Behind the Curtain

1. Computer organization
2. Computer Instructions
3. Memory concepts
4. Where should code go?
5. Computers as systems

Introduce TAs
Friday's lab is posted.
CS logins.
Do Prelab ahead before class
1st Problem set next Wed

Computers Everywhere

- The computers we are used to
  - Desktops
  - Laptops
  - Embedded processors
    - Cars
    - Mobile phones
    - Toasters, irons, wristwatches, happy-meal toys

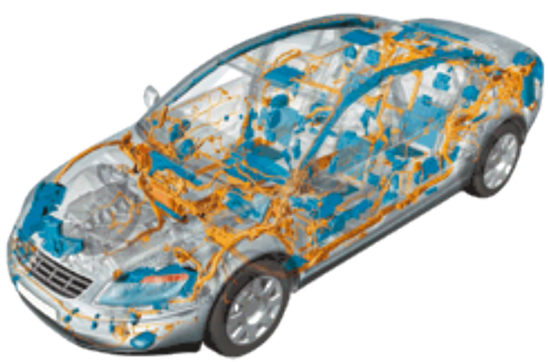

Computer Organization

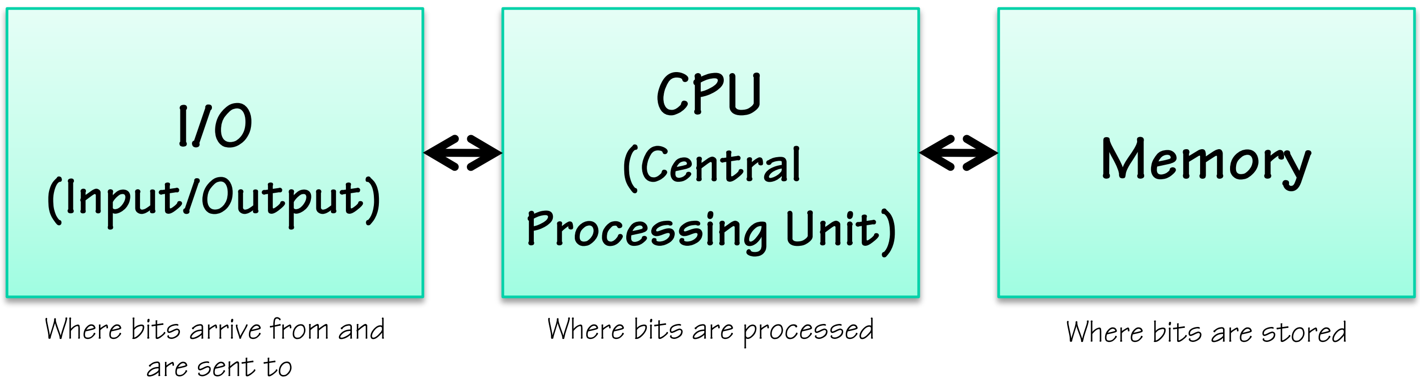

- Every computer has at least three basic units
  - Input/Output
    - where data arrives from the outside world
    - where data is sent to the outside world
    - where data is archived for the long term (i.e. when the lights go out)
  - Memory
    - where data is stored (numbers, text, lists, arrays, data structures)
  - Central Processing Unit
    - where data is manipulated, analyzed, etc.

Computer Organization (cont)

