

1 IR lowering

As described earlier, we express the IR lowering transformation using two syntax-directed translation functions, both named \mathcal{L} (for “lowering”).

$\mathcal{L}[[s]]$ translates an IR statement s to a sequence $s_1; \dots; s_n$ of canonical IR statements that have the same effect. We write $\mathcal{L}[[s]] = s_1; \dots; s_n$, or as a shorthand, $\mathcal{L}[[s]] = \vec{s}$.

$\mathcal{L}[[e]]$ translates an IR expression e to a sequence of canonical IR statements \vec{s} that have the same effect, and an expression e' that has the same value if evaluated after the whole sequence of statements \vec{s} . We write $\mathcal{L}[[e]] = \vec{s}; e'$ to denote this.

2 Lowering statements

We translate a sequence of statements by lowering each statement, and concatenating the results:

$$\mathcal{L}[[\mathbf{SEQ}(s_1, \dots, s_n)]] = \mathcal{L}[[s_1]]; \dots; \mathcal{L}[[s_n]]$$

To lower an **EXP** node, we just throw away the expression. We write the translation rule as an inference rule:

$$\frac{\mathcal{L}[[e]] = \vec{s}; e'}{\mathcal{L}[[\mathbf{EXP}(e)]] = \vec{s}}$$

For statements such as **JUMP** and **CJUMP**, we flatten the expression to obtain a sequence of statements that are done before the jump:

$$\frac{\mathcal{L}[[e]] = \vec{s}; e'}{\mathcal{L}[[\mathbf{JUMP}(e)]] = \vec{s}; \mathbf{JUMP}(e')} \quad \frac{\mathcal{L}[[e]] = \vec{s}; e'}{\mathcal{L}[[\mathbf{CJUMP}(e, l_1, l_2)]] = \vec{s}; \mathbf{CJUMP}(e', l_1, l_2)}$$

Some statements we can leave alone:

$$\mathcal{L}[[\mathbf{LABEL}(l)]] = \mathbf{LABEL}(l)$$

A tricky case is **MOVE**. We'd like the following translation, but it's not safe in general:

$$\frac{\mathcal{L}[[e]] = \vec{s}; e' \quad \mathcal{L}[[dest]] = \vec{s}'; dest'}{\mathcal{L}[[\mathbf{MOVE}(dest, e)]] = \vec{s}; \vec{s}'; \mathbf{MOVE}(dest', e')}$$

The problem is that if the statements \vec{s}' change the value of e' , the meaning of the translated program is not the same as the original. We'll say that e' *commutes* with s' if s' can be evaluated either before or after e' with the same effect and the same value resulting from e' . The **MOVE** rule above can only be used if s' and e' are guaranteed to commute.

What if they don't commute? In that case we can capture the value of e' in a fresh temporary. This is not as good a translation in general, because the simpler **MOVE** node at the end may eliminate some opportunities to generate efficient code. As we'll see later when we do instruction selection, big expression trees are helpful for generating efficient assembly code.

$$\frac{\mathcal{L}[[e]] = \vec{s}; e' \quad \mathcal{L}[[dest]] = \vec{s}'; dest' \quad e' \text{ and } s' \text{ do not commute}}{\mathcal{L}[[\mathbf{MOVE}(dest, e)]] = \vec{s}; \mathbf{MOVE}(\mathbf{TEMP}(t), e'); \vec{s}'; \mathbf{MOVE}(dest', t)}$$

Deferring the question of how to decide when expressions and statements commute, let's look at how to lower IR expressions.

3 Lowering expressions

Recall that our goal is to convert an expression into one that has no side effects, the side effects being factored out into a sequence of statements. If we use \bullet to represent an empty sequence of statements, expressions that already have no side effects are trivial to lower:

$$\frac{e = \mathbf{CONST}(i) \vee e = \mathbf{NAME}(l) \vee e = \mathbf{TEMP}(t)}{\mathcal{L}[[e]] = \bullet; e}$$

