

## CS61C : Machine Structures

### Lecture 3.2.2

#### Floating Point II & Linking

2004-07-06

Kurt Meinz

inst.eecs.berkeley.edu/~cs61c



CS 61C L3.2.2 Floating Point 2 (1)

K. Meinz, Summer 2004 © UCB

### FP Review

- Floating Point numbers approximate values that we want to use.
- IEEE 754 Floating Point Standard is most widely accepted attempt to standardize interpretation of such numbers
  - Every desktop or server computer sold since ~1997 follows these conventions
- Summary (single precision):  

31 30	23 22	0
S	Exponent	Significand
1 bit	8 bits	23 bits

 $\bullet (-1)^S \times (1 + \text{Significand}) \times 2^{(\text{Exponent}-127)}$



CS 61C L3.2.2 Floating Point 2 (2)

K. Meinz, Summer 2004 © UCB

### Representation for Denorms (1/3)

- Problem: There's a gap among representable FP numbers around 0
  - Smallest representable pos num:  
 $a = 1.0 \dots 2 * 2^{-126} = 2^{-126}$
  - Second smallest representable pos num:  
 $b = 1.000\dots 1_2 * 2^{-126} = 2^{-126} + 2^{-149}$   
 $a - 0 = 2^{-126}$   
 $b - a = 2^{-149}$   
Gaps!  
  
Normalization and implicit 1 is to blame!



CS 61C L3.2.2 Floating Point 2 (3)

K. Meinz, Summer 2004 © UCB

### Representation for Denorms (2/3)

- Solution:
  - We still haven't used Exponent = 0, Significand nonzero
  - Denormalized number: no leading 1, **implicit exponent = -126**.
  - Smallest representable pos num:  
 $a = 2^{-149}$
  - Second smallest representable pos num:  
 $b = 2^{-148}$   




CS 61C L3.2.2 Floating Point 2 (4)

K. Meinz, Summer 2004 © UCB

### Representation for Denorms (3/3)

- Normal FP equation:  
 $\bullet (-1)^S \times (1 + \text{Significand}) \times 2^{(\text{Exponent}-127)}$
- If (fp.exp == 0 and fp.significant != 0)
  - Denorm
  - $(-1)^S \times (0 + \text{Significand}) \times 2^{(-126)}$



CS 61C L3.2.2 Floating Point 2 (5)

K. Meinz, Summer 2004 © UCB

### IEEE Four Rounding Modes

- Math on real numbers  $\Rightarrow$  we worry about rounding to fit result in the significant field.
- FP hardware carries 2 extra bits of precision, and rounds for proper value
- Rounding occurs when converting...
  - double to single precision
  - floating point # to an integer



CS 61C L3.2.2 Floating Point 2 (6)

K. Meinz, Summer 2004 © UCB

## IEEE Four Rounding Modes

- Round towards  $+\infty$ 
  - ALWAYS round “up”:  $2.1 \Rightarrow 3, -2.1 \Rightarrow -2$
- Round towards  $-\infty$ 
  - ALWAYS round “down”:  $1.9 \Rightarrow 1, -1.9 \Rightarrow -2$
- Truncate
  - Just drop the last bits (round towards 0)
- Round to (nearest) even (default)
  - Normal rounding, almost:  $2.5 \Rightarrow 2, 3.5 \Rightarrow 4$
  - Like you learned in grade school
  - Insures fairness on calculation
  - Half the time we round up, other half down



CS 61C L3.2.2 Floating Point 2 (7)

K. Meinz, Summer 2004 © UCB

## Integer Multiplication (1/3)

- Paper and pencil example (unsigned):

$$\begin{array}{r} \text{Multiplicand} & 1000 & 8 \\ \text{Multiplier} & \times 1001 & 9 \\ & \hline & 1000 \\ & 0000 & \\ & 0000 & \\ & +1000 & \\ \hline & 01001000 & \end{array}$$

•  $m$  bits  $\times n$  bits =  $m + n$  bit product



CS 61C L3.2.2 Floating Point 2 (8)

K. Meinz, Summer 2004 © UCB

## Integer Multiplication (2/3)

- In MIPS, we multiply registers, so:
  - 32-bit value  $\times$  32-bit value = 64-bit value
- Syntax of Multiplication (signed):
  - `mult register1, register2`
  - Multiplies 32-bit values in those registers & puts 64-bit product in special result regs:
    - puts product upper half in **hi**, lower half in **lo**
  - **hi** and **lo** are 2 registers separate from the 32 general purpose registers
  - Use **mfhi register** & **mflo register** to move from **hi**, **lo** to another register