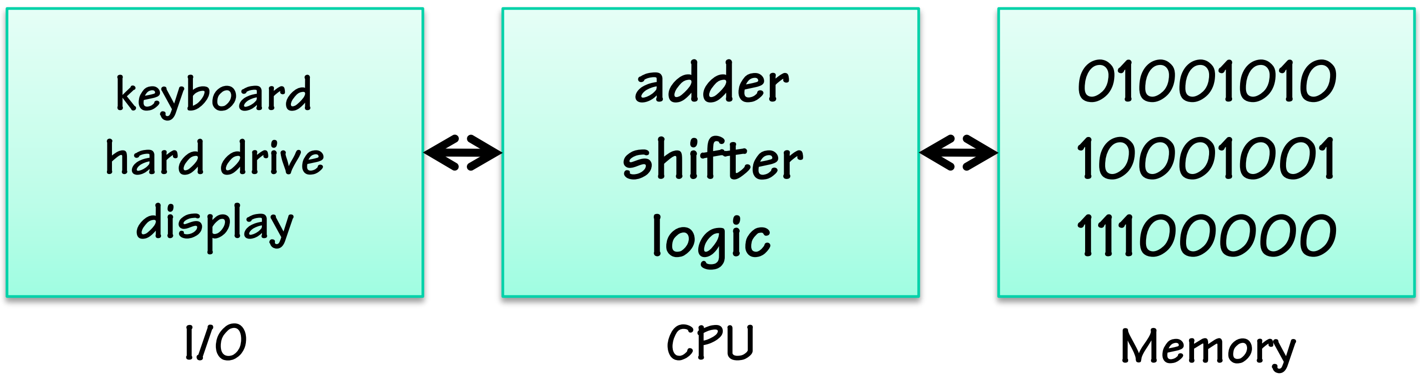

- Properties of units
  - Input/Output
    - must convert symbols to bits and vice versa
    - where the analog "real world" meets the digital "computer world"
    - must somehow synchronize to the CPU's clock
  - Memory
    - stores bits in "addressable" units, such as bytes or words
    - every memory unit has an "address" and "contents", like a mailbox
  - Central Processing Unit
    - besides processing, it also coordinates data's movements between units

What Sorts of Processing?

A CPU performs low-level operations called INSTRUCTIONS

Arithmetic

- ADD X to Y then put the result in Z
- SUBTRACT X from Y then put the result back in Y

Logical

- Set Z to 1 if X AND Y are 1, otherwise set Z to 0
  (AND X with Y then put the result in Z)
- Set Z to 1 if X OR Y are 1, otherwise set Z to 0
  (OR X with Y then put the result in Z)

Comparison

- Set Z to 1 if X is EQUAL to Y, otherwise set Z to 0
- Set Z to 1 if X is GREATER THAN OR EQUAL to Y, otherwise set Z to 0

Control

- Skip the next INSTRUCTION if Z is EQUAL to 0

Anatomy of an Instruction

Nearly all instructions can be made to fit a common template

OPCODE DESTINATION, OPERAND_1, OPERAND_2

Issues remaining …

- Which operations to include?
- Where to get variables and constants?
- Where to store the results?

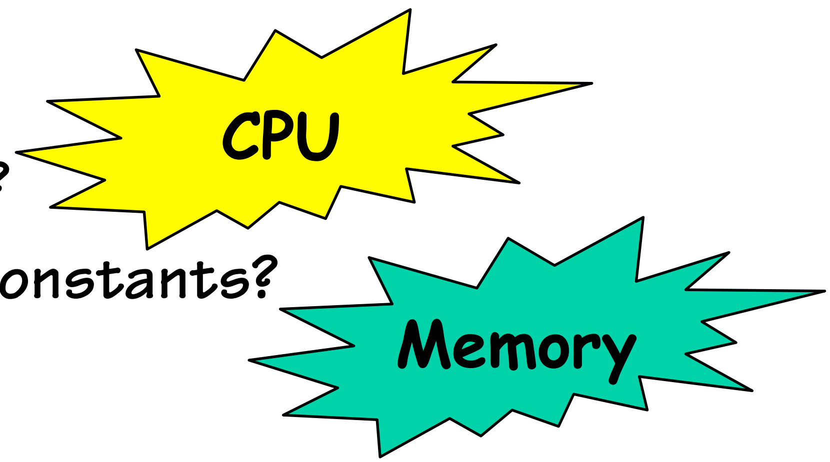

Memory Concepts

- Memory is divided into "addressable" blocks, each with an address (like an array with indices)
- Addressable blocks are usually larger than a bit, typically 8, 16, 32, or 64 bits
- Each address has variable "contents"
- Contents might be:
  - Integers in 2's complement
  - Floats in IEEE format
  - Strings in ASCII or Unicode
  - Data structure de jour
  - ADDRESSES
  - Nothing distinguishes the difference

Address	Contents
0	42
1	3.141592
2	"Lee "
3	"Hart"
4	"Bud "
5	"Levi"
6	"le "
7	2
8	0c3c1d7fff
9	0x37bdfffc
10	0x24040090
11	0x0c00000e
12	0x1000ffff
13	-100
14	0x00004020
15	0x20090001

