Lecture 25
CIS 341: COMPILERS

Announcements

- HW6: Dataflow Analysis
 - Due: Weds. April 26th

NOTE: See Piazza for an update... TLDR: "simple" regalloc should not suffice. Change gradedtests.ml >= to >

• FINAL EXAM

- Thursday, May 4th noon 2:00p.m.
- Location: DRLB A4

LOOPS AND DOMINATORS

Zdancewic CIS 341: Compilers

Loops in Control-flow Graphs

- Taking into account loops is important for optimizations.
 - The 90/10 rule applies, so optimizing loop bodies is important
- Should we apply loop optimizations at the AST level or at a lower representation?
 - Loop optimizations benefit from other IR-level optimizations and vice-versa, so it is good to interleave them.
- Loops may be hard to recognize at the quadruple / LLVM IR level.
 - Many kinds of loops: while, do/while, for, continue, goto...
- Problem: *How do we identify loops in the control-flow graph?*

Definition of a Loop

- A *loop* is a set of nodes in the control flow graph.
 - One distinguished entry point called the *header*
- Every node is reachable from the header & the header is reachable from every node.
 - A loop is a strongly connected component
- No edges enter the loop except to the header
- Nodes with outgoing edges are called loop exit nodes



Nested Loops

- Control-flow graphs may contain many loops
- Loops may contain other loops:





The *control tree* depicts the nesting structure of the loops in the program.

Control-flow Analysis

- Goal: Identify the loops and nesting structure of the CFG.
- Control flow analysis is based on the idea of *dominators*:
- Node A *dominates* node B if the only way to reach B from the start node is through node A.
- An edge in the graph is a *back edge* if the target node dominates the source node.
- A loop contains at least one back edge.



Dominator Trees

- Domination is transitive:
 - if A dominates B and B dominates C then A dominates C
- Domination is anti-symmetric:
 - if A dominates B and B dominates A then A = B
- Every flow graph has a dominator tree
 - The Hasse diagram of the dominates relation



Dominator Dataflow Analysis

- We can define Dom[n] as a forward dataflow analysis.
 - Using the framework we saw earlier: Dom[n] = out[n] where:
- "A node B is dominated by another node A if A dominates *all* of the predecessors of B."
 - in[n] := $\bigcap_{n' \in pred[n]} out[n']$
- "Every node dominates itself."
 - $out[n] := in[n] \cup \{n\}$
- Formally: $\mathcal{L} = \text{set of nodes ordered by } \subseteq$
 - $T = \{all nodes\}$
 - $\ F_n(x) = x \ U \ \{n\}$
 - \square is \cap
- Easy to show monotonicity and that F_n distributes over meet.
 - So algorithm terminates and is MOP

Improving the Algorithm

- Dom[b] contains just those nodes along the path in the dominator tree from the root to b:
 - e.g. $Dom[8] = \{1, 2, 4, 8\}, Dom[7] = \{1, 2, 4, 5, 7\}$
 - There is a lot of sharing among the nodes
- More efficient way to represent Dom sets is to store the dominator *tree*.
 - doms[b] = immediate dominator of b
 - doms[8] = 4, doms[7] = 5
- To compute Dom[b] walk through doms[b]
- Need to efficiently compute intersections of Dom[a] and Dom[b]
 - Traverse up tree, looking for least common ancestor:
 - Dom[8] \cap Dom[7] = Dom[4]



• See: "A Simple, Fast Dominance Algorithm" Cooper, Harvey, and Kennedy

Completing Control-flow Analysis

- Dominator analysis identifies *back edges*:
 - Edge n \rightarrow h where h dominates n
- Each back edge has a *natural loop*:
 - h is the header
 - All nodes reachable from h that also reach n without going through h
- For each back edge $n \rightarrow h$, find the natural loop:
 - $\{n' \mid n \text{ is reachable from } n' \text{ in } G \{h\}\} \cup \{h\}$
- Two loops may share the same header: merge them
- Nesting structure of loops is determined by set inclusion
 - Can be used to build the control tree





Example Natural Loops



Control Tree:

The control tree depicts the nesting structure of the program.

Natural Loops

Uses of Control-flow Information

- Loop nesting depth plays an important role in optimization heuristics.
 - Deeply nested loops pay off the most for optimization.
- Need to know loop headers / back edges for doing
 - loop invariant code motion
 - loop unrolling
- Dominance information also plays a role in converting to SSA form
 - Used internally by LLVM to do register allocation.