CS 61C L3.2.2 Floating Point 2 (9)

K. Meinz, Summer 2004 © UCB

## Integer Multiplication (3/3)

- Example:

• in C: `a = b * c;`

• in MIPS:

```
- let b be $s2; let c be $s3; and let a be $s0  
and $s1 (since it may be up to 64 bits)  
mult $s2,$s3 # b*c  
mfhi $s0 # upper half of  
# product into $s0  
mflo $s1 # lower half of  
# product into $s1
```

• Note: Often, we only care about the lower half of the product.



CS 61C L3.2.2 Floating Point 2 (10)

K. Meinz, Summer 2004 © UCB

## Integer Division (1/2)

- Paper and pencil example (unsigned):

$$\begin{array}{r} \text{Divisor} & 1001 & \text{Quotient} \\ 1000 | 1001010 & \text{Dividend} \\ -1000 \\ \hline & 10 \\ & 101 \\ & 1010 \\ & -1000 \\ \hline & 10 & \text{Remainder} \\ & & (\text{or Modulo result}) \end{array}$$

• Dividend = Quotient  $\times$  Divisor + Remainder



CS 61C L3.2.2 Floating Point 2 (11)

K. Meinz, Summer 2004 © UCB

## Integer Division (2/2)

- Syntax of Division (signed):

• `div register1, register2`

• Divides 32-bit register 1 by 32-bit register 2:

• puts remainder of division in **hi**, quotient in **lo**

- Implements C division (/) and modulo (%)

• Example in C: `a = c / d;`  
`b = c % d;`

• in MIPS: `a↔$s0;b↔$s1;c↔$s2;d↔$s3`

```
div $s2,$s3 # lo=c/d, hi=c%d  
mflo $s0 # get quotient  
mfhi $s1 # get remainder
```



CS 61C L3.2.2 Floating Point 2 (12)

K. Meinz, Summer 2004 © UCB

### Unsigned Instructions & Overflow

- MIPS also has versions of mult, div for **unsigned operands**:

multu  
divu

- Determines whether or not the product and quotient are changed if the operands are signed or unsigned.

- MIPS does not check overflow on ANY signed/unsigned multiply, divide instr

- Up to the software to check hi



CS 61C L3.2.2 Floating Point 2 (13)

K. Meinz, Summer 2004 © UCB

### FP Addition & Subtraction

- Much more difficult than with integers (can't just add significands)
- How do we do it?
  - De-normalize to match larger exponent
  - Add significands to get resulting one
  - Normalize (& check for under/overflow)
  - Round if needed (may need to renormalize)
- If signs ≠, do a subtract. (Subtract similar)
  - If signs ≠ for add (or = for sub), what's ans sign?
- Question: How do we integrate this into the integer arithmetic unit? [Answer: We don't!]



CS 61C L3.2.2 Floating Point 2 (14)

K. Meinz, Summer 2004 © UCB

### MIPS Floating Point Architecture (1/4)

- Separate floating point instructions:
  - Single Precision:  
add.s, sub.s, mul.s, div.s
  - Double Precision:  
add.d, sub.d, mul.d, div.d
- These are far more complicated than their integer counterparts
  - Can take much longer to execute



CS 61C L3.2.2 Floating Point 2 (15)

K. Meinz, Summer 2004 © UCB

### MIPS Floating Point Architecture (2/4)

#### • Problems:

- Inefficient to have different instructions take vastly differing amounts of time.
- Generally, a particular piece of data will not change FP ⇔ int within a program.
  - Only 1 type of instruction will be used on it.
- Some programs do no FP calculations
- It takes lots of hardware relative to integers to do FP fast



CS 61C L3.2.2 Floating Point 2 (16)

K. Meinz, Summer 2004 © UCB

### MIPS Floating Point Architecture (3/4)

- 1990 Solution: Make a completely separate chip that handles only FP.
- Coprocessor 1: FP chip
  - contains 32 32-bit registers: \$f0, \$f1, ...
  - most of the registers specified in .s and .d instruction refer to this set
  - separate load and store: lwc1 and swc1 ("load word coprocessor 1", "store ...")
  - Double Precision: by convention, even/odd pair contain one DP FP number:  
\$f0/\$f1, \$f2/\$f3, ... , \$f30/\$f31
    - Even register is the name