One More Thing…

- INSTRUCTIONS for the CPU are stored in memory along with data
- CPU fetches instructions, decodes them and then performs their implied operation
- Mechanism inside the CPU directs which instruction to get next.
- They appear in memory as a string of bits that are typically uniform in size
- Their encoding as "bits" is called "machine language." ex: 0c3c1d7fff
- We assign "mnemonics" to particular bit patterns to indicate meanings. These mnemonics are called assembly language. ex: lui $sp, 0x7fff

Address	Contents
0	42
1	3.141592
2	"Lee "
3	"Hart"
4	"Bud "
5	"Levi"
6	"le "
7	2
8	lui $sp,0x7fff
9	ori $sp,$sp,0x7fff
10	addiu $a0,$0,144
11	jal 0x0000000e
12	beq $0,$0,0x0c
13	-100
14	add $t0,$0,$0
15	addi $t1,$0,1

A Bit of History

- There is a commonly reoccurring debate over whether "data" and "instructions" should be mixed. Leads to two common flavors of computer architectures

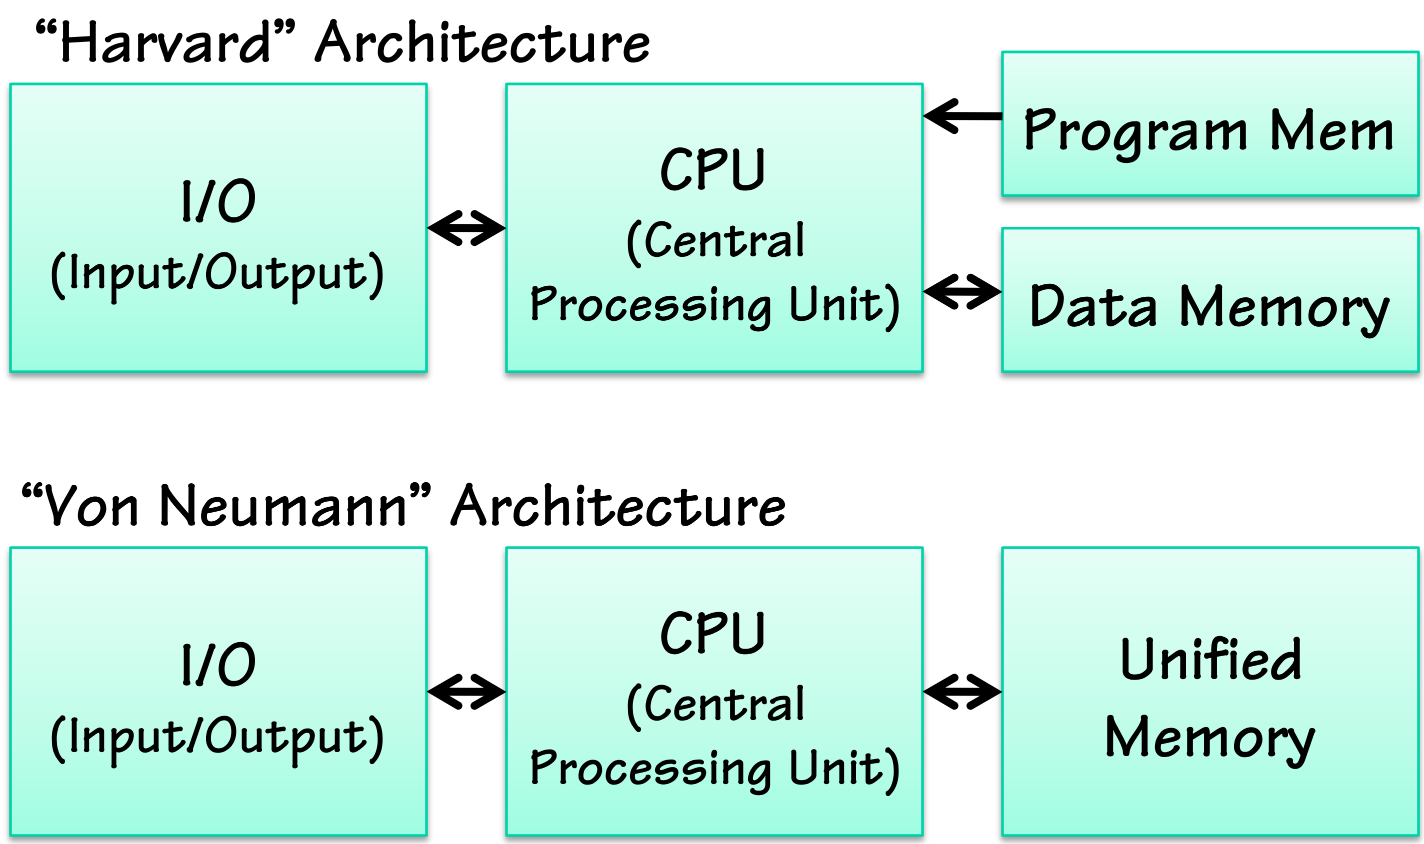

A Bit of History

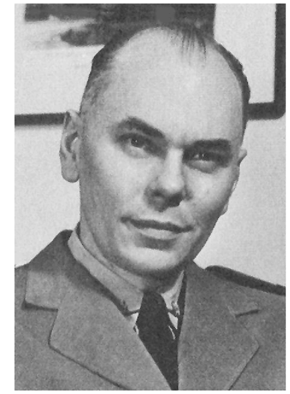

Howard Aiken:
Architect of the
Harvard Mark 1

- Harvard Architecture
  - Instructions and data do not interact, they can have different "word sizes" and exist in different "address spaces"
  - Advantages:
    - No self-modifying code (a common hacker trick)
    - Optimize word-lengths of instructions for control and data for applications
    - Higher Throughput (i.e. you can fetch data and instructions from their memories simultaneously)
  - Disadvantages:
    - The H/W designer decides the trade-off between how big of a program and how large are data
    - Hard to write "Native" programs that generate new programs (i.e. assemblers, compliers, etc.)
    - Hard to write "Operating Systems" which are programs that at various points treat other programs as data (i.e. loading them from disk into memory, swapping out processes that are idle)

A Bit of History

John Von Neumann: Proponent of unified memory architecture

- Von Neumann Architecture
  - Instructions and data are indistinguishable bits in a common memory that share a common "word size" and "address space"
  - Most common model used today, and what we assume in 411
  - Advantages:
    - S/W designer decides how to allocate memory between data and programs
    - Can write "Native" programs to create new programs (assemblers and compliers)
    - Programs and subroutines can be loaded, relocated, and modified by other programs (dangerous, but powerful)
  - Disadvantages:
    - Word size must suit both common data types and instructions
    - Slightly lower performance due to memory bottleneck (mediated in modern computers by the use of separate program and data caches)
    - We need to be very careful when treading on memory. Folks have taken advantage of the program-data unification to introduce viruses.

Computer Systems

- What is a computer system?
- Where does it start?
- Where does it end?

