

COMP 412 FALL 2010

# Local Register Allocation & Lab 1 COMP 412

Extra slides at the end on live ranges may prove helpful.

Copyright 2010, Keith D. Cooper & Linda Torczon, all rights reserved.

Students enrolled in Comp 412 at Rice University have explicit permission to make copies of these materials for their personal use.

Faculty from other educational institutions may use these materials for nonprofit educational purposes, provided this copyright notice is preserved.

Where are we (and why)?

Local Register Allocation

- Chapter 13 in EAC covers Register Allocation
  - Read Sections 13.1, 13.2, & 13.3
  - Look at questions 1 and 2 in the exercises for Chapter 13
- Read Chapter 1 (for context) , look at Appendix A (for ILOC)
- Lab specs are on the class website
  - http://www.clear.rice.edu/comp412

Why are we here?

- Making time for scanning and parsing
- Providing implicit motivation for the start of the course
- And, allocation is both challenging and fun ...

Now, back to the lecture





Part of the compiler's back end



Critical properties

- Produce <u>correct</u> code that uses no more than k registers
- Minimize added work from loads and stores that *spill* values
- Minimize space used to hold spilled values
- Operate efficiently O(n), O(n log<sub>2</sub>n), maybe O(n<sup>2</sup>), but not O(2<sup>n</sup>)

Notation: The literature on register allocation consistently uses k as the number of registers available on the target system.

Comp 412, Fall 2010

r0 holds base address for local variables

Register Allocation / @x is constant offset of x from r0



Consider a fragment of assembly code (or ILOC)

loadI	2	$\Rightarrow$ r1	// r1 ← 2
loadAI	r0, @b	$\Rightarrow$ r2	// r2 ← b
mult	r1, r2	$\Rightarrow$ r3	// r3 ← 2 · b
loadAI	r0, @a	⇒r4	// r4 ← a
sub	r4, r3	$\Rightarrow$ r5	// r5 $\leftarrow$ a - (2 $\cdot$ b)

From the allocation perspective, these registers are virtual or pseudo-registers

The Problem

- At each instruction, decide which values to keep in registers
  - Note: each pseudo-register in the example is a value
- Simple if |values| ≤ |registers|
- Harder if |values| > |registers|
- The compiler must automate this process

# **Register** Allocation

The Task

- At each point in the code, pick the values to keep in registers
- Insert code to move values between registers & memory

   No transformations (leave that to optimization & scheduling)
- Minimize inserted code both dynamic & static measures
- Make good use of any *extra* registers

Allocation versus assignment

- Allocation is deciding which values to keep in registers
- Assignment is choosing specific registers for values
- This distinction is often lost in the literature

The compiler must perform both allocation & assignment and your lab



For straight-line code

dynamic  $\cong$  static

Basic Blocks in Assembly Code (or ILOC)

#### Definition

 A basic block is a maximal length segment of straight-line (i.e., branch free) code

Importance

- Strongest facts are provable for branch-free code
- If any statement executes, they all execute
- Execution is totally ordered

Role of Basic Blocks in Optimization

- Many techniques for improving basic blocks
- Simplest problems
- Strongest methods



Ignore, for the moment, exceptions

Local Register Allocation

- What is "local" ? (different from "regional" or "global")
  - A local transformation operates on basic blocks
  - Many optimizations are done on a local scale or scope
- Does local allocation solve the problem?
  - It produces good register use inside a block
  - Inefficiencies can arise at boundaries between blocks
  - Your lab assumes that the block is the entire program
    - Assumption significantly simplifies allocation
- How many passes can the allocator make?
  - This is an off-line problem
  - As many passes as it takes, within reason
    - You can do a fine job in a couple of passes



Blocks in a Controlflow Graph (CFG)

Comp 412, Fall 2010

# Register Allocation

Optimal register allocation is hard

#### Local Allocation

- Simplified cases  $\Rightarrow O(n)$
- Real cases  $\Rightarrow$  NP-Complete

#### <u>Global Allocation</u>

- NP-Complete for 1 register
- NP-Complete for k registers (most sub-problems are NPC, too)

#### Real compilers face real problems

#### Local Assignment

- Single size, no spilling  $\Rightarrow O(n)$
- Two sizes  $\Rightarrow$  NP-Complete

#### <u>Global Assignment</u>

• NP-Complete

#### Recent Results:

Optimal allocation on a procedure in SSA Form can be done in low-order polynomial time.

This result does not solve the entire problem, but it does offer insight into the structure of the problem & where the complexity lies.