CS 61C L3.2.2 Floating Point 2 (17)

K. Meinz, Summer 2004 © UCB

### MIPS Floating Point Architecture (4/4)

- 1990 Computer actually contains multiple separate chips:
  - Processor: handles all the normal stuff
  - Coprocessor 1: handles FP and only FP;
  - more coprocessors?... Yes, later
  - Today, FP coprocessor integrated with CPU, or cheap chips may leave out FP HW
- Instructions to move data between main processor and coprocessors:
  - mfc0, mtc0, mfc1, mtc1, etc.
- Appendix pages A-70 to A-74 contain many, many more FP operations.



CS 61C L3.2.2 Floating Point 2 (18)

K. Meinz, Summer 2004 © UCB

## Questions

1. Converting float  $\rightarrow$  int  $\rightarrow$  float  
produces same float number  
**No!**       $3.14 \rightarrow 3 \rightarrow 3$
2. Converting int  $\rightarrow$  float  $\rightarrow$  int  
produces same int number  
**No!**       $(2^{30} + 2^1)$
3. FP add is associative:  
 $(x+y)+z = x+(y+z)$   
**No!**       $2^{30} + -2^{30} + 1$



CS 61C L3.2.2 Floating Point 2 (19)

K. Meinz, Summer 2004 © UCB

## FP/Math Summary

- Reserve exponents, significands:

Exponent	Significand	Object
0	0	0
0	nonzero	Denorm
1-254	anything	+/- fl. pt. #
255	0	+/- infinity
255	nonzero	NaN

- Integer mult, div uses hi, lo regs  
• mfhi and mflo copies out.
- Four rounding modes (to even default)
- MIPS FL ops complicated, expensive



CS 61C L3.2.2 Floating Point 2 (20)

K. Meinz, Summer 2004 © UCB

## CLL Overview

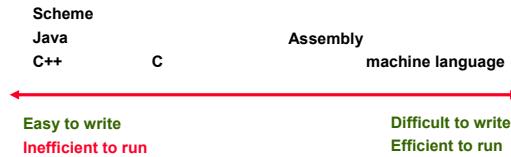
- Interpretation vs Translation
- Translating C Programs
  - Compiler
  - Assembler
  - Linker
  - Loader
- An Example



CS 61C L3.2.2 Floating Point 2 (21)

K. Meinz, Summer 2004 © UCB

## Language Continuum



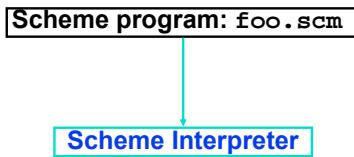
- Interpret a high level language if efficiency is not critical
- Translate (compile) to a lower level language to improve performance



CS 61C L3.2.2 Floating Point 2 (22)

K. Meinz, Summer 2004 © UCB

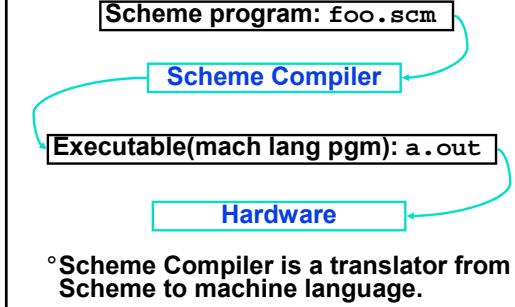
## Interpretation



CS 61C L3.2.2 Floating Point 2 (23)

K. Meinz, Summer 2004 © UCB

## Translation



CS 61C L3.2.2 Floating Point 2 (24)

K. Meinz, Summer 2004 © UCB

## Interpretation

- Any good reason to interpret machine language in software?
- SPIM – useful for learning / debugging
- Apple Macintosh conversion
  - Switched from Motorola 680x0 instruction architecture to PowerPC.
  - Could require all programs to be re-translated from high level language
  - Instead, let executables contain old and/or new machine code, interpret old code in software if necessary



CS 61C L3.2.2 Floating Point 2 (28)