For other simple expressions, we just hoist the statements out of subexpressions:

$$\frac{\mathcal{L}[[e]] = \vec{s}; e'}{\mathcal{L}[[\mathbf{MEM}(e)]] = \vec{s}; \mathbf{MEM}(e')}$$

$$\frac{\mathcal{L}[[e]] = \vec{s}; e'}{\mathcal{L}[[\mathbf{JUMP}(e)]] = \vec{s}; \mathbf{JUMP}(e')}$$

$$\frac{\mathcal{L}[[e]] = \vec{s}; e'}{\mathcal{L}[[\mathbf{CJUMP}(e, l_1, l_2)]] = \vec{s}; \mathbf{CJUMP}(e', l_1, l_2)}$$

Since we can hoist statements, we can eliminate **ESEQ** nodes completely:

$$\frac{\mathcal{L}[[s]] = \vec{s} \quad \mathcal{L}[[e]] = \vec{s}'; e'}{\mathcal{L}[[\mathbf{ESEQ}(s, e)]] = \vec{s}; \vec{s}'; e'}$$

Binary operations create the same commutation problem we saw with **MOVE**:

$$\frac{\mathcal{L}[[e_1]] = \vec{s}_1; e'_1 \quad \mathcal{L}[[e_2]] = \vec{s}_2; e'_2 \quad s_2 \text{ and } e'_1 \text{ commute}}{\mathcal{L}[[\mathbf{OP}(e_1, e_2)]] = \vec{s}_1; \vec{s}_2; \mathbf{OP}(e'_1, e'_2)}$$

Call nodes must be hoisted too because they can cause side effects. And we have to prevent the side effects of computing arguments from changing the computation of other arguments.

$$\frac{\mathcal{L}[[e_i]] = \vec{s}_i; e'_i \quad \forall i \in 0..n}{\mathcal{L}[[\mathbf{CALL}(e_0, e_1, \dots, e_n)]] = \vec{s}_0; \mathbf{MOVE}(\mathbf{TEMP}(t_0), e'_0); \vec{s}_1; \mathbf{MOVE}(\mathbf{TEMP}(t_1), e'_1); \dots; \vec{s}_n; \mathbf{MOVE}(\mathbf{TEMP}(t_n), e'_n); \mathbf{MOVE}(\mathbf{TEMP}(t), \mathbf{CALL}(t_0, t_1, \dots, t_n)); \mathbf{TEMP}(t)}$$

4 Commuting statements and expressions

The rules for **OP** and **MOVE** both rely on interchanging the order of a statement s and an expression e . This can be done safely when the statement cannot alter the value of the expression. There are two ways in which this could happen: the statement could change the value of a temporary variable used by the expression, and the statement could change the value of a memory location used by the expression. It is easy to determine whether the statement updates a temporary used by the expression, because temporaries have unique names. Memory is harder because two memory location can be *aliases*. If the statement s

uses a memory location $\text{MEM}(e_1)$ as a destination, and the expression reads from the memory location $\text{MEM}(e_2)$, and e_1 might have the same value as e_2 at run time, we cannot safely interchange the operations.

A simple, conservative approach is to assume that all MEM nodes are potentially aliases, so a statement and an expression that both use memory are never reordered. We can do better by exploiting the observation that two nodes of the form $\text{MEM}(\text{TEMP}(t) + \text{CONST}(k_1))$ and $\text{MEM}(\text{TEMP}(t) + \text{CONST}(k_2))$ cannot be aliases if $k_1 \neq k_2$ ¹.

To do a good job of reordering memory accesses, we want more help from analysis of the program at source level. Otherwise there is not enough information available at the IR level for alias analysis to be tractable.

A simple observation we can exploit is that if the source language is strongly typed, two MEM nodes with different types cannot be aliases. If we keep around source-level type information for each MEM node, which we might denote as $\text{MEM}_t(e)$, then these type annotations t can help identify opportunities to reorder operations. Accesses to $\text{MEM}_t(e_1)$ and to $\text{MEM}_{t'}(e_2)$ cannot conflict if t is not compatible with t' .

Sophisticated compilers incorporate some form of *pointer analysis* to determine which memory locations might be aliases. The typical output of pointer analysis is a label for each distinct MEM node, such that accesses to differently labeled MEM nodes cannot interfere with each other. The label could be as simple as an integer index, where MEM_i and MEM_j cannot be aliases unless $i = j$. We'll see how to do such a pointer analysis in a later lecture.

5 Basic blocks and traces

See the slides for material on basic blocks and traces.

¹We also need to know that the language run-time system will never map two virtual addresses to the same physical address.