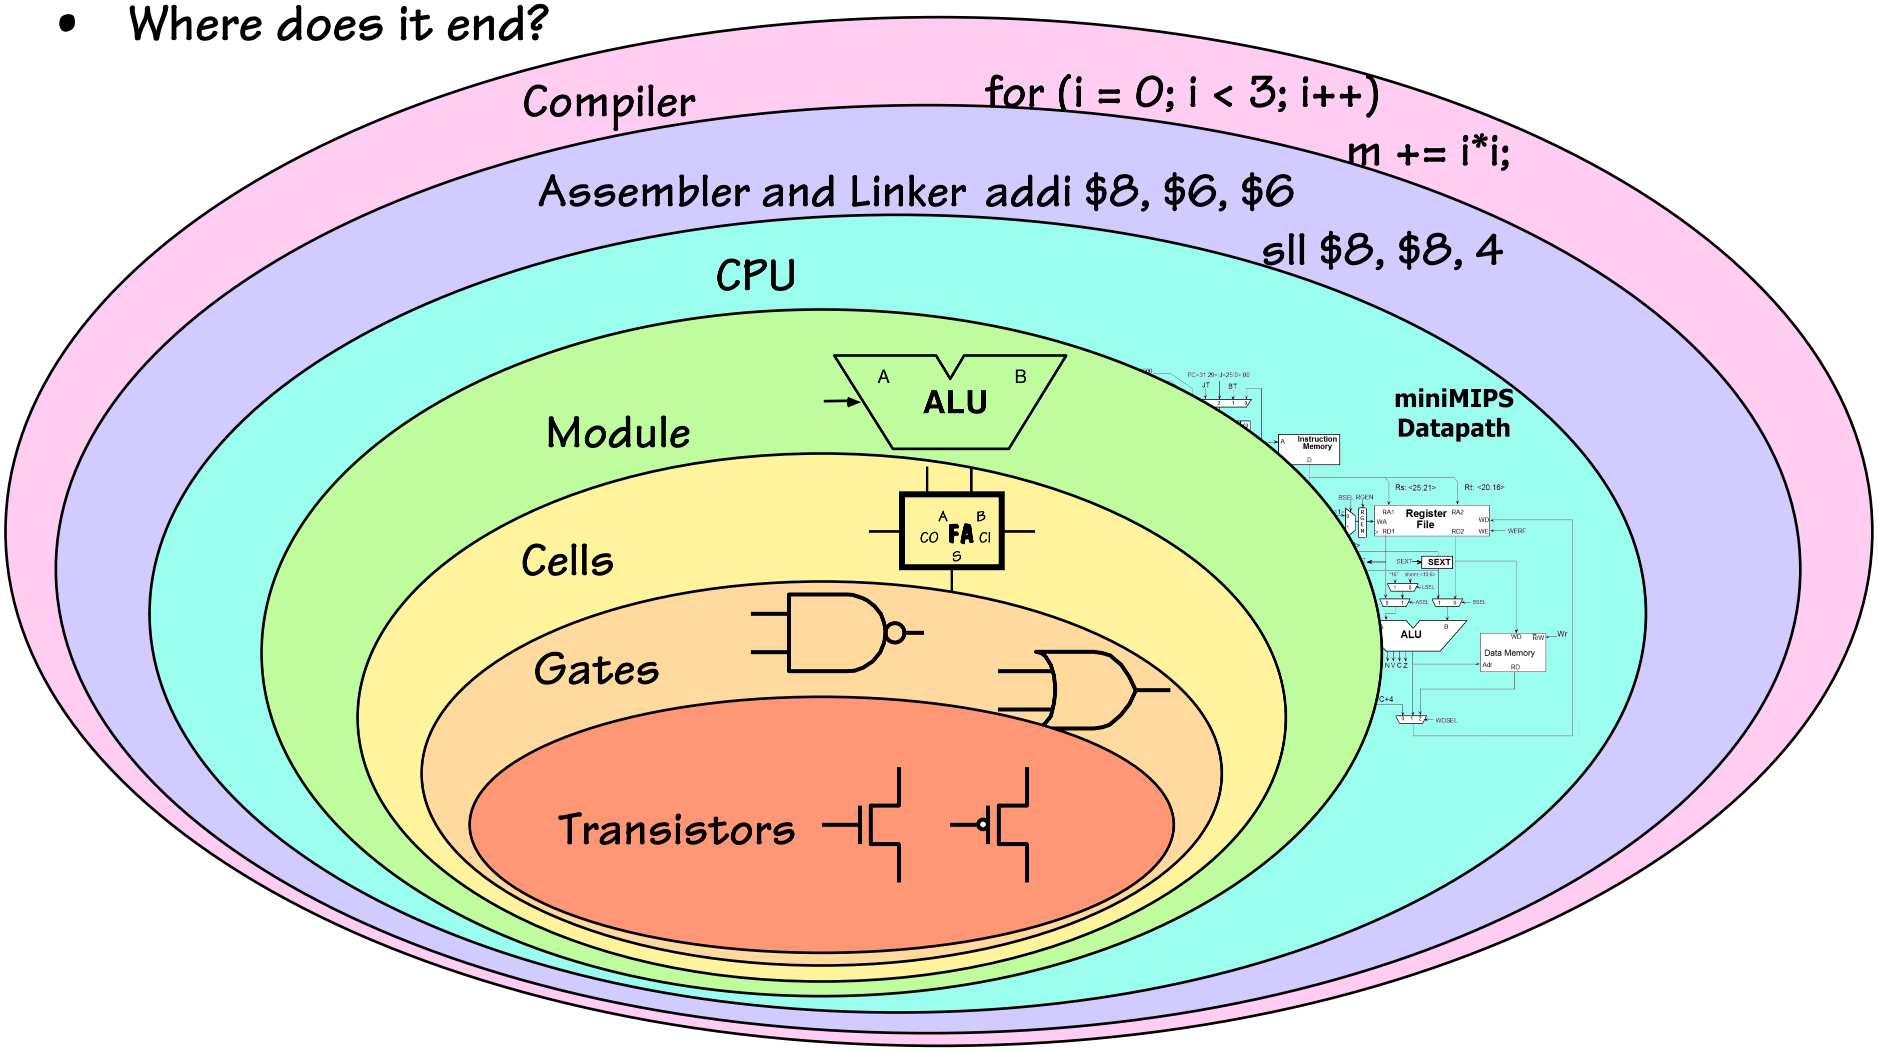

Computers as Translators

Much of what computers do is run programs that interpret a "High-level" problem specification and converts it to a "lower-level" problem that is closer the simple instructions that it understands

- High-Level Languages
  - Compilers
  - Interpreters
- Assembly Language

```
int x, y;
y = (x-3)*(y+123456)
```

→

```
x:      .word 0
y:      .word 0
c:      .word 123456
    ...
    lw      $t0, x
    addi    $t0, $t0, -3
    lw      $t1, y
    lw      $t2, c
    add     $t1, $t1, $t2
    mul     $t0, $t0, $t1
    sw      $t0, y
```

Computers as Translators

Much of what computers do is run programs that interpret a "High-level" problem specification and converts it to a "lower-level" problem that is closer the simple instructions that it understands