We will come back to this issue later.



# ILOC



Your lab will do register allocation on basic blocks in "ILOC"

- Pseudo-code for a simple, abstracted RISC machine — generated by the instruction selection process
- Simple, compact data structures
- You will use a tiny subset of ILOC

a - 2 x b	loadI	2	$\Rightarrow$ <b>r</b> <sub>1</sub>	<u>ILOC:</u>
	loadAI	r <sub>0</sub> , @b	$\Rightarrow r_2$	• simple three-address code
	add	<b>r</b> <sub>1</sub> , <b>r</b> <sub>2</sub>	$\Rightarrow$ r <sub>3</sub>	• RISC-like addressing modes $\rightarrow T$ AT AO
	loadAI	r <sub>0</sub> , @a	$\Rightarrow$ r <sub>4</sub>	<ul> <li>unlimited virtual registers</li> </ul>
	sub	<b>r</b> <sub>4</sub> , <b>r</b> <sub>3</sub>	$\Rightarrow$ r <sub>5</sub>	5

ILOC is described in Appendix A of EAC

The subset used in lab 1 & 3 is described in the lab handout and on the last slide of this lecture

Comp 412, Fall 2010

# ILOC



Your lab will do register allocation on basic blocks in "ILOC"

- Pseudo-code for a simple, abstracted RISC machine
   generated by the instruction selection process
  - generated by the instruction selection pr
- Simple, compact data structures
- You will use a tiny subset of ILOC

a - 2 x b	loadI	2		r <sub>1</sub>
	loadAI	r <sub>0</sub>	@b	r <sub>2</sub>
	add	r <sub>1</sub>	r <sub>2</sub>	r <sub>3</sub>
	loadAI	r <sub>0</sub>	@a	r <sub>4</sub>
	sub	r <sub>4</sub>	r <sub>3</sub>	r <sub>5</sub>

Quadruples:

- table of k x 4 small integers
- simple record structure
- easy to reorder
- all names are explicit

ILOC is described in Appendix A of EAC

The subset used in lab 1 & 3 is described in the lab handout

### Observations



The Register Allocator does not need to "understand" the code

- It needs to distinguish definitions from uses
  - Definitions might need to store a spilled value
  - Uses might need to load a spilled value
- ILOC makes definitions and uses pretty clear
  - The assignment arrow,  $\Rightarrow$ , separates uses from definitions
    - Except on the store operation, which uses all its register operands
  - That is the point of the arrow!
- Your allocator needs to know, by opcode, how many definitions and how many uses it should see
  - Beyond that, the meaning of the ILOC is somewhat irrelevant to the allocator

### Observations



A value is live between its definition and its uses

- Find definitions (x  $\leftarrow$  ...) and uses (y  $\leftarrow$  ... x ...)
- From definition to <u>last</u> use is its live range
   How does a second definition affect this?
- Can represent live range as an interval [i,j] (in block)

Let MAXLIVE be the maximum, over each instruction *i* in the block, of the number of values (pseudo-registers) live at *i*.

- If MAXLIVE  $\leq k$ , allocation should be easy
- If MAXLIVE  $\leq k$ , no need to reserve F registers for spilling
- If MAXLIVE > k, some values must be spilled to memory
- If MAXLIVE > k, need to reserve F registers for spilling Finding live ranges is harder in the global case



Sample code sequence

loadI	1028	$\Rightarrow$ r1	// r1 ← 1028	
load	r1	$\Rightarrow$ r2	// r2 ← MEM	(r1)
mult	r1, r2	$\Rightarrow$ r3	// r3 ← 1028	·y
load	X	$\Rightarrow$ r4	// r4 ← x	
sub	r4, r2	$\Rightarrow$ r5	// $r5 \leftarrow x - y$	
load	Z	$\Rightarrow$ r6	// r6 ← z	
mult	r5, r6	$\Rightarrow$ r7	// r7 $\leftarrow z \cdot (x)$	- y)
sub	r7, r3	$\Rightarrow$ r8	// $r5 \leftarrow z \cdot (x)$	- y) - (1028 · y)
store	r8	⇒(r1)_	// MEM(r1) ←	$-z \cdot (x - y) - (1028 \cdot y)$
			<u> </u>	
				Store uses this register & defines a memory location.

The code uses 1028 as both the address of y and as a constant in the computation.

The intent is to create a long live range for pedagogical purposes. Remember, the allocator does not need to understand the computation. It just needs to preserve the computation.

