

Code optimizations

Advanced Compiler Construction
Michel Schinz – 2016-03-24

Optimizations

The goal of optimizations is to rewrite the program being compiled to a new program that is simultaneously:

- behaviorally equivalent to the original one,
- better in some respect – e.g. faster, smaller, more energy-efficient, etc.

Optimizations can be broadly split in two classes:

- **machine-independent optimizations**, like constant folding or dead code elimination, are high-level and do not depend on the target architecture,
- **machine-dependent optimizations**, like register allocation or instruction scheduling, are low-level and depend on the target architecture.

Rewriting optimizations

In this lesson, we will examine a simple set of machine-independent rewriting optimizations.

Most of them are relatively simple rewrite rules that transform the source program to a shorter, equivalent one.

IRs and optimizations

The importance of IRs

The intermediate representation (IR) on which rewriting optimizations are performed can have a dramatic impact on their ease of implementation.

A rewriting optimization generally works in two steps:

1. the program is analyzed to find rewrite opportunities,
2. the program is rewritten based on the opportunities identified in the first step.

A good IR should make both steps as easy as possible. The following examples illustrate the importance of using the right IR.

Case 1: constant propagation

To illustrate the impact of IR on the analysis step, consider the following program fragment in some imaginary IR:

```
x ← 7
```

...

Is it legal to perform constant propagation and blindly replace all later occurrences of x by 7? It depends on the IR:

- If multiple assignments to the same variable are allowed, then additional (data-flow) analyses are required to answer the question, as x might be re-assigned later.
- However, if multiple assignments to the same variable are prohibited, then yes, all occurrences of x can be unconditionally replaced by 7!

Other simple optimizations

Apart from constant propagation, many simple optimizations are made hard by the presence of multiple assignments to a single variable:

- **common-subexpression elimination**, which consists in avoiding the repeated evaluation of expressions,
- (simple) **dead code elimination**, which consists in removing assignments to variables whose value is not used later,
- etc.

In all cases, analyses are required to distinguish the various “versions” of a variable that appear in the program.

Conclusion: a good IR should not allow multiple assignments to a variable!

Case 2: inlining


Inlining (or **in-line expansion**) consists in replacing a call to a function by a copy of the body of that function, with parameters replaced by the actual arguments. It is a very important compiler optimization, as it often opens the door to further optimizations.

Some aspects of the intermediate representation can have an important impact on the implementation of inlining. To illustrate this, let us examine some problems that can occur when performing inlining directly on the AST – a choice that might seem reasonable at first sight.

Naïve inlining: problem #1

```
(def print/ret (fun (x) (int-print x) x))  
(def twice (fun (y) (+ y y)))  
(def f (fun (z) (twice (print/ret z))))
```

incorrect inlining
of twice in f



```
(def f (fun (z)  
          (+ (print/ret z)  
            (print/ret z))))
```


z is
printed
twice!

Possible solution: bind actual parameters to variables (using a `let`) to ensure that they are evaluated *at most* once.

Naïve inlining: problem #2

```
(def first (fun (x y) x))  
(def print/ret  
  (fun (z) (first z (int-print z))))
```

incorrect inlining of
first in print/ret



```
(def print/ret (fun (z) z))
```

z
isn't printed
at all!

Possible solution: bind actual parameters to variables (using a `let`) to ensure that they are evaluated *at least* once.

Easy inlining

The two pitfalls presented earlier can be avoided by bindings actual arguments to variables (using a `let`) before using them in the body of the inlined function.

However, a properly-designed IR can also avoid the problems altogether by ensuring that actual parameters are always atoms, i.e. variables or constants.

Conclusion: a good IR should only allow atomic arguments to functions.

IR comparison

Using the two very basic criteria we identified, we can evaluate the various classes of IRs we have seen:

- standard RTL/CFG is bad in that its variables are mutable; however, it allows only atoms as function arguments, which is good,
- RTL/CFG in SSA form, CPS/L₃ and similar functional IRs are good in that their variables are immutable, and they only allow atoms as function arguments.

Simple CPS/L₃ optimizations

Rewriting optimizations

The rewriting optimizations for CPS/L₃ are specified as a set of rewriting rules of the form $T \rightsquigarrow_{\text{opt}} T'$. These rules rewrite a CPS/L₃ term T to an equivalent – but hopefully more efficient – term T' .

Optimization contexts

Optimization rewriting rules cannot be applied anywhere, but only in specific locations. For example, it would be incorrect to try to rewrite the parameter list of a function.