Phi nodes Alloc "promotion" Register allocation

REVISITING SSA

Zdancewic CIS 341: Compilers

Single Static Assignment (SSA)

- LLVM IR names (via **%uids**) *all* intermediate values computed by the program.
- It makes the order of evaluation explicit.
- Each **%uid** is assigned to only once
 - Contrast with the mutable quadruple form
 - Note that dataflow analyses had these kill[n] sets because of updates to variables...
- Naïve implementation of backend: map **%uids** to stack slots
- Better implementation: map **%uids** to registers (as much as possible)
- Question: How do we convert a source program to make maximal use of **%uids**, rather than alloca-created storage?
 - two problems: control flow & location in memory
- Then: How do we convert SSA code to x86, mapping **%uids** to registers?
 - Register allocation.

Alloca vs. %UID

• Current compilation strategy:





• Directly map source variables into **%uids**?



• Does this always work?

What about If-then-else?

• How do we translate this into SSA?



```
entry:
  %y1 = ...
  %x1 = ...
  %z1 = ...
  %p = icmp ...
  br i1 %p, label %then, label %else
then:
  %x2 = add i64 %y1, 1
  br label %merge
else:
  %x3 = mult i64 %y1, 2
merge:
  %z2 = %add i64 ???, 3
```

• What do we put for ???

Phi Functions

- Solution: φ functions
 - Fictitious operator, used only for analysis
 - implemented by Mov at x86 level
 - Chooses among different versions of a variable based on the path by which control enters the phi node.

 $\texttt{%uid} = \texttt{phi} < \texttt{ty} > \texttt{v}_1, < \texttt{label}_1 >, \dots, \texttt{v}_n, < \texttt{label}_n >$



Phi Nodes and Loops

- Importantly, the **%uids** on the right-hand side of a phi node can be defined "later" in the control-flow graph.
 - Means that **%uids** can hold values "around a loop"
 - Scope of **%uids** is defined by dominance (discussed soon)

```
entry:
  %y1 = ...
  %x1 = ...
  br label %body
body:
  %x2 = phi i64 %x1, %entry, %x3, %body
  %x3 = add i64 %x2, %y1
  %p = icmp slt i64, %x3, 10
  br i1 %p, label %body, label %after
after:
  ...
```

Alloca Promotion

- Not all source variables can be allocated to registers
 - If the address of the variable is taken (as permitted in C, for example)
 - If the address of the variable "escapes" (by being passed to a function)
- An alloca instruction is called promotable if neither of the two conditions above holds

```
entry:
%x = alloca i64 // %x cannot be promoted
%y = call malloc(i64 8)
%ptr = bitcast i8* %y to i64**
store i65** %ptr, %x // store the pointer into the heap
```

- Happily, most local variables declared in source programs are promotable
 - That means they can be register allocated

Converting to SSA: Overview

- Start with the ordinary control flow graph that uses allocas
 - Identify "promotable" allocas
- Compute dominator tree information
- Calculate def/use information for each such allocated variable
- Insert φ functions for each variable at necessary "join points"
- Replace loads/stores to alloc'ed variables with freshly-generated %uids
- Eliminate the now unneeded load/store/alloca instructions.

Where to Place ϕ functions?

- Need to calculate the "Dominance Frontier"
- Node A *strictly dominates* node B if A dominates B and $A \neq B$.
 - Note: A does not strictly dominate B if A does not dominate B or A = B.
- The *dominance frontier* of a node B is the set of all CFG nodes y such that B dominates a predecessor of y but does not strictly dominate y
 - Intuitively: starting at B, there is a path to y, but there is another route to y that does not go through B
- Write DF[n] for the dominance frontier of node n.

