

Lecture Notes on Memory Optimizations

15-411: Compiler Design
Frank Pfenning and Jan Hoffmann

Lecture 18
March 21, 2024

1 Introduction

Even on modern architectures with hierarchical memory caches, memory access, on average, is still significantly more expensive than register access or even most arithmetic operations. Therefore, memory optimizations play a significant role in generating fast code. As we will see, whether certain memory optimizations are possible or not depends on properties of the whole language. For example, whether or not we can obtain pointers to the middle of heap-allocated objects will be a crucial question to answer.

2 Examples

We have seen many optimizations that involve temps or variables. It is natural to ask if these optimizations still apply if we replace some or all of the temps with memory locations.

Constant propagation. First, we consider constant propagation.

$$\begin{aligned}l_1 : t_1 &\leftarrow 4 \\l_2 : t_2 &\leftarrow 8128 \\l_3 : x &\leftarrow t_1\end{aligned}$$

In this program we can replace t_1 with 4 in line l_3 . Of course, we have to make sure that the definition in line l_1 is the only one that reaches line l_3 . However, this is always the case when the program is in SSA form.

Can we still perform the optimization if we replace t_1 and t_2 with memory locations?

$$\begin{aligned} l_1 : M[q] &\leftarrow 4 \\ l_2 : M[p] &\leftarrow 8128 \\ l_3 : x &\leftarrow M[q] \end{aligned}$$

Before replacing line l_3 with $x \leftarrow 4$, we have to ensure is that the temp q has the same value in lines l_1 and l_3 . Again, this can be ensured by SSA form. However, even then it's not sound to perform the optimization. We also need to ensure that $M[q]$ is not modified in any other line on a path that reaches l_3 . For example, we have to ensure that $p \neq q$. If $q = p$ then x would have the value 8128 in the original program.

If we assume that this code snippet appears in a basic block (no jumps) then we can perform the optimization if p and q do not *alias*, that is, $p \neq q$. Alias analysis is a central problem of optimization of memory locations.

Common Subexpression Elimination Next, we consider the equivalent of common subexpression elimination. Consider a function, $\text{mult}(A, p, q)$ that multiplies matrix A with vector p and returns the result in vector q .

```
struct point {
    int x;
    int y;
};
typedef struct point pt;

void mult(int[] A, pt* p, pt* q) {
    q->x = A[0] * p->x + A[1] * p->y;
    q->y = A[2] * p->x + A[3] * p->y;
    return;
}
```

Below is the translation into abstract assembly, with the small twist that we have allowed memory reference to be of the form $M[\text{base} + \text{offset}]$. The memory optimization question we investigate is whether some load instructions $t \leftarrow M[s]$ can

be avoided because the corresponding value is already held in a temp.

```

mult(A, p, q) :
    t0 ← M[A + 0]
    t1 ← M[p + 0]
    t2 ← t0 * t1
    t3 ← M[A + 4]
    t4 ← M[p + 4]
    t5 ← t3 * t4
    t6 ← t2 + t5
    M[q + 0] ← t6
    t8 ← M[A + 8]
    t9 ← M[p + 0]    # redundant load?
    t10 ← t8 * t9
    t11 ← M[A + 12]
    t12 ← M[p + 4]   # redundant load?
    t13 ← t11 * t12
    t14 ← t10 + t13
    M[q + 4] ← t14
    return

```

We see that the source refers to $p \rightarrow x$ and $p \rightarrow y$ twice, and those are reflected in the two, potentially redundant loads above. Before you read on, consider if we could replace the lines with $t_9 \leftarrow t_1$ and $t_{12} \leftarrow t_4$. We can do that if we can be assured that memory at the addresses $p + 0$ and $p + 4$, respectively, has not changed since the previous load instructions.

It turns out that in C0 the second load is definitely redundant, but the first one may not be.

The first load is not redundant because when this function is called, the pointers p and q might be the same (they might *aliased*). When this is the case, the store to $M[q+0]$ will likely change the value stored at $M[p+0]$, leading to a different answer than expected for the second line.

On the other hand, this cannot happen for the first line, because $M[q + 0]$ could never be the same as $M[p + 4]$ since one accesses the x field and the other the y field of a struct.

Of course, the answer is most likely wrong when $p = q$. One could either rewrite the code, or require that $p \neq q$ in the precondition to the function.

In C, the question is more delicate because the use of the address-of (&) operator could obtain pointers to the middle of objects. For example, the argument `int [] A` would be `int* A` in C, and such a pointer might have been obtained with `&q->x`.

