# EECS 583 – Class 15 Register Allocation

University of Michigan

*November 2, 2011* 

# Announcements + Reading Material

- Midterm exam: Monday, Nov 14?
  - » Could also do Wednes Nov 9 (next week!) or Wednes Nov 16 (2 wks from now)
  - » Class vote
- Today's class reading
  - » "Register Allocation and Spilling Via Graph Coloring," G. Chaitin, Proc. 1982 SIGPLAN Symposium on Compiler Construction, 1982.
- Next class reading More at the end of class
  - » "Revisiting the Sequential Programming Model for Multi-Core," M. J. Bridges, N. Vachharajani, Y. Zhang, T. Jablin, and D. I. August, *Proc 40th IEEE/ACM International Symposium on Microarchitecture*, December 2007.

#### Homework Problem – Answers in Red

latencies: add=1, mpy=3, ld = 2, st = 1, br = 1

LC = 99

1: $r3 = load(r1)$
2: r4 = r3 * 26
3: store (r2, r4)
4: $r1 = r1 + 4$
5: $r^2 = r^2 + 4$
7: brlc Loop

How many resources of each type are required to achieve an II=1 schedule? For II=1, each operation needs a dedicated resource, so: 3 ALU, 2 MEM, 1 BR

If the resources are non-pipelined, how many resources of each type are required to achieve II=1 Instead of 1 ALU to do the multiplies, 3 are needed, and instead of 1 MEM to do the loads, 2 are needed. Hence: 5 ALU, 3 MEM, 1 BR

Assuming pipelined resources, generate the II=1 modulo schedule. See next few slides

#### HW continued

Assume II=1 so resources are: 3 ALU, 2 MEM, 1 BR



#### HW continued

resources: 3 alu, 2 mem, 1 br latencies: add=1, mpy=3, Id = 2, st = 1, br = 1

#### LC = 99

Schedule op7 at time 5

Loop:	1:1
	2: 1
	3:



Unrolled

#### HW continued

The final loop consists of a single MultiOp containing 6 operations, each predicated on the appropriate staging predicate. Note register allocation still needs to be performed.

LC = 99

Loop:

 $r_{3}[-1] = load(r_{1}[0])$  if  $p_{1}[0]$ ;  $r_{4}[-1] = r_{3}[-1] * 26$  if  $p_{1}[2]$ ; store ( $r_{2}[0], r_{4}[-1]$ ) if  $p_{1}[5]$ ;  $r_{1}[-1] = r_{1}[0] + 4$  if  $p_{1}[0]$ ;  $r_{2}[-1] = r_{2}[0] + 4$  if  $p_{1}[5]$ ; brf Loop

# Register Allocation: Problem Definition

- Through optimization, assume an infinite number of virtual registers
  - » Now, must allocate these infinite virtual registers to a limited supply of hardware registers
  - » Want most frequently accessed variables in registers
    - Speed, registers much faster than memory
    - Direct access as an operand
  - » Any VR that cannot be mapped into a physical register is said to be <u>spilled</u>
- Questions to answer
  - » What is the minimum number of registers needed to avoid spilling?
  - » Given n registers, is spilling necessary
  - » Find an assignment of virtual registers to physical registers
  - » If there are not enough physical registers, which virtual registers get spilled?

# Live Range

- Value = definition of a register
- Live range = Set of operations
  - » 1 more or values connected by common uses
  - » A single VR may have several live ranges
  - » Very similar to the web being constructed for HW3
- Live ranges are constructed by taking the intersection of reaching defs and liveness
  - Initially, a live range consists of a single definition and all ops in a function in which that definition is live

#### Example – Constructing Live Ranges



Each definition is the seed of a live range. Ops are added to the LR where <u>both the defn reaches</u> and the variable is live

> LR1 for def  $1 = \{1,3,4\}$ LR2 for def  $2 = \{2,4\}$ LR3 for def  $5 = \{5,7,8\}$ LR4 for def  $6 = \{6,7,8\}$

### Merging Live Ranges

- If 2 live ranges for the same VR overlap, they must be merged to ensure correctness
  - » LRs replaced by a new LR that is the union of the LRs
  - » Multiple defs reaching a common use
  - » Conservatively, all LRs for the same VR could be merged
    - Makes LRs larger than need be, but done for simplicity
    - We will not assume this



#### Example – Merging Live Ranges



#### Class Problem



#### Compute the LRs

- ) for each def
- ) merge overlapping

#### Interference

- Two live ranges interfere if they share one or more ops in common
  - » Thus, they cannot occupy the same physical register
  - » Or a live value would be lost
- Interference graph
  - » Undirected graph where
    - Nodes are live ranges
    - There is an edge between 2 nodes if the live ranges interfere
  - » What's not represented by this graph
    - Extent of interference between the LRs
    - Where in the program is the interference

Example – Interference Graph



# Graph Coloring

- A graph is <u>n-colorable</u> if every node in the graph can be colored with one of the n colors such that 2 adjacent nodes do not have the same color
  - » Model register allocation as graph coloring
  - » Use the fewest colors (physical registers)
  - » Spilling is necessary if the graph is not n-colorable where n is the number of physical registers
- Optimal graph coloring is NP-complete for n > 2
  - » Use heuristics proposed by compiler developers
    - "Register Allocation Via Coloring", G. Chaitin et al, 1981
    - "Improvement to Graph Coloring Register Allocation", P. Briggs et al, 1989
  - » <u>Observation</u> a node with degree < n in the interference can always be successfully colored given its neighbors colors

# Coloring Algorithm

- ✤ 1. While any node, x, has < n neighbors</p>
  - » Remove x and its edges from the graph
  - » Push x onto a stack
- ✤ 2. If the remaining graph is non-empty
  - » Compute cost of spilling each node (live range)
    - For each reference to the register in the live range
      - Cost += (execution frequency \* spill cost)
  - > Let NB(x) = number of neighbors of x
  - » Remove node x that has the smallest cost(x) / NB(x)
    - Push x onto a stack (mark as spilled)
  - » Go back to step 1
- While stack is non-empty
  - » Pop x from the stack
  - » If x's neighbors are assigned fewer than R colors, then assign x any unsigned color, else leave x uncolored

#### Example – Finding Number of Needed Colors

How many colors are needed to color this graph?