K. Meinz, Summer 2004 © UCB

## Interpretation vs. Translation?

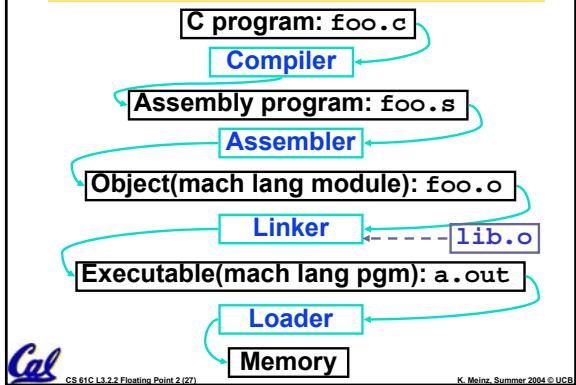
- Easier to write interpreter
- Interpreter closer to high-level, so gives better error messages (e.g., SPIM)
  - Translator reaction: add extra information to help debugging (line numbers, names)
- Interpreter slower (10x?) but code is smaller (1.5X to 2X?)



CS 61C L3.2.2 Floating Point 2 (28)

K. Meinz, Summer 2004 © UCB

## Steps to Starting a Program



CS 61C L3.2.2 Floating Point 2 (27)

K. Meinz, Summer 2004 © UCB

## Compiler

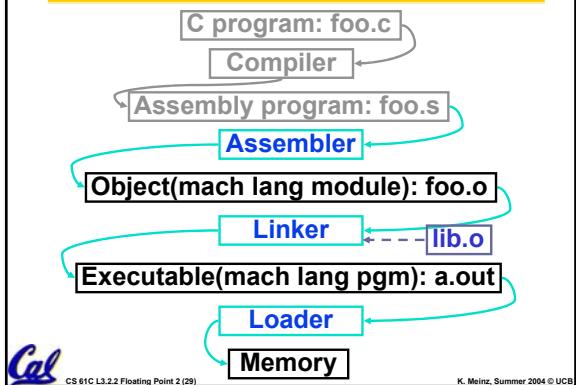
- Input: High-Level Language Code (e.g., C, Java such as foo.c)
- Output: MAL Assembly Language Code (e.g., foo.s for MIPS)
- Note: Output **may** contain pseudoinstructions
  - (btw: hardest stage by far!)



CS 61C L3.2.2 Floating Point 2 (28)

K. Meinz, Summer 2004 © UCB

## Where Are We Now?



CS 61C L3.2.2 Floating Point 2 (29)

K. Meinz, Summer 2004 © UCB

## Assembler

- Input: MAL Assembly Language Code (e.g., foo.s for MIPS)
- Output: Object Code, information tables (e.g., foo.o for MIPS)
- Reads and Uses **Directives**
- Replace Pseudoinstructions
- Produce Machine Language
- Creates **Object File**



CS 61C L3.2.2 Floating Point 2 (30)

K. Meinz, Summer 2004 © UCB

### Assembler Directives (p. A-51 to A-53)

- Give directions to assembler, but do not produce machine instructions
  - .text: Subsequent items put in user text segment
  - .data: Subsequent items put in user data segment
  - .globl sym: declares sym global and can be referenced from other files
  - .asciiz str: Store the string str in memory and null-terminate it
  - .word w1...wn: Store the n 32-bit quantities in successive memory words



CS 61C L3.2.2 Floating Point 2 (31)

K. Meinz, Summer 2004 © UCB

### Pseudoinstruction Replacement

- Asm. treats convenient variations of machine language instructions as if real instructions

Pseudo:

```
subu $sp,$sp,32      addiu $sp,$sp,-32
sd $a0, 32($sp)    sw $a0, 32($sp)
mul $t7,$t6,$t5     mult $t6,$t5
mflo $t7             mflo $t7
addu $t0,$t6,1       addiu $t0,$t6,1
ble $t0,100,loop    slti $at,$t0,101
                    bne $at,$0,loop
la $a0, str          lui $at, left(str)
                     ori $a0,$at,right(str)
```



CS 61C L3.2.2 Floating Point 2 (32)

K. Meinz, Summer 2004 © UCB

### Producing Machine Language (1/3)

- Constraint on Assembler:
  - The object file output (foo.o) may be only one of many object files in the final executable:
    - C: #include "my\_helpers.h"
    - C: #include <stdio.h>
- Consequences:
  - Object files won't know their base addresses until they are linked/loaded!
  - References to addresses will have to be adjusted in later stages



CS 61C L3.2.2 Floating Point 2 (33)

K. Meinz, Summer 2004 © UCB

### Producing Machine Language (2/3)

- Simple Case
  - Arithmetic, Logical, Shifts, and so on.
  - All necessary info is within the instruction already.
- What about Branches?
  - PC-Relative and in-file
  - In TAL, we know by how many instructions to branch.
- So these can be handled easily.