This constraint can be captured by specifying all the **contexts** in which it is valid to perform a rewrite, where a context is a term with a single **hole** denoted by \square .

The hole of a context C can be plugged with a term T , an operation written as $C[T]$.

For example, if C is $(\text{if } \square \text{ ct } \text{cf})$, then $C[(< x y)]$ is $(\text{if } (< x y) \text{ ct } \text{cf})$.

Optimization contexts

The optimization contexts for CPS/L₃ are generated by the following grammar:

$C_{\text{opt}} ::= \square$

| $(\text{let}_l ((n\ l))\ C_{\text{opt}})$

| $(\text{let}_p ((n\ (p\ n_1\ \dots)))\ C_{\text{opt}})$

| $(\text{let}_c ((c_1\ e_1)\ \dots\ (c_i\ (\text{cnt}\ (n_{i,1}\ \dots))\ C_{\text{opt}}))\ \dots\ (c_k\ e_k))\ e)$

| $(\text{let}_c ((c_1\ e_1)\ \dots)\ C_{\text{opt}})$

| $(\text{let}_f ((f_1\ e_1)\ \dots\ (f_i\ (\text{fun}\ (n_{i,1}\ \dots))\ C_{\text{opt}}))\ \dots\ (f_k\ e_k))\ e)$

| $(\text{let}_f ((f_1\ e_1)\ \dots)\ C_{\text{opt}})$

Optimization relation

Using the optimization rewriting rules – presented later – and the optimization contexts, it is possible to specify the optimization relation \Rightarrow_{opt} that rewrites a term to an optimized version:

$$C_{\text{opt}}[T] \Rightarrow_{\text{opt}} C_{\text{opt}}[T'] \text{ where } T \rightsquigarrow_{\text{opt}} T'$$

(Non-)shrinking rules

We can distinguish two classes of rewriting rules:

1. **shrinking rules** rewrite a term to an equivalent but smaller one,
2. **non-shrinking rules** rewrite a term to an equivalent but potentially larger one.

Shrinking rules can be applied at will, possibly until the term is fully reduced, while non-shrinking rules cannot, as they could result in infinite expansion. Heuristics must be used to decide when to apply non-shrinking rules.

Except for (non-linear) inlining, all optimizations we will see are shrinking.

Dead code elimination

Dead code elimination removes all code that neither performs side effects nor produces a value used later.

$(\text{let}_l ((n l)) e)$

$\rightsquigarrow_{\text{opt}} e$ [if n is not free in e]

$(\text{let}_p ((n (p n_1 \dots))) e)$

$\rightsquigarrow_{\text{opt}} e$ [if n is not free in e and

$p \notin \{\text{byte-read, byte-write, block-set!}\}$]

$(\text{let}_f ((n_1 f_1) \dots (n_i f_i) \dots (n_k f_k)) e)$

$\rightsquigarrow_{\text{opt}} (\text{let}_f ((n_1 f_1) \dots (n_k f_k)) e)$

[if n_i is not free in $\{f_1, \dots, f_{i-1}, f_{i+1}, \dots, f_k, e\}$]

The rule for continuations is similar to the one for functions.

Dead code elimination

The rewriting rules to eliminate dead functions does not eliminate dead but mutually-recursive functions.

To handle them, one must start from the main expression of the program, and identify all functions transitively reachable from it. All unreachable functions are dead.

The rule for continuations has the same problem, which admits a similar solution. The only difference is that, continuations being local to functions, the reachability analysis can start in function bodies.

CSE

Common subexpression elimination (CSE) avoids the repeated evaluation of a single expression.

$$(\text{let}_l ((n_1 l)) C_{\text{opt}}[(\text{let}_l ((n_2 l)) e])])$$
$$\rightsquigarrow_{\text{opt}} (\text{let}_l ((n_1 l)) C_{\text{opt}}[e\{n_2 \rightarrow n_1\}])$$
$$(\text{let}_p ((n_1 (+ a_1 a_2)))$$
$$C_{\text{opt}}[(\text{let}_p ((n_2 (+ a_1 a_2))) e])])$$
$$\rightsquigarrow_{\text{opt}} (\text{let}_p ((n_1 (+ a_1 a_2))) C_{\text{opt}}[e\{n_2 \rightarrow n_1\}])$$

etc.

CSE

These simple rules for common subexpression elimination are also a bit too naïve and miss some opportunities because of scoping.