Comp 412, Fall 2010



Live ranges in the example

loadI	1028	$\Rightarrow$ r1	// r1
load	r1	$\Rightarrow$ r2	// r1r2
mult	r1, r2	$\Rightarrow$ r3	// r1r2 r3
load	X	$\Rightarrow$ r4	// r1r2 r3 r4
sub	r4, r2	$\Rightarrow$ r5	// r1 r3 r5
load	Z	⇒r6	// r1 r3 r5 r6
mult	r5, r6	⇒r7	// r1 r3 r7
sub	r7, r3	$\Rightarrow$ r8	// r1 r8
store	r8	$\Rightarrow$ r1	//
			A pseudo-register is <u>live</u>
			after an operation if it has
			in the future
			in the juiure



Live ranges in the example



Top-down Versus Bottom-up Allocation



Top-down allocator

- Work from external notion of what is important
- Assign registers in priority order
- Save some registers for the values relegated to memory

Bottom-up allocator

- Work from detailed knowledge about problem instance
- Incorporate knowledge of partial solution at each step
- Handle all values uniformly

#### You will implement one of each



The idea

- Keep busiest values in a register
- Reserve registers for use in spills, say r registers

Algorithm

- Rank values by number of occurrences
- Allocate first k r values to registers
- Rewrite code to reflect these choices

Move values with no register into memory (add LOADs & STOREs)

Programmers applied this idea by hand in the 70's & early 80's

• C's register declaration

You will implement a variant (see Chapter 13, Question 1)

# Top-down Allocator



How many registers must the allocator reserve?

- Need registers to compute spill addresses & load values
- Number depends on target architecture
  - Typically, must be able to load 2 values
- Reserve these registers for spilling

What if k - r < |values| < k?

- Remember that the underlying problem is NP-Complete
- The allocator can either
  - Check for this situation
  - Adopt a more complex strategy

(iterate?)

- Accept the fact that the technique is an approximation



Fop down (3 re	egisters	, need 2	for ope	rands	)
	5		•		r1 is used more
loadI	1028	$\Rightarrow$ r1	// r1		often than r3
load	r1	$\Rightarrow$ r2	// r1r2		
mult	r1, r2	$\Rightarrow$ r3	// r1r2	r3	
load	X	$\Rightarrow$ r4	// r1r2	r3 r4	
sub	r4, r2	$\Rightarrow$ r5	// r1	r3	r5
load	Z	$\Rightarrow$ r6	// r1	r3	r5 r6
mult	r5, r6	$\Rightarrow$ r7	// r1	r3	r7
sub	r7, r3	⇒r8	// r1		r8
store	r8	$\Rightarrow$ r1	11		

Note that this assumes that no extra register is needed for spilling (Absolute location for each register — use load and store immediate)





Note that this assumes that no extra register is needed for spilling (Absolute location for each register — use load and store immediate)



#### Top down (3 registers, need 2 for operands)



"spill" and "restore" become stores and loads

- In practice, relative to  $r_{arp}$  ( $r_0$  in this lecture)
- In lab 1, to absolute addresses between 0 and 1023



#### Top down (3 registers, need 2 for operands)



The two short versions of r3 each overlap with fewer values, which simplifies the allocation problem. Such "spilling" will (eventually) create a code where the allocator can succeed.



Top down (3 registers, need 2 for operands)



This code is slower than the original, but it works correctly on a target machine with only three (available) registers. Correctness is a virtue. Bottom-up Allocator

Algorithm:

The idea:

- Start with empty register set
- Load on demand
- When no register is available, free one

Focus on replacement rather than allocation

Replacement:

- Spill the value whose next use is farthest in the future
- Prefer clean value to dirty value
- Sound familiar? Think page replacement ...





The algorithm should sound familiar

Decade algorithm

- Sheldon Best, 1955, for Fortran I
- Laslo Belady, 1965, for paging studies
- William Harrison, 1975, in ECS compiler work
- Chris Fraser, 1989, in the LCC compiler
- Forgotten student, 1995, COMP 412 at Rice
- Vincenzo Liberatore, 1997, Rutgers
- It will be reinvented again
- Many authors have argued for its optimality

(wrong)



Bottom up (3 registers; need 2 for operands)



Note that this assumes that no extra register is needed for spilling (Absolute location for each register — use load and store immediate)



