

- Control-flow graph and basic blocks
- Data-flow analysis
- Liveness analysis
- Register allocation.

Pseudo-code

To discuss code optimization we employ a (vaguely defined) pseudo-IR called **three-address code** which uses virtual registers but does not require SSA form.

Instructions

- $x := y \# z$ where x , y and z are register names or literals and $\#$ is an arithmetic operator.
- goto L where L is a label.
- if $x \# y$ then goto L where $\#$ is a relational operator.
- $x := y$
- return x

Example code

```
s := 0
i := 1
L1: if i > n goto L2
t := i * i
s := s + t
i := i + 1
goto L1
L2: return s
```

CHALMERS

CHALMERS

Control-flow graph


Code as graph

- Each instruction is a node.
- Edge from each node to its possible **successors**.

Example code

```
s := 0
i := 1
L1: if i > n goto L2
t := i * i
s := s + t
i := i + 1
goto L1
L2: return s
```

Example as graph



CHALMERS

Static vs dynamic analysis

Dynamic analysis

If in some execution of the program ...

Dynamic properties are in general undecidable.

Compare with the halting problem:

"P halts" vs "P reaches instruction I".

Static analysis

If there is a path in the control-flow graph ...

Basis for many forms of compiler analysis – but in general we don't know if that path will ever be taken during execution.

Results are approximations – we must make sure to err on the correct side.

CHALMERS

Dataflow analysis

A static analysis

- General approach to code analysis.
- Useful for many forms of **intraprocedural optimization**:
 - Common subexpression elimination,
 - Constant propagation,
 - Dead code elimination,
 - ...
- Within a basic block, simpler methods often suffice.

CHALMERS

Example: Liveness of variables

Definitions and uses

An instruction $x := y \# z$ **defines** x and **uses** y and z .

Liveness

A variable v is **live** at a point P in the control-flow graph (CFG) if there is a path from P to a use of v along which v is not defined.

Uses of liveness information

- Register allocation: a non-live variable need not be kept in register.
- Useless-store elimination: a non-live variable need not be stored to memory.
- Detecting uninitialized variables: a local variable that is live on function entry.
- Optimizing SSA form; non-live vars don't need Φ -functions.

CHALMERS

Liveness analysis: Concepts

Def sets

The **def set** $def(n)$ of a node n is the set of variables that are defined in n (a set with 0 or 1 elements).

Use sets

The **use set** $use(n)$ of a node n is the set of variables that are used in n .

Live-out sets

The **live-out set** $live-out(n)$ of a node n is the set of variables that are live at an out-edge of n .


Live-in sets

The **live-in set** $live-in(n)$ of a node n is the set of variables that are live at an in-edge of n .

CHALMERS

An example

1st example revisited



Live-in sets

| Instr # | Set |
|---------|----------------|
| 1 | { n } |
| 2 | { n, s } |
| 3 | { i, n, s } |
| 4 | { i, n, s } |
| 5 | { i, n, s, t } |
| 6 | { i, n, s } |
| 7 | { i, n, s } |
| 8 | { s } |

How can these be computed?

CHALMERS

The dataflow equations

For every node n , we have

$$\begin{aligned} \text{live-in}(n) &= \text{use}(n) \cup (\text{live-out}(n) - \text{def}(n)) \\ \text{live-out}(n) &= \bigcup_{s \in \text{succs}(n)} \text{live-in}(s). \end{aligned}$$

where $\text{succs}(n)$ denote the set of successor nodes to n .

Computation

Let live-in , def and use be arrays indexed by nodes.

foreach node n **do** $\text{live-in}[n] = \emptyset$

repeat

foreach node n **do**

$\text{out} = \bigcup_{s \in \text{succs}(n)} \text{live-in}[s]$

$\text{live-in}[n] = \text{use}[n] \cup (\text{out} - \text{def}[n])$

until no changes in iteration.