For example, the common subexpression `(+ y z)` below is not optimized by them:

```
(letc ((c1 (cnt ()
              (letp ((x1 (+ y z)))
                    ...)))
        (c2 (cnt ()
              (letp ((x2 (+ y z)))
                    ...))))
  ...)
```

η -reduction

Variants of the standard η -reduction can be performed to remove redundant definitions.

$$\begin{aligned} &(\text{let}_c \ ((c_1 e_1) \dots \\ &\quad (c_i (\text{cnt } (a_1 \dots) (\text{app}_c d a_1 \dots))) \dots \\ &\quad (c_k e_k)) \end{aligned}$$

e)

$$\rightsquigarrow_{\text{opt}} (\text{let}_c \ ((c_1 e_{1\{c_i \rightarrow d\}}) \dots (c_k e_{k\{c_i \rightarrow d\}})) e_{\{c_i \rightarrow d\}})$$
$$\begin{aligned} &(\text{let}_f \ ((n_1 f_1) \dots \\ &\quad (n_i (\text{fun } (c a_1 \dots) (\text{app}_f g c a_1 \dots))) \dots \\ &\quad (n_k f_k)) \end{aligned}$$

e)

$$\rightsquigarrow_{\text{opt}} (\text{let}_f \ ((n_1 f_{1\{n_i \rightarrow g\}}) \dots (n_k f_{k\{n_i \rightarrow g\}})) e_{\{n_i \rightarrow g\}})$$

Constant folding (1)

Constant folding replaces a constant expression by its value.

Example for addition:

$(\text{let}_l ((n_1 l_1))$

$C_{\text{opt}}[(\text{let}_l ((n_2 l_2))$

$C'_{\text{opt}}[(\text{let}_p ((n_3 (+ n_1 n_2))) e)])]$

$\rightsquigarrow_{\text{opt}} (\text{let}_l ((n_1 l_1))$

$C_{\text{opt}}[(\text{let}_l ((n_2 l_2))$

$C'_{\text{opt}}[(\text{let}_l ((n_3 l_1+l_2)) e)])]$

Similar rules exist for other arithmetic primitives.

Constant folding (2)

Primitives appearing in conditional expressions are also amendable to constant folding, for example:

$(\text{if } (= n n) c_t c_f)$

$\rightsquigarrow_{\text{opt}} (\text{app}_c c_t)$

$(\text{if } (\neq n n) c_t c_f)$

$\rightsquigarrow_{\text{opt}} (\text{app}_c c_f)$

$(\text{let}_1 ((n_1 l_1))$

$C_{\text{opt}}[(\text{let}_1 ((n_2 l_2))$

$C'_{\text{opt}}[(\text{if } (= n_1 n_2) c_t c_f)])])$

$\rightsquigarrow_{\text{opt}} (\text{let}_1 ((n_1 l_1))$

$C_{\text{opt}}[(\text{let}_1 ((n_2 l_2))$

$C'_{\text{opt}}[(\text{app}_c c_t)])])$ *[if $l_1 = l_2$]*

etc.

Neutral/absorbing elements

Uses of neutral and absorbing elements of arithmetic primitives can be simplified. For multiplication, this is expressed by the following rules:

$$\begin{aligned} & (\text{let}_1 ((n_1 \ 0)) \ C_{\text{opt}}[(\text{let}_p ((n \ (* \ n_1 \ n_2))) \ e)]) \\ & \rightsquigarrow_{\text{opt}} (\text{let}_1 ((n_1 \ 0)) \ C_{\text{opt}}[e\{n \rightarrow n_1\}]) \end{aligned}$$

$$\begin{aligned} & (\text{let}_1 ((n_2 \ 0)) \ C_{\text{opt}}[(\text{let}_p ((n \ (* \ n_1 \ n_2))) \ e)]) \\ & \rightsquigarrow_{\text{opt}} (\text{let}_1 ((n_2 \ 0)) \ C_{\text{opt}}[e\{n \rightarrow n_2\}]) \end{aligned}$$

$$\begin{aligned} & (\text{let}_1 ((n_1 \ 1)) \ C_{\text{opt}}[(\text{let}_p ((n \ (* \ n_1 \ n_2))) \ e)]) \\ & \rightsquigarrow_{\text{opt}} (\text{let}_1 ((n_1 \ 1)) \ C_{\text{opt}}[e\{n \rightarrow n_2\}]) \end{aligned}$$

$$\begin{aligned} & (\text{let}_1 ((n_2 \ 1)) \ C_{\text{opt}}[(\text{let}_p ((n \ (* \ n_1 \ n_2))) \ e)]) \\ & \rightsquigarrow_{\text{opt}} (\text{let}_1 ((n_2 \ 1)) \ C_{\text{opt}}[e\{n \rightarrow n_1\}]) \end{aligned}$$

Similar rules exist for other arithmetic operators.