3 Using the Results of Alias Analysis

In C0, the types of pointers are a powerful basis of alias analysis. The way alias analysis is usually phrased is as a *may-alias* analysis, because we try to infer which pointers in a program may alias. Then we know for optimization purposes that if two pointers are not in the *may-alias* relationship that they must be different. Writing to one address cannot change the value stored at the other.

Let's consider how we might use the results of alias analysis, embodied in a predicate $\text{may-alias}(a, b)$ for two temps (containing addresses) a and b . We assume we have a load instruction

$$l : t \leftarrow M[a]$$

and we want to infer if this is *available* at some other line $l' : t' \leftarrow M[a]$ so we could replace it with $l' : t' \leftarrow t$. Our optimization rule can be defined as follows.

$$\left. \begin{array}{l} l : t \leftarrow M[a] \\ \dots \\ k : t' \leftarrow M[a] \end{array} \right\} \longrightarrow \left\{ \begin{array}{l} l : t \leftarrow M[a] \\ \dots \\ k : t' \leftarrow t \end{array} \right. \quad \text{provided } l > k, \text{avail}(l, k)$$

The fact that l dominates k is sufficient here in SSA form to guarantee that the meaning of t and a remains unchanged. The predicate avail is supposed to check that $M[a]$ also remains unchanged and $l > k$ ensures that t is defined at line k .

It turns out to be simpler to first define *unavailability*, the negation of availability. We define a predicate $\text{unavail}(l, k)$ $l : t \leftarrow M[a]$ at other instructions k that are dominated by l . Domination is not really required in the rules. We only add for efficiency reasons. We need the predicate $\text{avail}(l, k)$ only for pairs of locations l, k such that $l > k$. So we only need to derive unavailability for such pairs.

For unavailability, $\text{unavail}(l, k)$, we have the seeding rule on the left and the general propagation rule on the right. Because we are in SSA, we know in the seeding rule that $l > k$ where k is the (unique) successor of l' .

$$\frac{\begin{array}{l} l : t \leftarrow M[a] \\ l > l' \\ l' : M[b] \leftarrow s \\ \text{may-alias}(a, b) \\ \text{succ}(l', k) \end{array}}{\text{unavail}(l, k)} \qquad \frac{\begin{array}{l} \text{unavail}(l, k) \\ \text{succ}(k, k') \\ l > k' \end{array}}{\text{unavail}(l, k')}$$

The rule on the right includes the cases of jumps or conditional jumps. This ensures that in a node with multiple predecessors, if a value is unavailable in just one of them, it will be unavailable at the node. Function calls can also seed unavailability. Unfortunately it is enough if one of the function parameters is a memory reference, because from one memory reference we may be able to get to another by following

pointers and offsets.

$$\frac{\begin{array}{l} l : t \leftarrow M[a] \\ l > l' \\ l' : d \leftarrow f(s_1, \dots, s_n) \\ \text{memref}(s_i) \\ \text{succ}(l', k) \end{array}}{\text{unavail}(l, k)}$$

With more information on the shape of memory this rule can be relaxed.

From unavailability we can deduce which memory values are still available, namely those that are not unavailable (restriction attention to those that are dominated by the load—otherwise the question is not asked).

$$\frac{\begin{array}{l} l : t \leftarrow M[a] \\ l > l' \\ \neg \text{unavail}(l, l') \end{array}}{\text{avail}(l, l')}$$

Note that stratification is required: we need to saturate $\text{unavail}(l, l')$ before applying this rule.

4 Type-Based Alias Analysis

One way of performing type-based alias analysis is to propagate type information from the semantic analysis to abstract assembly. The type information is used to associate with each memory location the type of data that is stored at that location. The idea is that data that is stored at locations of different types is not aliased. Such an analysis works for type-safe languages such as Java and C0.

Here, we will restore the type information directly at the assembly level. Our alias analysis is based on the type and offset of the address. We call this an *alias class*, with the idea that pointers in different alias classes cannot alias. More formally, we derive a predicate $\text{class}(a, \tau, \text{offset})$, which expresses that a is temp containing an address derived from a source of type τ and offset offset from the start of the memory of type τ .

Then the *may-alias* relation is defined by

$$\frac{\text{class}(a, \tau, k) \quad \text{class}(b, \tau, k)}{\text{may-alias}(a, b)}$$

There is a couple of special cases we do not treat explicitly. For example, the location of the array length (which is stored in safe mode at least) may be at offset -8 . But such a location can never be written to (array lengths never change, once

allocated), so a load of the array length is available at all locations dominated by the load.