Try n=1, no, cannot remove any nodes

Try n=2, no again, cannot remove any nodes

Try n=3,

Remove B Then can remove A, C Then can remove D, E Thus it is 3-colorable

#### Example – Do a 3-Coloring



$$lr(a) = \{1,2,3,4,5,6,7,8\}$$
  

$$refs(a) = \{1,6,8\}$$

$$lr(b) = \{2,3,4,6\}$$
  

$$refs(b) = \{2,4,6\}$$
  

$$lr(c) = \{1,2,3,4,5,6,7,8,9\}$$
  

$$refs(c) = \{3,4,7\}$$
  

$$lr(d) = \{4,5\}$$
  

$$lr(d) = \{4,5\}$$
  

$$lr(e) = \{5,7,8\}$$
  

$$refs(e) = \{5,7,8\}$$
  

$$lr(f) = \{6,7\}$$
  

$$refs(f) = \{6,7\}$$
  

$$refs(g) = \{8,9\}$$
  
Arefs(g) = \{8,9\}  

$$refs(g) = \{8,9\}$$

	a	b	С	d	e	f	g
cost	225	200	175	150	200	50	200
neighbors	6	4	5	4	3	4	2
cost/n	37.5	50	35	37.5	66.7	12.5	100

# Example – Do a 3-Coloring (2)

Remove all nodes < 3 neighbors

So, g can be removed

<u>Stack</u> g



#### Example – Do a 3-Coloring (3)

Now must spill a node

Choose one with the smallest cost/NB  $\rightarrow$  f is chosen

<u>Stack</u> f (spilled) g



# Example – Do a 3-Coloring (4)

Remove all nodes < 3 neighbors	<u>Stack</u>
So, e can be removed	e f (spilled)
	g



#### Example – Do a 3-Coloring (5)

Now must spill another node

Choose one with the smallest cost/NB  $\rightarrow$  c is chosen

Stack c (spilled) e f (spilled) g



# Example – Do a 3-Coloring (6)



# Example – Do a 3-Coloring (7)



Have 3 colors: red, green, blue, pop off the stack assigning colors only consider conflicts with non-spilled nodes already popped off stack

 $d \rightarrow red$ 

- $b \rightarrow$  green (cannot choose red)
- a  $\rightarrow$  blue (cannot choose red or green)
- $c \rightarrow$  no color (spilled)
- $e \rightarrow$  green (cannot choose red or blue)
- f  $\rightarrow$  no color (spilled)
- $g \rightarrow red$  (cannot choose blue)

# Example – Do a 3-Coloring (8)



#### Homework Problem



do a 2-coloring compute cost matrix draw interference graph color graph

#### It's not that easy – Iterative Coloring



You can't spill withoutcreating more live rangesNeed regs for the stackptr, value spilled, [offset]

Can't color before taking this into account

# Iterative Coloring (1)



# Iterative Coloring (2)





- 3. Update interference graph
  - Nuke edges between spilled LRs

# Iterative Coloring (4)



- 3. Add edges for new/spilled LRs
  - Stack ptr (almost) always interferes with everything so ISAs usually just reserve a reg for it.
- 4. Recolor and repeat until no new spill is generated

# Time to Switch Gears – Research Topics!

#### Topics We Will Cover

- ✤ 1. Automatic Parallelization
- Optimizing Streaming Applications for Multicore/GPUs
- ✤ 3. Automatic SIMDization
- ✤ 4. TBD

# Paper Reviews

- ✤ 1 per class for the rest of the semester
- Paper review purpose
  - » Read the paper <u>before</u> class up to this point reading has not been a point of emphasis, now it is!
  - » Put together a list of non-trivial observations think about it!
  - » Have something interesting to say in class
- Review content -2 parts
  - » 1. 3-4 line summary of paper
  - » 2. Your thoughts/observations it can be any of the following:
    - An interesting observation on some aspect of the paper
    - Raise one or more questions that a cursory reader would not think of
    - Identify a weakness of the approach and explain why
    - Propose an extension to the idea that fills some hole in the work
    - Pick a difficult part of the paper, decipher it, and explain it in your own words

#### Paper Reviews – continued

- Review format
  - » Plain text only .txt file
  - » <sup>1</sup>/<sub>2</sub> page is sufficient
- Reviews are due by the start of each lecture
  - » Copy file to andrew.eecs.umich.edu:/y/submit
  - » Put uniquename\_classXX.txt
- First reading due Monday Nov 7 (pdf on the website)
  - \* "Revisiting the Sequential Programming Model for Multi-Core," M. J. Bridges, N. Vachharajani, Y. Zhang, T. Jablin, and D. I. August, *Proc 40th IEEE/ACM International Symposium on Microarchitecture*, December 2007.