Dominance Frontiers

- Example of a dominance frontier calculation results
- $DF[1] = \{1\}, DF[2] = \{1,2\}, DF[3] = \{2\}, DF[4] = \{1\}, DF[5] = \{8,0\}, DF[6] = \{8\}, DF[7] = \{7,0\}, DF[8] = \{0\}, DF[9] = \{7,0\}, DF[0] = \{\}$



Algorithm For Computing DF[n]

- Assume that doms[n] stores the dominator tree (so that doms[n] is the *immediate dominator* of n in the tree)
- Adds each B to the DF sets to which it belongs

```
for all nodes B

if \#(pred[B]) \ge 2 // (just an optimization)

for each p \in pred[B] {

runner := p // start at the predecessor of B

while (runner \neq doms[B]) // walk up the tree adding B

DF[runner] := DF[runner] U {B}

runner := doms[runner]

}
```

Insert \$\$ at Join Points

- Lift the DF[n] to a set of nodes N in the obvious way: $DF[N] = U_{n \in N} DF[n]$
- Suppose that at variable x is defined at a set of nodes N.
- $DF_0[N] = DF[N]$ $DF_{i+1}[N] = DF[DF_i[N] \cup N]$
- Let J[N] be the *least fixed point* of the sequence: $DF_0[N] \subseteq DF_1[N] \subseteq DF_2[N] \subseteq DF_3[N] \subseteq ...$
 - That is, $J[N] = DF_k[N]$ for some k such that $DF_k[N] = DF_{k+1}[N]$
- J[N] is called the "join points" for the set N
- We insert ϕ functions for the variable x at each such join point.
 - $x = \phi(x, x, ..., x)$; (one "x" argument for each predecessor of the node)
 - In practice, J[N] is never directly computed, instead you use a worklist algorithm that keeps adding nodes for $DF_k[N]$ until there are no changes.
- Intuition:
 - If N is the set of places where x is modified, then DF[N] is the places where phi nodes need to be added, but those also "count" as modifications of x, so we need to insert the phi nodes to capture those modifications too...

Example Join-point Calculation

- Suppose the variable x is modified at nodes 3 and 6
 - Where would we need to add phi nodes?
- $DF_0[\{3,6\}] = DF[\{3,6\}] = DF[3] \cup DF[6] = \{2,8\}$
- DF₁[{3,6}]
 - $= \mathsf{DF}[\mathsf{DF}_{0}\{3,6\} \cup \{3,6\}]$
 - $= DF[\{2,3,6,8\}]$
 - = DF[2] U DF[3] U DF[6] U DF[8]
 - $= \{1,2\} \cup \{2\} \cup \{8\} \cup \{0\} = \{1,2,8,0\}$
- $\mathsf{DF}_2[\{3,6\}]$
 - $= \dots$ = {1,2,8,0}
- So J[{3,6}] = {1,2,8,0} and we need to add phi nodes at those four spots.

Phi Placement Alternative

- Less efficient, but easier to understand:
- Place phi nodes "maximally" (i.e. at every node with > 2 predecessors)
- If all values flowing into phi node are the same, then eliminate it:
 %x = phi t %y, %pred1 t %y %pred2 ... t %y %predK
 // code that uses %x
 ⇒
 // code with %x replaced by %y
- Interleave with other optimizations
 - copy propagation
 - constant propagation
 - etc.



- How to place phi nodes without breaking SSA?
 - Note: the "real" implementation combines many of these steps into one pass.
 - Places phis directly at the dominance frontier
- This example also illustrates other common optimizations:
 - Load after store/alloca
 - Dead store/alloca elimination





 How to place phi nodes without breaking SSA?

Insert

- Loads at the end of each block



 How to place phi nodes without breaking SSA?

Insert

- Loads at the end of each block
- Insert stores
 after φ -nodes



- For loads after stores (LAS):
 - Substitute all uses of the load by the value being stored
 - Remove the load



- For loads after stores (LAS):
 - Substitute all uses of the load by the value being stored
 - Remove the load



- For loads after stores (LAS):
 - Substitute all uses of the load by the value being stored
 - Remove the load



- For loads after stores (LAS):
 - Substitute all uses of the load by the value being stored
 - Remove the load



- For loads after stores (LAS):
 - Substitute all uses of the load by the value being stored
 - Remove the load



- For loads after stores (LAS):
 - Substitute all uses of the load by the value being stored
 - Remove the load



- For loads after stores (LAS):
 - Substitute all uses of the load by the value being stored
 - Remove the load



- Dead Store Elimination (DSE)
 - Eliminate all stores with no subsequent loads.
- Dead Alloca Elimination (DAE)
 - Eliminate all allocas with no subsequent loads/stores.



- Dead Store Elimination (DSE)
 - Eliminate all stores with no subsequent loads.
- Dead Alloca Elimination (DAE)
 - Eliminate all allocas with no subsequent loads/stores.