Since some temps can store multiple addresses because they are assigned in a loop or as a result of multiple jumps the a parameterized label. In this case, we are not able to assign a single offset and introduce \top , pronounced “top”, to represent *all addresses*. So you can think of it as being \mathbb{N} . This leads to the following two rules.

$$\frac{\text{class}(a, \tau, k) \quad \text{class}(b, \tau, \top)}{\text{may-alias}(a, b)} \qquad \frac{\text{class}(a, \tau, \top) \quad \text{class}(b, \tau, \top)}{\text{may-alias}(a, b)}$$

The seed of the class relation comes from function types and need to be propagated by the compiler. In our example,

```

mult(A, p, q) :
  t0 ← M[A + 0]
  t1 ← M[p + 0]
  t2 ← t0 + t1
  t3 ← M[A + 4]
  ...

```

the compiler would generate

```

class(A, int[], 0)
class(p, struct point*, 0)
class(q, struct point*, 0)

```

We now propagate the information through a forward dataflow analysis. For example:

$$\frac{l : b \leftarrow a \quad \text{class}(a, \tau, k)}{\text{class}(b, \tau, k)} \qquad \frac{l : b \leftarrow a + \$n \quad \text{class}(a, \tau, k)}{\text{class}(b, \tau, k + n)}$$


In the second case we have written $\$n$ to emphasize the second summand is a constant n . Unfortunately, if it is a variable, we cannot precisely calculate the offset. This may happen with arrays, but not with pointers, including pointers to structs. So we need to generalize the third argument to class to be either a variable or \top , which indicates any value may be possible. We then have, for example

$$\frac{l : b \leftarrow a + t \quad \text{class}(a, \tau, k)}{\text{class}(b, \tau, \top)}$$

Now \top behaves like an information sink. For example, $\top + k = k + \top = \top$. Since in SSA form a is defined only once, we should not have to change our mind about the class assigned to a variable. However, at parameterized jump targets (which is equivalent to Φ -functions), we need to “disjoin” the information so that

if the argument is known to be k at one predecessor but unknown at \top at another predecessor, the result should be \top .

Because of loops, we then need to generalize further and introduce \perp which means that we believe (for now) that the variable is never used. Because of the seeding by the compiler, this will mostly happen for loop variables. The values are arranged in a *lattice*



where at the bottom we have more information, at the top the least. The \sqcup operation between lattice elements finds the least upper bounds of its two arguments. For example, $0 \sqcup 4 = \top$ and $\perp \sqcup 2 = 2$. We use it in SSA form to combine information about offsets. We now read an assertion $\text{class}(a, \tau, k)$ as saying that the offset is at least k under the lattice ordering. Then we have

$$\begin{array}{l} \text{lab}(a_1) : \\ \text{class}(a_1, \tau, k_1) \\ l : \text{goto lab}(a_2) \\ \text{class}(a_2, \tau, k_2) \\ \hline \text{class}(a_1, \tau, k_1 \sqcup k_2) \end{array}$$

Written with Φ -functions, we would have

$$\begin{array}{l} \text{class}(a_0, \tau, k_0) \\ a_0 \leftarrow \Phi(a_1, \dots, a_n) \\ \text{class}(a_i, \tau, k_i) \quad (1 \leq k \leq n) \\ \hline \text{class}(a_0, \tau, k_0 \sqcup k_1 \sqcup \dots \sqcup k_n) \end{array}$$

Because of loops we might perform this calculation multiple times until we have reached a fixed point. In this case the fixed point is least upper bound of all the offset classes we compute, which is a little different than the saturated data base we considered before.

It is not immediately clear why the different classes of the addresses a_i in the last two rules have the same type τ . For example, it could be possible that we have $\text{class}(a_1, \tau_1, k_1)$ and $\text{class}(a_2, \tau_2, k_2)$ for $\tau_1 \neq \tau_2$ in the first rule. However, this should not be possible if the programs are translated from C0. It would mean that an expression can have two different pointer types in the code.

This is an example of *abstract interpretation*, which may be a subject of a future lecture. One can obtain a more precise alias analysis if one refines the *abstract domain*, which is lattice shown above.

5 Allocation-Based Alias Analysis

Another technique to infer that pointers may not alias is based on their allocation point. In brief, if two pointers are allocated with different calls to `alloc` or `alloc_array`, then they cannot be aliased. Because allocation may happen in a different function than we are currently compiling (and hopefully optimizing), this is an example of an *interprocedural analysis*.