

Lecture Notes on Memory Optimizations

15-411: Compiler Design
Frank Pfenning

Lecture 18
October 23, 2014

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 A Simple Example

We will use a simple running example to illustrate memory optimization and their conditions of applicability. In this example, $\text{mult}(A, p, q)$ will multiply matrix A with vector p and return 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 $M[base + offset]$. 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 mostly 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 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 (in the notation of linear inference from [Lecture 16](#)):

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

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

Reaching analysis for memory references is a simple forward dataflow analysis. If we have a node with two or more incoming control flow edges, it must be available along all of them. For the purposes of traversing loops we assume availability, essentially trying to find a counterexample in the loop. To express this concisely,

our analysis rules propagate *unavailability* of a definition $l : t \leftarrow M[a]$ an other instructions k that are dominated by l .

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' .

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

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.

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

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).

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

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

4 Type-Based Alias Analysis

The simplest form of 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. The basic predicate here is $\text{class}(a, \tau, \text{offset})$ which expresses that a is 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.

The seed of the class relation comes from the compiler, that annotates an address with this information. 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

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

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.

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*.