CS 61C L3.2.2 Floating Point 2 (34)

K. Meinz, Summer 2004 © UCB

### Producing Machine Language (3/3)

- What about jumps (j and jal)?
  - Jumps require **absolute address**.
- What about references to data?
  - la gets broken up into lui and ori
  - These will require the full 32-bit address of the data.
- These can't be determined yet, so we create two tables for use by linker/loader...



CS 61C L3.2.2 Floating Point 2 (35)

K. Meinz, Summer 2004 © UCB

### 1: Symbol Table

- List of "items" **provided** by this file.
- What are they?
  - Labels: function calling
  - Data: anything in the .data section; variables which may be accessed across files
- Includes base address of label in the file.



CS 61C L3.2.2 Floating Point 2 (36)

K. Meinz, Summer 2004 © UCB

## 2: Relocation Table

- List of “items” **needed** by this file.
  - Any label jumped to: `j` or `jal`
    - internal
    - external (including lib files)
  - Any named piece of data
    - Anything referenced by the `la` instruction
    - static variables
  - Contains base address of instruction w/dependency, dependency name



CS 61C L3.2.2 Floating Point 2 (37)

K. Meinz, Summer 2004 © UCB

## Question

- Which lines go in the symbol table and/or relocation table?

```
my_func:
    lui $a0 my_arrayh      # a (from la)
    ori $a0 $a0 my_arrayl  # b (from la)
    jal add_link            # c
    bne $a0,$v0, my_func   # d
```

- A: Symbol: `my_func` relocate: `my_array`  
 B: - relocate: `my_array`  
 C: - relocate: `add_link`  
 D: -



CS 61C L3.2.2 Floating Point 2 (38)

K. Meinz, Summer 2004 © UCB

## Object File Format

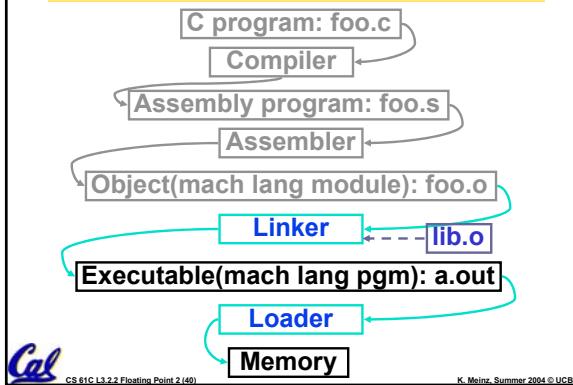
- **object file header:** size and position of the other pieces of the object file
- **text segment:** the machine code
- **data segment:** binary representation of the data in the source file
- **relocation information:** identifies lines of code that need to be “handled”
- **symbol table:** list of this file’s labels and data that can be referenced
- **debugging information**



CS 61C L3.2.2 Floating Point 2 (39)

K. Meinz, Summer 2004 © UCB

## Where Are We Now?



## Link Editor/Linker (1/3)

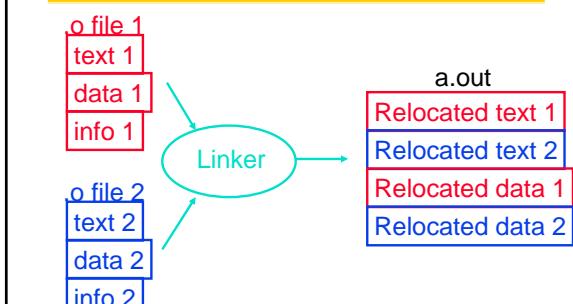
- Input: Object Code, information tables (e.g., `foo.o` for MIPS)
- Output: Executable Code (e.g., `a.out` for MIPS)
- Combines several object (.o) files into a single executable (“**linking**”)
- Enable Separate Compilation of files
  - Changes to one file do not require recompilation of whole program
    - Windows NT source is >40 M lines of code!
  - Link Editor name from editing the “links” in jump and link instructions



CS 61C L3.2.2 Floating Point 2 (41)

K. Meinz, Summer 2004 © UCB

## Link Editor/Linker (2/3)



CS 61C L3.2.2 Floating Point 2 (42)

