Datapath: performs operations on data
- Control: sequences datapath, memory, ...
 - Cache memory
 - Small fast SRAM memory for immediate access to data



Inside the Processor

AMD Barcelona: 4 processor cores





Chapter 1 — Computer Abstractions and Technology — 15

Abstractions

The BIG Picture

- Abstraction helps us deal with complexity
 Hide lower-level detail
- Instruction set architecture (ISA)
 - The hardware/software interface
- Application binary interface
 - The ISA plus system software interface
- Implementation
 - The details underlying and interface

A Safe Place for Data

- Volatile main memory
 - Loses instructions and data when power off
- Non-volatile secondary memory
 - Magnetic disk
 - Flash memory
 - Optical disk (CDROM, DVD)









Networks

- Communication and resource sharing
 Local area network (LAN): Ethernet
 Within a building
- Wide area network (WAN: the Internet)
- Wireless network: WiFi, Bluetooth





Technology Trends

- Electronics technology continues to evolve
 - Increased capacity and performance
 - Reduced cost



DRAM capacity

Year	Technology	Relative performance/cost
1951	Vacuum tube	1
1965	Transistor	35
1975	Integrated circuit (IC)	900
1995	Very large scale IC (VLSI)	2,400,000
2005	Ultra large scale IC	6,200,000,000

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 MIPS Instruction Set

- Used as the example throughout the book
- Stanford MIPS commercialized by MIPS Technologies (<u>www.mips.com</u>)
- Large share of embedded core market
 - Applications in consumer electronics, network/storage equipment, cameras, printers, ...
- Typical of many modern ISAs
 - See MIPS Reference Data tear-out card, and Appendixes B and E

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 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



Register Operands

- Arithmetic instructions use register operands
- MIPS has a 32 × 32-bit register file
 - Use for frequently accessed data
 - Numbered 0 to 31
 - 32-bit data called a "word"
- Assembler names
 - \$t0, \$t1, ..., \$t9 for temporary values
 - \$s0, \$s1, ..., \$s7 for saved variables
 - Design Principle 2: Smaller is faster
 - c.f. main memory: millions of locations

Register Operand Example

C code: f = (g + h) - (i + j);
f, ..., j in \$s0, ..., \$s4
Compiled MIPS code: add \$t0, \$s1, \$s2 add \$t1, \$s3, \$s4 sub \$s0, \$t0, \$t1



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 is Big Endian
 - Most-significant byte at least address of a word
 - *c.f.* Little Endian: least-significant byte at least address

Memory Operand Example 1

- C code:
 - g = h + A[8];
 - g in \$s1, h in \$s2, base address of A in \$s3
- Compiled MIPS code:
 - Index 8 requires offset of 32
 - 4 bytes per word

Memory Operand Example 2

C code: A[12] = h + A[8];h in \$s2, base address of A in \$s3 Compiled MIPS code: Index 8 requires offset of 32 lw \$t0, 32(\$s3) # load word add \$t0, \$s2, \$t0 sw \$t0, 48(\$s3) # store word

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 \$s3, \$s3, 4
- No subtract immediate instruction
 - Just use a negative constant addi \$s2, \$s1, -1
 - Design Principle 3: Make the common case fast
 - Small constants are common
 - Immediate operand avoids a load instruction

The Constant Zero

- MIPS register 0 (\$zero) is the constant 0
 - Cannot be overwritten
- Useful for common operations
 - E.g., move between registers add \$t2, \$s1, \$zero



Unsigned Binary Integers

Given an n-bit number

- Range: 0 to $+2^{n} 1$
- Example
 - 0000 0000 0000 0000 0000 0000 0000 1011₂ = 0 + ... + $1 \times 2^3 + 0 \times 2^2 + 1 \times 2^1 + 1 \times 2^0$
 - $= 0 + \dots + 8 + 0 + 2 + 1 = 11_{10}$
- Using 32 bits
 - 0 to +4,294,967,295

2s-Complement Signed Integers

Given an n-bit number

$$\mathbf{x} = -\mathbf{x}_{n-1}\mathbf{2}^{n-1} + \mathbf{x}_{n-2}\mathbf{2}^{n-2} + \dots + \mathbf{x}_{1}\mathbf{2}^{1} + \mathbf{x}_{0}\mathbf{2}^{0}$$

Range:
$$-2^{n-1}$$
 to $+2^{n-1} - 1$

Example

Using 32 bits

■ -2,147,483,648 to +2,147,483,647

2s-Complement Signed Integers

- Bit 31 is sign bit
 - I 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 $1 \rightarrow 0, 0 \rightarrow 1$

$$x + \overline{x} = 1111...111_{2} = -1$$

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

Sign Extension

- Representing a number using more bits
 - Preserve the numeric value
- In MIPS instruction set
 - addi: extend immediate value
 - Ib, Ih: extend loaded byte/halfword
 - beq, bne: extend the displacement
- 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

Representing Instructions

- Instructions are encoded in binary
 - Called machine code
- MIPS instructions
 - Encoded as 32-bit instruction words
 - Small number of formats encoding operation code (opcode), register numbers, ...
 - Regularity!
- Register numbers
 - \$t0 \$t7 are reg's 8 15
 - \$t8 \$t9 are reg's 24 25
 - \$s0 \$s7 are reg's 16 23