Solving the equations

Example revisited

| Instr | def | use | succs | live-in |
|-------|-----|-------|-------|---------|
| 1 | {s} | {} | {2} | {} |
| 2 | {i} | {} | {3} | {} |
| 3 | {} | {i,n} | {4,8} | {} |
| 4 | {t} | {i} | {5} | {} |
| 5 | {s} | {s,t} | {6} | {} |
| 6 | {i} | {i} | {7} | {} |
| 7 | {} | {} | {3} | {} |
| 8 | {} | {s} | {} | {} |

Initialization done above.

live-in updated from top to bottom in each iteration (to be completed in class).

But is there a better order?

CHALMERS

Liveness: A backwards problem

Fixpoint iteration

- We iterate until no live sets change during an iteration; we have reached a **fixpoint** of the equations.
- The number of iterations (and thus the amount of work) depends on the order in which we use the equations within an iteration.
- Since liveness info propagates from successors to predecessors in the CFG, we should start with the last instruction and work backwards.
(Since the program contains a loop, this is just a heuristic).

CHALMERS

Another node order

Working from bottom to top, we get

| Instr | def | use | succs | live-in ₀ | live-in ₁ | live-in ₂ |
|-------|-----|-------|-------|----------------------|----------------------|----------------------|
| 1 | {s} | {} | {2} | {} | {n} | {n} |
| 2 | {i} | {} | {3} | {} | {n,s} | {n,s} |
| 3 | {} | {i,n} | {4,8} | {} | {i,n,s} | {i,n,s} |
| 4 | {t} | {i} | {5} | {} | {i,s} | {i,n,s} |
| 5 | {s} | {s,t} | {6} | {} | {i,s,t} | {i,n,s,t} |
| 6 | {i} | {i} | {7} | {} | {} | {i,n,s} |
| 7 | {} | {} | {3} | {} | {} | {i,n,s} |
| 8 | {} | {s} | {} | {} | {s} | {s} |

CHALMERS

Implementing data flow analysis

Data structures

- Any standard data structure for graphs will work; one should arrange for *succs* to be fast.
- For sets of variables one may use bit arrays with one bit per variable. Then union is bit-wise or, intersection bit-wise and and complement bit-wise negation.

Termination

The live sets **grow monotonically** in each iteration, so the number of iterations is bounded by N^2 , where N is the number of variables. In practice, for realistic code, the number of iterations is much smaller.

Node ordering

A heuristically good order can be found by doing a depth-first search of the CFG and reversing the node ordering.

Basic blocks

Motivations

- Control-graph with instructions as nodes become big.
- Between jumps, graph structure is trivial (**straight-line code**).


Definition

- A **basic block** starts at a labelled instruction or after a conditional jump. (First basic block starts at beginning of function).
- A basic block ends at a (conditional) jump.

We ignore code where an unlabeled statement follows an unconditional jump (such code is **unreachable**).

Example

Testing if n is prime



Notes

- Edges correspond to branches.
- Jump destinations are now blocks, not instructions.
- We may insert empty blocks.
- Analysis of control-flow graphs often done on graph with basic blocks as nodes.

Liveness analysis for CFG graphs of basic blocks

We can easily modify data flow analysis to work on control flow graphs of basic blocks.

With knowledge of *live-in* and *live-out* for basic blocks it is easy to find the set of live variables at each instruction.

How do the basic concepts need to be modified to apply to basic blocks?

Modified definitions for CFG of basic blocks

Def sets

The **def set** $def(n)$ of a node n in a CFG is the set of variables that are defined in an instruction in n .

Use sets

The **use set** $use(n)$ of a node n is the set of variables that are used in an instruction in n **before** a possible redefinition of the variable.

Live-out sets

The **live-out set** $live-out(n)$ of a node n is the set of variables that are live at an out-edge of n .

Live-in sets

The **live-in set** $live-in(n)$ of a node n is the set of variables that are live at an in-edge of n .

CHALMERS

Another dataflow problem: dominators

Definition

In a CFG, node n **dominates** node m if every path from the start node to m passes through n .

Particular case: we consider each node to dominate itself.

Concept has many uses in compilation.

Prime test CFG



Questions

- Write dataflow equations for dominance.
- How would you solve the equations?

CHALMERS

Register allocation

An important code transformation

When translating an IR with (infinitely many) virtual registers to code for a real machine, we must

- assign virtual registers to physical registers.
- write register values to memory (**spill**), at program points when the number of live virtual registers exceeds the number of available registers.

