

More Register Allocation

Last time

- Register allocation
 - Global allocation via graph coloring

Today

- More register allocation
 - Procedure calls
 - Interprocedural

Register Allocation and Procedure Calls

Problem

- Register values may change across procedure calls
- The allocator must be sensitive to this

Two approaches

- Work within a well-defined calling convention } Make local decisions
- Use interprocedural allocation } Make global decisions

Calling Conventions

Goals

- Fast calls (pass arguments in registers, minimal register saving/restoring)
- Language-independent
- Support debugging, profiling, *etc.*

Complicating Issues

- Varargs
- Passing/returning aggregates
- Exceptions, non-local returns
 - `setjmp()`/`longjmp()`
- Non-LIFO activation records

April 15, 2015

More Register Allocation

3

Architecture Review: Caller- and Callee-Saved Registers

Partition registers into two categories

- Caller-saved
- Callee-saved

Caller-saved registers

- Caller must save/restore these registers when live across call
- Callee is free to use them

Example

```
foo ()
{
  rcaller = 4
  save rcaller
  goo ()
  restore rcaller
  use rcaller
}
```

caller

```
goo ()
{
  rcaller = 99
}
```

callee

goo () is free to modify r_{caller}

April 15, 2015

More Register Allocation


4

Architecture Review: Caller- and Callee-Saved Registers

Callee-saved registers

- Callee must save/restore these registers when it uses them
- Caller expects callee to not change them

Example



April 15, 2015

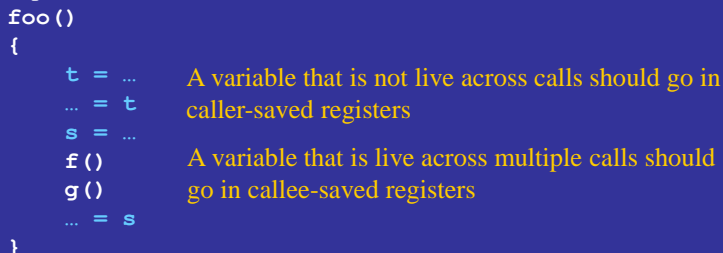
More Register Allocation

5

Register Allocation and Calling Conventions

Insensitive register allocation

- Save all live caller-saved registers before call; restore after
- Save all used callee-saved registers at procedure entry; restore at return
- Suboptimal



Sensitive register allocation

- Encode calling convention constraints in the IR and interference graph
- How? Use precolored nodes

April 15, 2015

More Register Allocation

6


Precolored Nodes

Add architectural registers to interference graph

- Precolored (mutually interfering)
- Not simplifiable
- Not spillable (infinite degree)

Express allocation constraints

- Integers usually can't be stored in floating point registers
- Some instructions can only store result in certain registers
- Caller-saved and callee-saved registers. . .



April 15, 2015

More Register Allocation

7

Precolored Nodes and Calling Conventions

Callee-saved registers

- Treat entry as def of all callee-saved registers
- Treat exit as use of them all
- Allocator must "spill" callee-saved registers to use them

Encourage use of callee-saved regs



Caller-saved registers

- Variables live across call interfere with all caller-saved registers
- Splitting can be used (before/during/after call segments)

Encourage use of callee-saved regs

April 15, 2015


More Register Allocation

8

Problem with Callee-Saved Registers

Run-time systems (e.g., `setjmp()`/`longjmp()` and debuggers) need to know register values in any stack frame

- Caller-saved registers are on stack frame at known location
- Callee-saved registers?



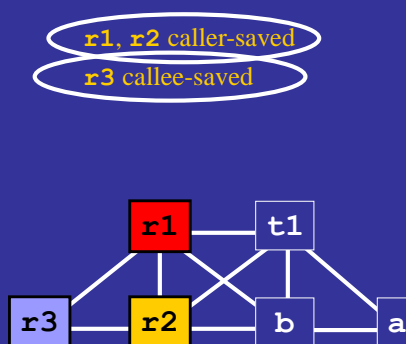
April 15, 2015

More Register Allocation

9

Example

```
foo() :
  def(r3)
  t1 := r3
  a := ...
  b := ...
  ... a ...
  call goo
  ... b ...
  r3 := t1
  use(r3)
  return
```



April 15, 2015

More Register Allocation

10

Tradeoffs

Callee-saved registers

- + Decreases code size: one procedure body may have multiple calls
- + Small procedures tend to need fewer registers than large ones; callee-save makes sense because procedure sizes are shrinking
- May increase execution time: For long-lived variables, may save and restore registers multiple times, once for each procedure, instead of a single end-to-end save/restore

The larger “problem”

- We’re making local decisions for policies that require global information

April 15, 2015

More Register Allocation

11

Interprocedural Register Allocation

Wouldn’t it be nice to . . .

