

Slide 1

Administrivia

- Reminder: Homework 1 due Wednesday. Hard copy, in class or in my mailbox by 5:30pm.

Slide 2

Minute Essay From Last Lecture

- (Review.) Most people got the point, but not everyone.

Measuring Performance — Recap/Review

Slide 3

- Many, many factors influence execution time for programs, from choice of algorithm to “processor speed” to system load, as discussed previously.
- Textbook chooses to focus in this chapter on “execution time” by which the authors mean processor time only, excluding delays caused by other factors. Might not be meaningful for comparing systems but seems like reasonable way to compare processors at least.
- Slides from last time updated to include additional example. (Review.)

What's Next — Overview

Slide 4

- Defining a representative architecture (MIPS): what “architecture” means in context, assembly language programming, machine language. (This is the “first half” of the course.)
- Implementing this architecture. (This is the “second half”.)

Slide 5

“Architecture” as Interface Definition

- “Architecture” here means “instruction set architecture” (ISA), a key abstraction.
- From software perspective, “architecture” defines lowest-level building blocks: what operations are possible, what kinds of operands, binary data formats, etc.
- From hardware perspective, “architecture” is a specification: Designers must build something that behaves the way the specification says.

Slide 6

Architecture — Key Abstractions

- Memory: Long long list of binary “numbers”, encoding all data (including programs!), each with “address” and “contents”.
When running a program, program itself is in memory; so is its data.
- Instructions: Primitive operations processor can perform.
- Fetch/execute cycle: What the processor does to execute a program — repeatedly get next instruction (from memory, location defined by “program counter”), increment program counter, execute instruction.
- Registers: Fast-access work space for processor, typically divided into “special-purpose” (e.g., program counter), “general-purpose” (integer and floating-point). Unlike memory, these are part of the processor.

Slide 7

Design Goals for Instruction Set

- From software perspective — expressivity.
- From hardware perspective — good performance, low cost.
- (Yes, these can sometimes be opposing forces!)

Slide 8

Why Study MIPS Architecture?

- Goal is not to become assembly-language programmers, but to understand how things work at this level. Once you understand basic principles, learning another assembly language is easier.
 - MIPS architecture is simple but representative.
- Aside: SPIM simulator will let you experiment (commands `spim` and `xspim`).

A Bit About Assembly Language Syntax

- Syntax for high-level languages can be complex. Allows for good expressivity, but translation into processor instructions is complicated.
- Syntax for assembly language, in contrast, is very simple. Less expressivity but much easier to translate into (binary form of) instructions.

Slide 9

Arithmetic Instructions — Addition

- Instruction for integer addition (in assembly-language form):

```
add    r1, r2, r3
```

Adds `r2` and `r3` giving `r1`.

(Notice the format — symbolic name, operands.)

- Is this expressive enough?
- Should we have more instructions (with different numbers of operands, e.g.)?
Basic principle: “Simplicity favors regularity.”
Easier to build simple hardware if ISA is “regular” — e.g., all arithmetic instructions have exactly three operands.
- `sub` (subtraction) is similar. Multiplication and division are more complicated, so punt for now.
- What are the operands? Registers. What are those? Well ...

Slide 10

Slide 11

Registers

- Access to main memory is slow compared to processor speed, so it's useful to have a within-the-chip memory — “registers”.
- MIPS architecture defines 32 “general-purpose” registers, each 32 bits.
- Would more be better?
Basic principle: “Smaller is faster.”
- In machine language, reference by number.
- In assembly language, useful to adopt conventions for which registers to use for what, define symbolic names indicating usage.
E.g., use registers 8 through 15 for “temporary” values (short-term), refer to as \$t0 through \$t7.

Slide 12

High-Level Languages Versus Assembly Language

- In a high-level language you work with “variables” — conceptually, names for memory locations. Can do arithmetic on them, copy them, etc.
- In machine/assembly language, what you can do may be more restricted — e.g., in MIPS architecture, must load data into a register before doing arithmetic.
- Compiler's job is to translate from the somewhat abstract HLL view to machine language. To do this, normally associate variables with registers — load data from memory into registers, calculate, store it back. A “good” compiler tries to minimize loads/stores.

Slide 13

Example

- Suppose we have this in C (and assuming all variables are 32-bit integers):

$$f = (g + h) - (i + j)$$

- What instructions should compiler produce? Assume we're using `$s0` for `f`, `$s1` for `g`, `$s2` for `h`, `$s3` for `i`, `$s4` for `j`.

(Symbolic register names starting `$s` are used for slightly longer-term storage than the ones starting `$t`.)

(Where do values come from? Next topic ...)

Slide 14

Memory, Revisited

- Usually we think of memory as big 1D array of 8-bit "bytes", each with address (index into array) and contents (value of array element).
- Often we operate on elements in larger units. For MIPS, natural unit is 32-bit "word". (Other architectures also often operate on words. 32 bits was common until recently; 64 bits probably more so now.)
- MIPS is a "load/store" architecture, meaning access to memory is limited to copying data between memory and registers. Alternatives include arithmetic instructions to operate on memory directly.

Memory-Access Instructions — Load

- Goal is to get one 32-bit word from memory and put in a register.
- How to specify location in memory? Seems most useful to have address in a register. For a little more flexibility, specify address in terms of “base” and “displacement”.

Slide 15

```
lw    r, d(b)
```

Address specified by contents of register b plus (constant) d . Loads word into register r .

- `sw` (“store word”) instruction is similar.

Example

- Suppose we have this in C (and assuming g and h are 32-bit integers and a is an array of same):

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

- What instructions should compiler produce? Assume we’re using $\$s3$ for starting (“base”) address of a , $\$s2$ for h , $\$s1$ for g .

Slide 16

Addition Using Constant

- “Add immediate”

```
addi r1, r2, c
```

adds constant c (16-bit signed integer, can be negative) to contents of $r2$, puts result in $r1$.

Slide 17

- Exists because often we need to use a small constant in a program.
Basic principle: “Make the common case fast.”

Representing (Integer) Data in Binary

- Remember that to the hardware “it’s all ones and zero” — any data you’re working with.
- As an example — representation of signed integers using two’s complement notation. Should have been covered in CSCI 1320, but read/skim 2.4 if you don’t remember.

Slide 18

A Little About the Simulator

Slide 19

- As mentioned, installed on our machines is a simulator you can use to try your programs. It simulates a MIPS processor running a *very* primitive operating system (just enough to load programs and do some simple console I/O). It assembles programs on the fly.
- Your code goes in a file with extension `.s`. (Sample starter code on “Sample programs” page. Contains many things we haven’t talked about yet but could still be useful for trying things out.)
- Start it with command `xspim` (`spim` for command-line version).
(Short demo.)

Minute Essay

Slide 20

- Write MIPS assembly code for the following C program fragment:

```
a = b + c + d + e
```

Assume we have `b, c, d, e` in `$s1` through `$s4` and want to have `a` in `$s0`.

Can you think of more than one way to do it? If you can, does one seem better than the other, and why?

Minute Essay Answer

- One way:

```
add    $s0, $s1, $s2
add    $s0, $s0, $s3
add    $s0, $s0, $s4
```

Slide 21

Another way (not as good since uses more registers?):

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