- Eliminate **φ** nodes:
 - Singletons
 - With identical values from each predecessor
 - See Aycock & Horspool, 2002



- Eliminate **φ** nodes:
 - Singletons
 - With identical values from each predecessor



LLVM Phi Placement

- This transformation is also sometimes called register promotion
 - older versions of LLVM called this "mem2reg" memory to register promotion
- In practice, LLVM combines this transformation with *scalar replacement of aggregates* (SROA)
 - i.e. transforming loads/stores of structured data into loads/stores on register-sized data
- These algorithms are (one reason) why LLVM IR allows annotation of predecessor information in the .ll files
 - Simplifies computing the DF

Thursday, May 4th noon – 2:00p.m. Location: DRLB A4

FINAL EXAM

Zdancewic CIS 341: Compilers

Final Exam

- Will cover material since the midterm almost exclusively
 - Starting from Lecture 14
 - Typechecking
 - Objects, inheritance, types, implementation of dynamic dispatch
 - Basic optimizations
 - Dataflow analysis (forward vs. backward, fixpoint computations, etc.)
 - Liveness
 - Graph-coloring Register Allocation
 - Control flow analysis
 - Loops, dominator trees
- Will focus more on the theory side of things
- Format will be similar to the midterm
 - Simple answer, computation, multiple choice, etc.
 - Sample exam from last time is on the web

What have we learned? Where else is it applicable? What next?

COURSE WRAP-UP

Why CIS 341?

- You will learn:
 - Practical applications of theory
 - Parsing
 - How high-level languages are implemented in machine language
 - (A subset of) Intel x86 architecture
 - A deeper understanding of code
 - A little about programming language semantics
 - Functional programming in OCaml
 - How to manipulate complex data structures
 - How to be a better programmer
- Did we meet these goals?

Stuff we didn't Cover

- We skipped stuff at every level...
- Concrete syntax/parsing:
 - Much more to the theory of parsing...
 - Good syntax is art not science!
- Source language features:
 - Exceptions, recursive data types (easy!), advanced type systems, type inference, concurrency
- Intermediate languages:
 - Intermediate language design, bytecode, bytecode interpreters, just-intime compilation (JIT)
- Compilation:
 - Continuation-passing transformation, efficient representations, scalability
- Optimization:
 - Scientific computing, cache optimization, instruction selection/ optimization
- Runtime support:
 - memory management, garbage collections

Related Courses

- CIS 500: Software Foundations
 - Prof. Pierce
 - Theoretical course about functional programming, proving program properties, type systems, lambda calculus. Uses the theorem prover Coq.
- CIS 501: Computer Architecture
 - Prof. Devietti
 - 371++: pipelining, caches, VM, superscalar, multicore,...
- CIS 552: Advanced Programming
 - Prof. Weirich
 - Advanced functional programming in Haskell, including generic programming, metaprogramming, embedded languages, cool tricks with fancy type systems
- CIS 670: Special topics in programming languages

Where to go from here?

- Conferences (proceedings available on the web):
 - Programming Language Design and Implementation (PLDI)
 - Principles of Programming Langugaes (POPL)
 - Object Oriented Programming Systems, Languages & Applications (OOPSLA)
 - International Conference on Functional Programming (ICFP)
 - European Symposium on Programming (ESOP)

- ...

- Technologies / Open Source Projects
 - Yacc, lex, bison, flex, ...
 - LLVM low level virtual machine
 - Java virtual machine (JVM), Microsoft's Common Language Runtime (CLR)
 - Languages: OCaml, F#, Haskell, Scala, Go, Rust, ...?

Where else is this stuff applicable?

- General programming
 - In C/C++, better understanding of how the compiler works can help you generate better code.
 - Ability to read assembly output from compiler
 - Experience with functional programming can give you different ways to think about how to solve a problem
- Writing domain specific languages
 - lex/yacc very useful for little utilities
 - understanding abstract syntax and interpretation
- Understanding hardware/software interface
 - Different devices have different instruction sets, programming models

Thanks!

- To the TAs: Dmitri, Richard, J.J. and Vivek
 - for doing an amazing job putting together the projects for the course.
- To *you* for taking the class!

- How can I improve the course?
 - Feedback survey posted to Piazza (soon!)