K. Meinz, Summer 2004 © UCB

### Link Editor/Linker (3/3)

- Step 1: Take text segment from each .o file and put them together.
- Step 2: Take data segment from each .o file, put them together, and concatenate this onto end of text segments.
- Step 3: Resolve References
  - Go through Relocation Table and handle each entry
  - That is, fill in all **absolute addresses**



CS 61C L3.2.2 Floating Point 2 (43)

K. Meinz, Summer 2004 © UCB

### Resolving References (1/2)

- Linker **assumes** first word of first text segment is at address 0x00000000.
- Linker knows:
  - length of each text and data segment
  - ordering of text and data segments
- Linker calculates:
  - absolute address of each label to be jumped to (internal or external) and each piece of data being referenced



CS 61C L3.2.2 Floating Point 2 (44)

K. Meinz, Summer 2004 © UCB

### Resolving References (2/2)

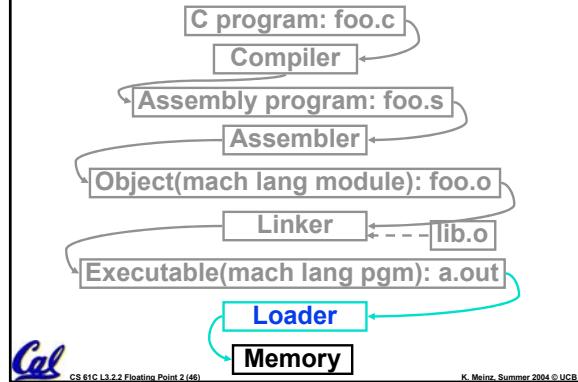
- To resolve references:
  - search for reference (data or label) in all symbol tables
  - if not found, search library files (for example, for `printf`)
  - once absolute address is determined, fill in the machine code appropriately
- Output of linker: executable file containing text and data (plus header)



CS 61C L3.2.2 Floating Point 2 (45)

K. Meinz, Summer 2004 © UCB

### Where Are We Now?



### Loader (1/3)

- Input: Executable Code (e.g., a.out for MIPS)
- Output: (program is run)
- Executable files are stored on disk.
- When one is run, loader's job is to load it into memory and start it running.
- In reality, loader is the operating system (OS)
  - loading is one of the OS tasks



CS 61C L3.2.2 Floating Point 2 (47)

K. Meinz, Summer 2004 © UCB

### Loader (2/3)

- So what does a loader do?
- Reads executable file's header to determine size of text and data segments
- Creates new address space for program large enough to hold text and data segments, along with a stack segment
- Copies instructions and data from executable file into the new address space (this may be anywhere in memory)



CS 61C L3.2.2 Floating Point 2 (48)

K. Meinz, Summer 2004 © UCB

### Loader (3/3)

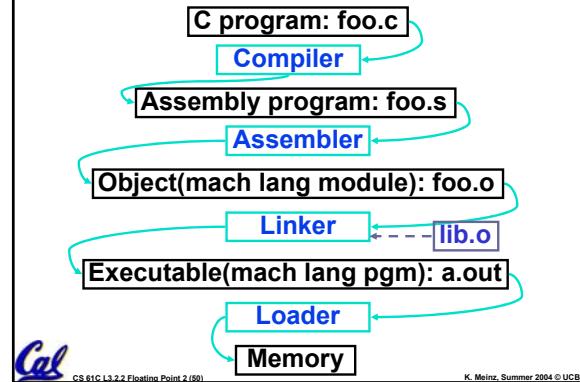
- Copies arguments passed to the program onto the stack
- Initializes machine registers
  - Most registers cleared, but stack pointer assigned address of 1st free stack location
- Jumps to start-up routine that copies program's arguments from stack to registers and sets the PC
  - If main routine returns, start-up routine terminates program with the exit system call



CS 61C L3.2.2 Floating Point 2 (49)

K. Meinz, Summer 2004 © UCB

### Things to Remember (1/3)



CS 61C L3.2.2 Floating Point 2 (60)

K. Meinz, Summer 2004 © UCB

### Things to Remember (2/3)

- Compiler converts a single HLL file into a single assembly language file.
- Assembler removes pseudoinstructions, converts what it can to machine language, and creates a checklist for the linker (relocation table). This changes each .s file into a .o file.
- Linker combines several .o files and resolves absolute addresses.
- Loader loads executable into memory and begins execution.