Block primitives

Block primitives are harder to optimize, because block elements can be modified.

However, some blocks used by the compiler, e.g. to implement closures, are known to be constant once initialized. This makes rewritings like the following possible:

$$\begin{aligned} & (\text{let}_p ((b (\text{block-alloc-} k \text{ s}))) \\ & \quad C_{\text{opt}}[(\text{let}_p ((t (\text{block-set! } b \text{ } i \text{ } v))) \\ & \quad \quad C'_{\text{opt}}[(\text{let}_p ((n (\text{block-get } b \text{ } i))) e])])]) \\ \rightsquigarrow_{\text{opt}} & (\text{let}_p ((b (\text{block-alloc-} k \text{ s}))) \\ & \quad C_{\text{opt}}[(\text{let}_p ((t (\text{block-set! } b \text{ } i \text{ } v))) \\ & \quad \quad C'_{\text{opt}}[e\{n \rightarrow v\}])]) \end{aligned}$$

[when tag k identifies a block that is not modified after initialization, e.g. a closure block]

Exercise

CPS/L₃ contains the following block primitives:

- `block-alloc-n` size
- `block-tag` block
- `block-size` block
- `block-get` block index
- `block-set!` block index value

Informally describe three rewriting optimizations that could be performed on these primitives, and in which conditions they could be performed.

CPS/L₃ inlining

Shrinking inlining

A non-recursive continuation or function that is applied exactly once – i.e. used in a linear fashion – can always be inlined without making the code grow:

$$(\text{let}_f ((f_1 e_1) \dots (f_i (\text{fun } (c_i n_{i,1} \dots) e_i)) \dots (f_k e_k))$$
$$C_{\text{opt}}[(\text{app}_f f_i c m_1 \dots)])$$
$$\rightsquigarrow_{\text{opt}} (\text{let}_f ((f_1 e_1) \dots (f_k e_k))$$
$$C_{\text{opt}}[e_i\{c_i \rightarrow c\}\{n_{i,1} \rightarrow m_1\}\dots])$$

[when f_i is not free in $C_{\text{opt}}, e_1, \dots, e_n$]

The rule for continuations is similar.

General inlining

Non-linear inlining can also be performed trivially in CPS/L₃, either for continuation or for functions (illustrated here):

$$\begin{aligned} & (\text{let}_f \dots (f_i (\text{fun } (c_i \ n_{i,1} \dots) e_i)) \dots) \\ & \quad C_{\text{opt}}[(\text{app}_f f_i \ c \ m_1 \dots)] \\ \rightsquigarrow_{\text{opt}} & (\text{let}_f \dots (f_i (\text{fun } (c_i \ n_{i,1} \dots) e_i)) \dots) \\ & \quad C_{\text{opt}}[e_i\{c_i \rightarrow c\}\{n_{i,1} \rightarrow m_1\}\dots] \end{aligned}$$

(To preserve the uniqueness of names, fresh versions of bound names should be created during inlining.)

The problem of these rules is that they are not shrinking and rewriting does not even terminate with recursive continuations or functions.

Inlining heuristics (1)

Since non-shrinking inlining cannot be performed indiscriminately, heuristics are used to decide whether a candidate function should be inlined at a given call site.

These heuristics typically combine several factors, like:

- the size of the candidate function – smaller ones should be inlined more eagerly than bigger ones,
- the number of times the candidate is called in the whole program – a function called only a few times should be inlined,

(continued on next slide)

Inlining heuristics (2)

- the nature of the candidate – not much gain can be expected from the inlining of a recursive function,
- the kind of arguments passed to the candidate, and/or the way these are used in the candidate – constant arguments could lead to further reductions in the inlined candidate, especially if it combines them with other constants,
- etc.

Exercise

Imagine an imperative intermediate language equipped with a `return` statement to return from the current function to its caller.

1. Describe the problem that would appear when inlining a function containing such a `return` statement.
2. Explain how a `return` statement could be encoded in CPS/L₃ and why such an encoding would not suffer from the above problem.

CPS/L₃

“contification”

Contification

Contification is the name generally given to an optimization that transforms functions into (local) continuations.

When applicable, this transformation is interesting because it transforms expensive functions – compiled as closures – to inexpensive continuations – compiled as code blocks.

Contification example

Contification can for example transform the `loop` function in the L_3 example below to a local continuation, leading to efficient compiled code.

```
(def fact
  (fun (x)
    (rec loop ((i 1) (r 1))
      (if (> i x)
        r
        (loop (+ i 1) (* r i))))))
```

Contifiability

A CPS/L₃ function is contifiable if and only if it always returns to the same location, because then it does not need a return continuation.

