

For previous labs, we have always introduced a new set of features to the source language that you compile. In L5, you are not adding new features, but optimizing your compiler to output more efficient assembly. We will be discussing some optimizations that you can write and challenges you may face.

Easy Optimizations First

We spend so much time talking about interesting and hard-to-implement optimizations that it's easy to get ahead of oneself and jump straight into implementing something complex like alias analysis. Don't underestimate how much simple things like peephole optimizations and improving instruction selection can impact your performance. Easy-to-implement optimizations such as eliminating self-moves or making better use of x86's parentheses syntax for memory loads could yield drastic performance improvements for relatively little work. We recommend starting L5 by reflecting on your compiler's weak suits and looking for easy improvements to make before moving on to more advanced optimizations.

Peephole Optimizations

- **Constant Folding** is a code transformation that replaces a binary operation with a constant. It works very well in tandem with constant propagation. Be careful to preserve side effects of operations though, e.g. division, modulo, shift.

Constant folding can also be applied to branches. If the condition of a branch is constant (or can be folded into a constant), the conditional jump can be replaced with an unconditional jump. This may cause entire basic blocks of code to become dead code.

- **Strength Reduction** replaces expensive operations with simpler ones. Arithmetic operations that interact with 0 or 1 can often be replaced with moves, to which constant or copy propagation can be applied.
- **Null Sequences** are sequences of operations that do not have any effects and can be replaced with a single nop. For example, self moves (moves where dest and src are statically the same) are always safe to remove since they do not change the machine state. Moving a to b and then immediately back from b to a means that the second move can be replaced with a nop and not change its behavior. Another common example is an unconditional jump to the next instruction.

Checkpoint 0

Identify opportunities for peephole optimizations in the following function.

```
1 foo:
2   t1 ← 40 + 2
3   t2 ← t1 * 1
4   if 3 < 2 then l1 else l2
5 l1:
6   t3 ← t2
7   goto l3
8 l2:
9   t4 ← t1 - 0
10  t3 ← t4
11  goto l3
12 l3:
13  t5 ← phi(t3,t4)
14  t6 ← t5 * 128
15  t5 ← t6 - t6
```

Common Subexpression Elimination

When the same operation on the same data is performed multiple times in code, common subexpression elimination replaces the latter instances of the same operation with the result of the earlier operation. This can potentially save repeating the same calculation at runtime. However, in order for this replacement to be correct, we must be sure that the same operation will yield the same result at the point of the replacement, otherwise we risk replacing subexpressions with outdated results.

Doing CSE in SSA form ensures that syntactically equal subexpressions are indeed equal in value, but we still need to make sure that the result being reused is defined along every path from the beginning of the function to the line of reuse. This can be determined by looking at the dominance relation between basic blocks.

Checkpoint 1

Identify opportunities for common subexpression elimination in the following code. Make sure to check the dominance relations.

```
1 l1:
2   c ← b
3   x ← a ⊕ b
4   if a < b then l2 else l3
5 l2:
6   y1 ← a ⊕ b
7   goto l4
8 l3:
9   y2 ← z ⊕ b
10  goto l4
11 l4:
12  y3 ← phi(y1,y2)
13  u ← z ⊕ b
```

Checkpoint 2

If line 6 above was $y1 \leftarrow a \oplus c$ instead, $a \oplus b$ would no longer be a syntactic common subexpression. They are clearly semantically equivalent though. How can we identify these semantically but not syntactically equivalent subexpressions to help CSE do better?

Checkpoint 3

One of the pairs of common subexpressions from checkpoint 1 could not have CSE applied due to the lack of a dominance relation. However, no matter what path we take through the CFG, that subexpression is computed at least once, and possibly twice. How can we move some instructions around (code motion) so that the subexpression in question only gets computed once?

Re-motivating SSA

Many optimizations become significantly simpler in SSA form. This is because in SSA, each variable is assigned exactly once, every use has a single, well-defined definition, and definitions dominate uses.

Checkpoint 4

Explain how SSA simplifies:

- (a) Constant propagation
- (b) Common subexpression elimination

(c) Dead code elimination

Conditional Constant Propagation

Many people often implement a simple constant propagation pass in their compilers in the initial labs. However, consider:

```
1 x ← 1
2 if (x < 2) then l1 else l2
3 l1:
4   y ← 5
5 l2:
6   y ← 10
7 return y
```

Standard constant propagation can't eliminate the branch since it doesn't track which branches execute, so it can't prove that the false branch is unreachable.

Conditional Constant Propagation (CCP) tracks which blocks are executable and whether variables are constant, unknown, or varying. It can also help eliminate unreachable branches.

Checkpoint 5

Apply conditional constant propagation (CCP) to the following program. Identify which blocks are executable, as well as the final simplified program after CCP.

```
1 L1:
2   x ← 1
3   y ← 2
4   if x < y then L2 else L3
5
6 L2:
7   z1 ← x + y
8   goto L4
9
10 L3:
11  z2 ← y - x
12  goto L4
13
14 L4:
15  z3 ← phi(z1, z2)
16  w ← z3 * 2
17  return w
```

Aggressive Dead Code Elimination (ADCE)

Standard dead code elimination removes instructions whose results are unused. However, it can miss opportunities because it assumes all code is reachable.

Aggressive Dead Code Elimination (ADCE) takes a more aggressive approach than standard dead code elimination, where we attempt to identify pure statements that are not used previously, where we now assume all instructions are dead unless proven otherwise, and we only keep instructions that are necessary for observable behavior.

Checkpoint 6

- What kinds of instructions must always be considered live initially?
- How does liveness propagate backwards through the program?

Checkpoint 7

Apply aggressive dead code elimination (ADCE) to the following program. Which instructions remain live, and what simplified program do we get?

```
1  L1:
2  x ← 5
3  y ← 10
4  if x < y then L2 else L3
5
6  L2:
7  a ← x + y
8  b ← a * 2
9  goto L4
10
11 L3:
12 c ← x - y
13 d ← magic(c)
14 goto L4
15
16 L4:
17 return 0
```

Optimization Interaction

As a hint, it can often be useful to run optimizations multiple times, in various orders!

Checkpoint 8

Why?