- Assembly Language
- Machine Language

```
x:      .word 0
y:      .word 0
c:      .word 123456
    ...
    lw        $t0, x
    addi      $t0, $t0, -3
    lw        $t1, y
    lw        $t2, c
    add       $t1, $t1, $t2
    mul       $t0, $t0, $t1
    sw        $t0, y
```

→

```
0x04030201
0x08070605
0x00000001
0x00000002
0x00000003
0x00000004
0x706d6f43
```

Why So Many Languages?

- Application Specific
  - Pre-historically: COBOL vs. Fortran
  - Middle ages: C++ vs. Objective C
  - Recent Past: C# vs. Java
  - Today: Python vs. Matlab
- Code Maintainability
  - High-level specifications are easier to understand and modify
- Code Reuse
- Code Portability
- Virtual Machines

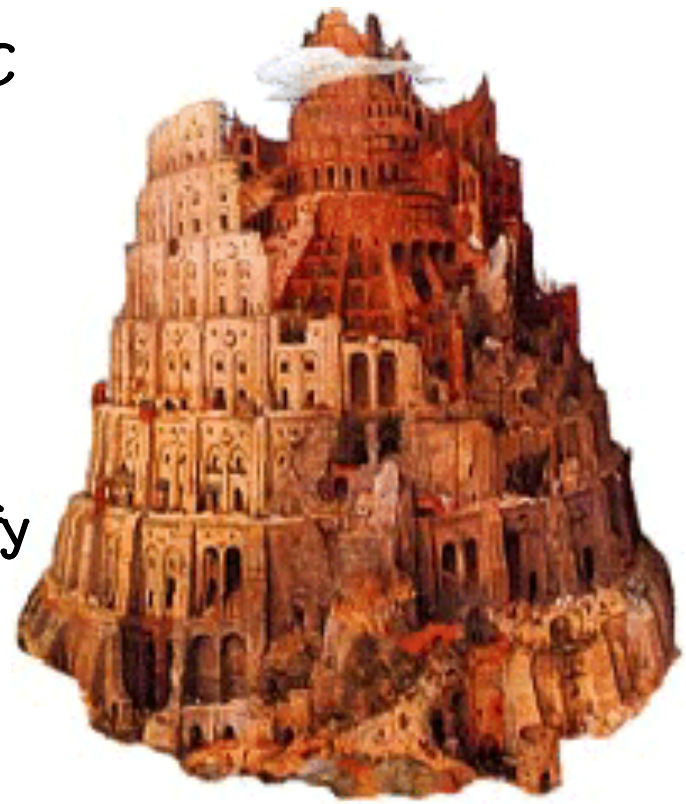

Next Time

- A complete Instruction Set
- Assembly Language
- Machine Language