CS 61C L3.2.2 Floating Point 2 (61)

K. Meinz, Summer 2004 © UCB

### Things to Remember 3/3

- Stored Program concept mean instructions just like data, so can take data from storage, and keep transforming it until load registers and jump to routine to begin execution
  - Compiler ⇒ Assembler ⇒ Linker (⇒ Loader )
- Assembler does 2 passes to resolve addresses, handling internal forward references
- Linker enables separate compilation, libraries that need not be compiled, and resolves remaining addresses



CS 61C L3.2.2 Floating Point 2 (62)

CS 61C L3.2.2 Floating Point 2 (62)

### Example: C ⇒ Asm ⇒ Obj ⇒ Exe ⇒ Run

```
#include <stdio.h>
int main (int argc, char *argv[]) {
    int i;
    int sum = 0;

    for (i = 0; i <= 100; i = i + 1)
        sum = sum + i * i;
    printf ("The sum from 0 .. 100 is %d\n",
           sum);
}
```



CS 61C L3.2.2 Floating Point 2 (63)

K. Meinz, Summer 2004 © UCB

### Example: C ⇒ Asm ⇒ Obj ⇒ Exe ⇒ Run

```
.text
.align 2
.globl main
main:
    subu $sp,$sp,32
    sw $ra, 20($sp)
    sd $a0, 32($sp)
    sw $0, 24($sp)
    sw $0, 28($sp)
loop:
    lw $t6, 28($sp)
    mul $t7, $t6,$t6
    lw $t8, 24($sp)
    addu $t9,$t8,$t7
    sw $t9, 24($sp)
str:
    .data
    .align 0
    .ascii "The sum
            from 0 .. 100 is
            %d\n"
```



CS 61C L3.2.2 Floating Point 2 (64)

CS 61C L3.2.2 Floating Point 2 (64)

K. Meinz, Summer 2004 © UCB

**Example: C ⇒ Asm ⇒ Obj ⇒ Exe ⇒ Run**

```

.text
.align 2
.globl main
main:
    subu $sp,$sp,32
    sw $ra, 20($sp)
    sd $a0, 32($sp)
    sw $0, 24($sp)
    sw $0, 28($sp)
loop:
    lw $t6, 28($sp)
    mul $t7, $t6,$t6
    lw $t8, 24($sp)
    addu $t9,$t8,$t7
    sw $t9, 24($sp)

```

*(Note: The original code has a bug where it adds the value of \$t7 to \$t8 instead of multiplying them.)*



CS 61C L3.2.2 Floating Point 2 (58)

```

    addu $t0, $t6, 1
    sw $t0, 28($sp)
    ble $t0,100, loop
    la $a0, str
    lw $a1, 24($sp)
    jal printf
    move $v0, $0
    lw $ra, 20($sp)
    addiu $sp,$sp,32
    j $ra

```

*(Note: The original code has a bug where it adds the value of \$t7 to \$t8 instead of multiplying them.)*

.data  
.align 0  
str:  
-asciiiz "The sum  
-from 0 .. 100 is  
%d\n"

K. Meinz, Summer 2004 © UCB

**Example: C ⇒ Asm ⇒ Obj ⇒ Exe ⇒ Run**

• Remove pseudoinstructions, assign addresses

```

00 addiu $29,$29,-32
04 sw    $31,20($29)
08 sw    $4, 32($29)
0c sw    $5, 36($29)
10 sw    $0, 24($29)
14 sw    $0, 28($29)
18 lw    $14, 28($29)
1c multu $14, $14
20 mflo $15
24 lw    $24, 24($29)
28 addu $25,$24,$15
2c sw    $25, 24($29)
30 addiu $8,$14, 1
34 sw    $8,28($29)
38 slti $1,$8, 101
3c bne $1,$0, -10
40 lui   $4, 1.str
44 ori   $4,$4,r.str
48 lw    $5,24($29)
4c jal   printf
50 add   $2, $0, $0
54 lw    $31,20($29)
58 addiu $29,$29,32
5c jr   $31

```



CS 61C L3.2.2 Floating Point 2 (68)

K. Meinz, Summer 2004 © UCB

**Example: C ⇒ Asm ⇒ Obj ⇒ Exe ⇒ Run**

- Example.o contains these tables:

- Symbol Table

Label	Address
main:	text+0x00000000
loop:	text+0x00000018
str:	data+0x00000000

- Relocation Information

Address	Instr.	Type	Dependency
text+0x00040	lui	l.str	
text+0x0044	ori	r.str	
text+0x004c	jal	printf	



CS 61C L3.2.2 Floating Point 2 (67)

K. Meinz, Summer 2004 © UCB

**Example: C ⇒ Asm ⇒ Obj ⇒ Exe ⇒ Run**

- Linker sees all the .o files.