For a non-recursive function, this condition is satisfied if and only if that function is only used in app_f nodes, in function position, and always passed the same return continuation.

For recursive functions, the condition is slightly more involved – see later.

Non-recursive contification

The contification of the non-recursive function f is given by:

$$\begin{aligned} & (\text{let}_f ((f (\text{fun } (c \ a_1 \ \dots) \ e))) \\ & \quad C_{\text{opt}}[C'_{\text{opt}}[(\text{app}_f \ f \ c_0 \ n_{1,1} \ \dots), (\text{app}_f \ f \ c_0 \ n_{2,1} \ \dots), \dots]]) \\ & \rightsquigarrow_{\text{opt}} C_{\text{opt}}[(\text{let}_c ((m (\text{cnt } (a_1 \ \dots) \ e\{c \rightarrow c_0\}))) \\ & \quad \quad C'_{\text{opt}}[(\text{app}_c \ m \ n_{1,1} \ \dots), (\text{app}_c \ m \ n_{2,1} \ \dots), \dots]])] \end{aligned}$$

where f does not appear free in either C_{opt} or C'_{opt} , and C'_{opt} is the smallest (multi-hole) context enclosing all applications of f . It ensures that the scope of m is as small as possible, and therefore that m obeys the scoping rules for continuations.

In this rule, c_0 is the (single) return continuation that is passed to function f .

Recursive contifiability

A set of mutually-recursive functions $F = \{ f_1, \dots, f_n \}$ is contifiable – which we write $\text{Cnt}(F)$ – if and only if every function in F is always used in one of the following two ways:

1. applied to a common return continuation, or
2. called in tail position by a function in F .

Intuitively, this ensures that all functions in F eventually return through the common continuation.

Example

As an example, functions `even` and `odd` in the CPS/ L_3 translation of the following L_3 term are contifiable:

```
(letrec
  ((even (fun (x)
          (if (= 0 x) #t (odd (- x 1)))))
  (odd (fun (x)
         (if (= 0 x) #f (even (- x 1)))))
  (even 12))
```

$\text{Cnt}(F = \{\text{even}, \text{odd}\})$ is satisfied since:

- the single use of `odd` is a tail call from `even` $\in F$,
- `even` is tail-called from `odd` $\in F$ and called with the continuation of the `letrec` statement – the common return continuation c_0 for this example.

Recursive contification

Given a set of mutually-recursive functions

$$(\text{let}_f ((f_1 e_1) (f_2 e_2) \dots (f_n e_n)) \\ e)$$

the condition $\text{Cnt}(F)$ for some $F \subseteq \{f_1, \dots, f_n\}$ can only be true if all the non tail calls to functions in F appear either:

- in the term e , or
- in the body of exactly one function $f_i \notin F$.

Therefore, two separate rewriting rules must be defined, one per case.

Recursive contification #1

When all non tail calls to functions in $F = \{ f_1, \dots, f_i \}$ appear in the body of the let_f , and $\text{Cnt}(F)$ holds, contification is performed by the following rewriting:

$$\begin{aligned} & (\text{let}_f ((f_1 (\text{fun } (c_1 a_{1,1} \dots) e_1)) \dots (f_n \dots)) \\ & \quad C_{\text{opt}}[e]) \\ & \rightsquigarrow_{\text{opt}} (\text{let}_f ((f_{i+1} (\text{fun } (c_{i+1} a_{i+1,1} \dots) e_{i+1})) \dots (f_n \dots)) \\ & \quad C_{\text{opt}}[(\text{let}_c ((m_1 (\text{cnt } (a_{1,1} \dots) \\ & \quad \quad e_1^*\{c_1 \rightarrow c_0\})) \dots) \\ & \quad \quad e^*)]) \end{aligned}$$

where f_1, \dots, f_i do not appear free in C_{opt} and e is minimal.

Note: the term t^* is t with all applications of contified functions transformed to continuation applications.

Contifiable subsets

Given a λet_f term defining a set of functions $F = \{f_1, \dots, f_n\}$, the subsets of F of potentially contifiable functions are obtained by:

1. building the tail-call graph of its functions, identifying the functions that call each-other in tail position,
2. extracting the strongly-connected components of that graph.

A given set of strongly-connected functions $F_i \subseteq F$ is then either contifiable together, i.e. $\text{Cnt}(F_i)$, or not contifiable at all – i.e. none of its subsets of functions are contifiable.