Lecture 6: Assembly Programs

- Today's topics:
 - Large constants
 - The compilation process
 - A full example
 - Intro to the MARS simulator

 Caller saved: Temp registers \$t0-\$t9 (the callee won't bother saving these, so save them if you care), \$ra (it's about to get over-written), \$a0-\$a3 (so you can put in new arguments)

 Callee saved: \$s0-\$s7 (these typically contain "valuable" data)

- Immediate instructions can only specify 16-bit constants
- The lui instruction is used to store a 16-bit constant into the upper 16 bits of a register... combine this with an OR instruction to specify a 32-bit constant
- The destination PC-address in a conditional branch is specified as a 16-bit constant, relative to the current PC
- A jump (j) instruction can specify a 26-bit constant; if more bits are required, the jump-register (jr) instruction is used

Starting a Program



- Convert pseudo-instructions into actual hardware instructions – pseudo-instrs make it easier to program in assembly – examples: "move", "blt", 32-bit immediate operands, etc.
- Convert assembly instrist into machine instrist a separate object file (x.o) is created for each C file (x.c) – compute the actual values for instruction labels – maintain info on external references and debugging information

• Stitches different object files into a single executable

- patch internal and external references
- determine addresses of data and instruction labels
- organize code and data modules in memory
- Some libraries (DLLs) are dynamically linked the executable points to dummy routines – these dummy routines call the dynamic linker-loader so they can update the executable to jump to the correct routine

Full Example – Sort in C (pg. 133)

```
void sort (int v[], int n)
{
    int i, j;
    for (i=0; i<n; i+=1) {
        for (j=i-1; j>=0 && v[j] > v[j+1]; j==1) {
            swap (v,j);
        }
    }
}
```

```
void swap (int v[], int k)
{
    int temp;
    temp = v[k];
    v[k] = v[k+1];
    v[k+1] = temp;
}
```

- Allocate registers to program variables
- Produce code for the program body
- Preserve registers across procedure invocations

 Register allocation: \$a0 and \$a1 for the two arguments, \$t0 for the temp variable – no need for saves and restores as we're not using \$s0-\$s7 and this is a leaf procedure (won't need to re-use \$a0 and \$a1)

swap:	sll	\$t1, \$a1, 2
	add	\$t1, \$a0, \$t1
	W	\$t0, 0(\$t1)
	W	\$t2, 4(\$t1)
	SW	\$t2, 0(\$t1)
	SW	\$t0, 4(\$t1)
	jr	\$ra

void swap (int v[], int k)
{
 int temp;
 temp = v[k];
 v[k] = v[k+1];
 v[k+1] = temp;
}

- Register allocation: arguments v and n use \$a0 and \$a1, i and j use \$s0 and \$s1; must save \$a0 and \$a1 before calling the leaf procedure
- The outer for loop looks like this: (note the use of pseudo-instrs)

move \$s0, \$zero # initialize the loop loopbody1: bge \$s0, \$a1, exit1 # will eventually use slt and beq ... body of inner loop ... addi \$s0, \$s0, 1 loopbody1

for (i=0; i<n; i+=1) { for (j=i-1; j>=0 && v[j] > v[j+1]; j-=1) { swap (v,j); 9

The inner for loop looks like this:

```
$s1, $s0, -1 # initialize the loop
           addi
                  $s1, $zero, exit2 # will eventually use slt and beq
loopbody2: blt
              $t1, $s1, 2
           sl
           add $t2, $a0, $t1
                  $t3, 0($t2)
           W
                  $t4, 4($t2)
           W
                  $t3, $t4, exit2
           bgt
           ... body of inner loop ...
                  $s1, $s1, -1
           addi
                   loopbody2
                                 for (i=0; i<n; i+=1) {
exit2:
                                   for (j=i-1; j>=0 && v[j] > v[j+1]; j-=1) {
                                       swap (v,j);
                                                                   10
```

- Since we repeatedly call "swap" with \$a0 and \$a1, we begin "sort" by copying its arguments into \$s2 and \$s3 – must update the rest of the code in "sort" to use \$s2 and \$s3 instead of \$a0 and \$a1
- Must save \$ra at the start of "sort" because it will get over-written when we call "swap"
- Must also save \$s0-\$s3 so we don't overwrite something that belongs to the procedure that called "sort"

Saves and Restores

sort: a s s s s s r r r	addi sw sw	\$sp, \$sp, -20 \$ra, 16(\$sp) \$s3, 12(\$sp)		
	SW SW SW MOVE MOVE	\$s2, 8(\$sp) \$s1, 4(\$sp) \$s0, 0(\$sp) \$s2, \$a0 \$s3, \$a1	9 lines of C code \rightarrow 35 lines of assem	bly
	move move jal	\$a0, \$s2 \$a1, \$s1 swap	# the inner loop body starts here	
exit1:	lw	\$s0, 0(\$sp)		
	addi jr	\$sp, \$sp, 20 \$ra	12	

- Intel's IA-32 instruction set has evolved over 20 years old features are preserved for software compatibility
- Numerous complex instructions complicates hardware design (Complex Instruction Set Computer – CISC)
- Instructions have different sizes, operands can be in registers or memory, only 8 general-purpose registers, one of the operands is over-written
- RISC instructions are more amenable to high performance (clock speed and parallelism) – modern Intel processors convert IA-32 instructions into simpler micro-operations.

Two major formats for transferring values between registers and memory

Memory: low address 45 7b 87 7f high address

Little-endian register: the first byte read goes in the low end of the register Register: 7f 87 7b 45 Most-significant bit / Least-significant bit (x86)

Big-endian register: the first byte read goes in the big end of the register Register: 45 7b 87 7f Most-significant bit / Least-significant bit (MIPS, IBM)



- MARS is a simulator that reads in an assembly program and models its behavior on a MIPS processor
- Note that a "MIPS add instruction" will eventually be converted to an add instruction for the host computer's architecture – this translation happens under the hood
- To simplify the programmer's task, it accepts pseudo-instructions, large constants, constants in decimal/hex formats, labels, etc.
- The simulator allows us to inspect register/memory values to confirm that our program is behaving correctly



main:

addi \$t0, \$zero, 5 addi \$t1, \$zero, 7 add \$t2, \$t0, \$t1

If we inspect the contents of \$t2, we'll find the number 12



Bullet