Register allocation is very important; good allocation can make a program run an order of magnitude faster (or more) as compared to poor allocation.

CHALMERS

The interference graph

Live sets and register usage

If two variables are live at the same point in the CFG, they must be in different registers.

Conversely, two variables that are never live at the same time can share a register.

Interfering variables

We say that variables x and y **interfere** if they are both in $live-in(n)$ or $live-out(n)$ for some node n .


The **interference graph** has variables as nodes and edges between interfering variables.

Our example

Since $live-out(4) = \{i, n, s, t\}$, all four variables interfere pair-wise and we need four registers for this code sequence.

An example

How many registers are needed?



Answer: Two!

Use one register for a, c and d,
the other for b, e and f.

Reformulation

To assign K registers to variables given an interference graph can be seen as colouring the nodes of the graph with K colours, with adjacent nodes getting different colours.

CHALMERS

Register allocation by graph colouring

The algorithm (K colours available)

- 1 Find a node n with less than K edges. Remove n and its edges from the graph and put on a stack.
- 2 Repeat with remaining graph until either
 - only K nodes remain or
 - all remaining nodes have at least K adjacent edges.

In the first case, give each remaining node a distinct colour and pop nodes from the stack, inserting them back into the graph with their edges and colouring them.

In the second case, we may need to **spill** a variable to memory. Optimistic algorithm: Choose one variable and push on the stack. Later, when popping the stack, we **may** be lucky and find that the neighbours use at most K-1 colours.

CHALMERS

Complexity

A hard problem

The problem to decide whether a graph can be K-coloured is NP-complete.

The simplify/select algorithm on the previous slide works well in practice; its complexity is $O(n^2)$, where n is the number of virtual registers used.

When optimistic algorithm fails, memory store and fetch instructions must be added and algorithm restarted.

Heuristics to choose variable to spill:

- Little use+def within loop;
- Interference with many variables.

CHALMERS

Move instructions

An example

```

t := s
x := s + 1
y := t + 2
...
  
```

s and t interfere,
but if t is not later redefined, they
may share a register.

Coalescing

Move instructions $t := s$ can sometimes be removed and the nodes s and t merged in the interference graph.

Conditions:

- No interference between s and t for other reasons.
- The graph must remain colourable. Safe strategies exist.

CHALMERS

Linear scan register allocation

Compilation time vs code quality

Register allocation based on graph colouring produces good code, but requires significant compilation time.

For e.g. JIT compiling, allocation time is a problem.

The Java HotSpot compiler uses a **linear scan** register allocator.

Much faster and in many cases only 10% slower code.

CHALMERS

The linear scan algorithm

Preliminaries

- Number all the instructions 1, 2, ... in some way (for now, think of numbering them from top to bottom). (Other instruction orderings improves the algorithm; also here depth first ordering is recommended.)
- Do a simplified liveness analysis, assigning a **live range** to each variable.
A live range is an interval of integers starting with the number of the instruction where the variable is first defined and ending with the number where it is last used.
- Sort live ranges in order of increasing start points into list L.

CHALMERS

The linear scan algorithm

The algorithm

- Maintain a list, called *active*, of live ranges that have been assigned registers. *active* is sorted by increasing end points and initially empty. Traverse L and for each interval I:
 - Traverse *active* and remove intervals with end points before start point of I.
 - If length of *active* is smaller than number of registers, add I to *active*; otherwise spill either I or the last element of *active*.

In the latter case, the choice of interval to spill is usually to keep interval with longest remaining range in *active*.

CHALMERS

More algorithms

Still a hot topic

Register allocation is still an active research area, an indication of its importance in practice.

Puzzle solving

Recent work by Pereira and Palsberg views register allocation as a puzzle solving problem.



Chordal graphs

Hack, Grund and Goos exploit the fact that the interference graph is **chordal** to get an $O(n^2)$ optimal algorithm. Care is needed when destructing SSA form.

CHALMERS

Next week

Last two lectures:

- Monday: Optimizations in LLVM, SSA form.
- Thursday: More code optimization (using data flow analysis).