- One of these (example.o) provides main and needs printf.
- Another (stdio.o) provides printf.

• 1) Linker decides order of text, data segments

• 2) This fills out the symbol tables

• 3) This fills out the relocation tables



CS 61C L3.2.2 Floating Point 2 (68)

K. Meinz, Summer 2004 © UCB

**Example: C ⇒ Asm ⇒ Obj ⇒ Exe ⇒ Run**

- Linker first stage:

• Set text=0x0400 0000; data=0x1000 0000

- Symbol Table

Label	Address
main:	0x04000000
loop:	0x04000018
str:	0x10000000

- Relocation Information

Address	Instr.	Type	Dependency
text+0x0040	lui	l.str	
text+0x0044	ori	r.str	
text+0x004c	jal	printf	



CS 61C L3.2.2 Floating Point 2 (69)

K. Meinz, Summer 2004 © UCB

**Example: C ⇒ Asm ⇒ Obj ⇒ Exe ⇒ Run**

- Linker second stage:

• Set text=0x0400 0000; data=0x1000 0000

- Symbol Table

Label	Address
main:	0x04000000
loop:	0x04000018
str:	0x10000000

- Relocation Information

Address	Instr.	Type	Dependency
text+0x0040	lui	l.str=0x1000	
text+0x0044	ori	r.str=0x0000	
text+0x004c	jal	printf=04440000	



CS 61C L3.2.2 Floating Point 2 (68)

K. Meinz, Summer 2004 © UCB

**Example: C  $\Rightarrow$  Asm  $\Rightarrow$  Obj  $\Rightarrow$  Exe  $\Rightarrow$  Run**

•Edit Addresses: start at 0x0400000

00 addiu \$29,\$29,-32	30 addiu \$8,\$14, 1
04 sw \$31,20(\$29)	34 sw \$8,28(\$29)
08 sw \$4, 32(\$29)	38 slti \$1,\$8, 101
0c sw \$5, 36(\$29)	3c bne \$1,\$0, -10
10 sw \$0, 24(\$29)	40 lui \$4, <b>1000</b>
14 sw \$0, 28(\$29)	44 ori \$4, <b>\$4,0000</b>
18 lw \$14, 28(\$29)	48 lw \$5,24(\$29)
1c multu \$14, \$14	4c jal <b>01110000</b>
20 mflo \$15	50 add \$2, \$0, \$0
24 lw \$24, 24(\$29)	54 lw \$31,20(\$29)
28 addu \$25,\$24,\$15	58 addiu \$29,\$29,32
2c sw \$25, 24(\$29)	5c jr \$31



CS 61C L3.2.2 Floating Point 2 (61)

K. Meinz, Summer 2004 © UCB

**Example: C  $\Rightarrow$  Asm  $\Rightarrow$  Obj  $\Rightarrow$  Exe  $\Rightarrow$  Run**

0x004000	0010011110111101111111111100000
0x004008	1010111110100100000000000000101000
0x00400c	1010111110100100000000000000100100
0x004010	1010111110100000000000000000111000
0x004014	1010111110100000000000000000111000
0x004018	1000111110101110000000000000111000
0x00401c	1000111110111100000000000000111000
0x004020	0000000111001110000000000000110001
0x004024	001001011100100000000000000000000001
0x004028	00101001000000001000000000001100101
0x00402c	1010111110101000000000000000111000
0x004030	0000000000000000011110000000100100
0x004034	000000110000111110010000010000100001
0x004038	000101000010000011111111110111
0x00403c	101011111011100100000000000011000
0x004040	001111000000010000001000000000000000
0x004044	1000111110100101000000000000110000
0x004048	000011000001000000000000000011101100
0x00404c	001001001000100000000100000110000
0x004050	1000111110111110000000000000101000
0x004054	0010011110111101000000000000100000
0x004058	00000011110000000000000000000000000010000
0x00405c	000000000000000000000000000000000000100001



CS 61C L3.2.2 Floating Point 2 (62)

K. Meinz, Summer 2004 © UCB