Bottom up (3 registers; need 2 for operands)

loadI	1028	$\Rightarrow$ r1	// r1			
load	r1	$\Rightarrow$ r2	// r1 ri	2		chill n1
mult	r1, r2	$\Rightarrow$ r3	// <u>r1 ri</u>	2 r3		spiirit
loadI	X	$\Rightarrow$ r4	// r1 ra	2 r3 r4	•	
sub	r4, r2	$\Rightarrow$ r5	// r1	r3	r5	
loadI	Z	$\Rightarrow$ r6	// r1	r3	r5 r6	
mult	r5, r6	$\Rightarrow$ r7	// r1	r3		r7
sub	r7, r3	$\Rightarrow$ r8	// <u>r1</u>			r8
store	r8	$\Rightarrow$ r1	//			restore r1

Note that this assumes that no extra register is needed for spilling (Absolute location for each register — use load and store immediate)



Bottom up (3 registers; need 2 for operands)

The two short versions of r1 each overlap with fewer values, which simplifies the allocation problem. Such "spilling" will (eventually) create a code where the allocator can succeed.

# Lab One



Your Task

- Implement a version of the top-down allocator (See EaC Section 13.3.1 & Exercise 2a, Section 13.3)
- Implement a version of the bottom-up allocator (See EaC Section 13.3.2)
- Run them on a collection of test blocks
- Write up a report
  - Describe the experience
  - Compare the two allocators

Due date: Wednesday, September 15, 2010, 11:59 PM Documentation: Friday, September 19, 2010, 11:59 PM Test & report blocks are available on the web site



Your allocator is not bound by the names used in its input

- Every computed value is part of some live range
   Even if it has no name in the source code (e.g., 2 \* y in x 2 \* y)
- A live range usually has a single name, such as  $r_{17}$
- A single name with multiple values can be renamed into distinct live ranges

Live	Ranges
------	--------

# (From Figure 13.3 in EaC)



	Operati	ion		Live Rai	nges
loadI	@base	$\Rightarrow$	<b>r</b> <sub>arp</sub>	none	
loadAI	r <sub>arp</sub> , @a	$\Rightarrow$	r <sub>a</sub>	r <sub>arp</sub> ,r <sub>a</sub>	
loadI	2	$\Rightarrow$	r <sub>2</sub>	r <sub>arp</sub> ,r <sub>a</sub> ,r <sub>2</sub>	
loadAI	r <sub>arp</sub> /@b	$\Rightarrow$	r <sub>b</sub>	r <sub>arp</sub> ,r <sub>a</sub> ,r <sub>2</sub>	
loadAI	r <sub>arp</sub> , @c	$\Rightarrow$	r <sub>c</sub>	r <sub>arp</sub> ,r <sub>a</sub> ,r <sub>b</sub> ,r <sub>2</sub>	
loadAI	r <sub>arp</sub> , @d	$\Rightarrow$	r <sub>d</sub>	r <sub>arp</sub> ,r <sub>a</sub> ,r <sub>c</sub> ,r <sub>b</sub> ,r <sub>2</sub>	
mult	r <sub>a</sub> , r <sub>2</sub>	$\Rightarrow$	ra	r <sub>arp</sub> ,r <sub>a</sub> ,r <sub>d</sub> ,r <sub>c</sub> ,r <sub>b</sub> ,r	2
mult	r <sub>a</sub> , r <sub>b</sub>	$\Rightarrow$	ra	r <sub>arp</sub> ,r <sub>a</sub> ,r <sub>d</sub> ,r <sub>c</sub> ,r <sub>b</sub>	
mult	r <sub>a</sub> , r <sub>c</sub>	$\Rightarrow$	ra	r <sub>arp</sub> ,r <sub>a</sub> ,r <sub>d</sub> ,r <sub>c</sub>	I here are five distinct values, or live ranges.
mult	r <sub>a</sub> , r <sub>d</sub>	$\Rightarrow$	ra	r <sub>arp</sub> ,r <sub>a</sub> ,r <sub>d</sub>	named r <sub>a</sub>
storeAI	r	$\Rightarrow$	r <sub>arp</sub> , @w		The last 4 fit in F.



Your allocator is not bound by the names used in its input

- Every computed value is part of some live range
   Even if it has no name in the source code (e.g., 2 \* y in x 2 \* y)
- A live range usually has a single name, such as  $r_{17}$
- A single name with multiple values can be renamed into distinct live ranges
- Renaming distinct live ranges with distinct names can simplify the implementation of the allocator
   It can also help with debugging the allocator



ILOC subset:



These same operations, with different latencies, will appear in lab 3 Assume a register-to-register memory model, with 1 class of registers