- Allocate registers across calls to minimize unnecessary saves/restores?
- Allocate global variables to registers over entire program?

Compile-time interprocedural register allocation?

- + Could have great performance
- Might be expensive
- Might require lots of recompilation after changes (no separate compilation?)

Link-time interprocedural re-allocation?

- + Low compile-time cost
- + Little impact on separate compilation
- Link-time cost

April 15, 2015

More Register Allocation

12

Wall's Link-time Register Allocator [Wall 86]

Overall strategy

- Compiler uses 8 registers for local register allocation
- Linker controls allocation of remaining 52 registers

Compiler does local allocation & planning for linker

- Load all values at beginning of each basic block; store all values at end of each basic block
- Generate call graph information
- Generate variable usage information for each procedure
- Generate **register actions**

Linker does interprocedural allocation & patches compiled code

- Generates "interference graph" among variables
- Picks best variables to allocate to registers
- Executes register actions for allocated variables to patch code

April 15, 2015

More Register Allocation

13

Register Actions

Describe code patch if particular variable allocated to a register

- **REMOVE(var)**: Delete instruction if **var** allocated to a register
- **OPx(var)**: Replace op x with register that was allocated to **var**
- **RESULT(var)**: Replace result with register allocated to **var**

Usage

- **r := load var: REMOVE(var)**
- **ri := rj op rk:**
 - OP1(var)** if **var** loaded into **rj**
 - OP2(var)** if **var** loaded into **rk**
 - RESULT(var)** if **var** stored from **ri**
- **store var := r: REMOVE(var)**

April 15, 2015

More Register Allocation

14

Example

w := (x + y) * z

r1 := load x	REMOVE(x)
r2 := load y	REMOVE(y)
r3 := r1 + r2	OP1(x), OP2(y)
r4 := load z	REMOVE(z)
r5 := r3 * r4	OP2(z), RESULT(w)
store w := r5	REMOVE(w)

April 15, 2015

More Register Allocation

15

Another Example

w := y++ * z

Suppose **y** is allocated to register **r5**

	REMOVE(y)
r5 := r5 + 1	OP1(y), RESULT(y)
	REMOVE(y)
r2 := load z	REMOVE(z)
r1 := r5 * r2	OP1(y), OP2(z), RESULT(w)
	REMOVE(w)

Problem

- Loaded value is still live after store overwrites it
- Post-incremented value of **y** is lost if **y** is allocated to register
- We need two registers to hold the two values of **y**

April 15, 2015

More Register Allocation

16

Extension

More actions

- **LOAD(var)**: Replace load with move from the register holding **var**
- **STORE(var)**: Replace store with move to the register holding **var**

LOAD(var)

- Use instead of **REMOVE(var)** if **var** is stored into while result of load is still live

STORE(var)

- Use instead of **REMOVE(var)** if source is stored into more than one variable

April 15, 2015

More Register Allocation

17

Example Revisited

w := x := y++ * z

r1 := load y	REMOVE(y)	LOAD(y)
r2 := r1 + 1	OP1(y), RESULT(y)	RESULT(y)
store y := r2	REMOVE(y)	REMOVE(y)
r2 := load z	REMOVE(z)	REMOVE(z)
r1 := r1 * r2	OP1(y), OP2(z), RESULT(w)	OP2(z), RESULT(w)
store x := r1		STORE(x), OP1(w)
store w := r1	REMOVE(w)	REMOVE(w)

April 15, 2015

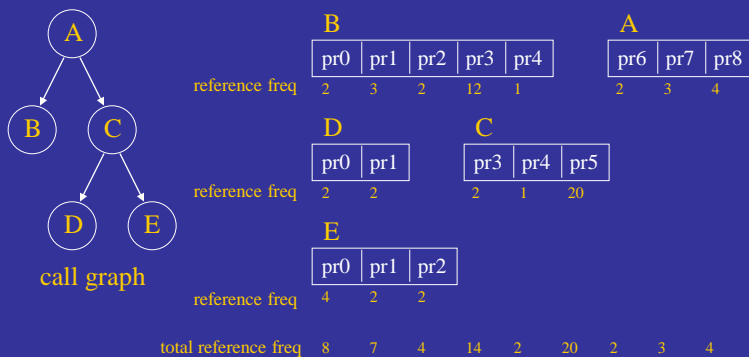
More Register Allocation

18

Deciding Which Variables to Promote to Registers

Steps

- Use bottom-up algorithm to assign pseudo registers
- Allocate pseudo registers to non-simultaneously live variables
- Allocate real registers to most frequently used pseudo registers



April 15, 2015

More Register Allocation

19

Possible Improvements

Use profile data to construct weights

Do global register allocation at compile-time

Track liveness information for variables at each call site

Track intraprocedural interference graph

Use real interference graph at link-time

April 15, 2015

More Register Allocation

20

Performance Summary

Machine: DEC WRL Titan RISC processor (64 registers)

Basic experiment

- Local compile-time allocator uses 8 registers
 - Link-time allocator uses 52 registers
 - Simple static frequency estimates
 - Small benchmarks
- ⇒ 10-25% speed-up over local allocation alone

Improvements

- 0-6% with profile data
- 0-5% with compile-time global allocation

Benefit decreases with number of link-time registers

Link-time better than global register allocation

April 15, 2015

More Register Allocation

21

Link-Time Register Allocation: The Big Picture

Delayed decision making

- Make decisions when more information is available, eg. link time
- Requires communication among different system components, in this case the compiler and the linker
- Leads to staged compilation
- Intuitively, more information is better, but effectively using this information can require cleverness

April 15, 2015


More Register Allocation

22

JIT Environment

Dynamic compilation requires fast register allocation

- Linear Scan Register Allocation [Poletto & Sarkar99]
- Not based on graph coloring
- Greedy algorithm based on live intervals
 - Spill the variable whose interval ends furthest in the future



- What if we had spilled a or b instead of c?

April 15, 2015

More Register Allocation

23

Linear Scan Register Allocation

Performance results

- Linear scan is linear in number of variables
- Graph coloring is $O(n^2)$
- Code quality is within 12% of graph coloring

April 15, 2015

More Register Allocation

24


Register Allocation Using SSA Form [Hack, et al 2006]

Graph Coloring

- NP-complete
- Register allocation based on graph coloring is NP-complete

Use heuristics

- Coalesce, Color, Spill, repeat . . .



New approach

- Use graph theory and SSA



April 15, 2015

More Register Allocation

25

The Highlights

SSA form

- The interference graphs of programs in SSA form are **chordal**, which are perfect graphs

Perfect graphs

- Perfect graphs are k -colorable, where k is the size of the largest clique

Spilling

- Can determine exactly where spills must occur without using an interference graph

Coloring

- In general, can color chordal graphs in $O(|V|^2)$
- With dominance information, can color chordal graphs in $O(\omega(G) \cdot n)$ time, where $n = \#$ of instructions and $\omega(G) =$ size of largest live set

April 15, 2015

More Register Allocation

26

Evaluation

Benefits

- No interference graph
- Decouples spilling and coalescing

Good way to trade off coalescing vs. copying

- Coalescing can reduce copies
 - eg. `mov t1, t2`
- Coalescing can increase spills



- In traditional register allocation, removing ϕ -functions is equivalent to prematurely coalescing SSA variables
- By keeping ϕ -functions, we coalesce as much as possible without incurring spills

April 15, 2015

More Register Allocation

27

Concepts

Register allocation and procedure calls

Calling conventions

- Caller- vs. callee-saved registers
- Precoloring
- Finding register values in stack can be hard

Interprocedural analysis

- Link-time register allocation
 - Register actions

Register allocation in a JIT

- Linear Scan Register Allocation

April 15, 2015

More Register Allocation

28

Next Time

Course projects

- How are projects coming along? Feel free to send me updates

Reading

- Read SoftBound paper
- Response due Sunday night