

# Optimising Typed Programs

Martin Elsmann\*  
University of Copenhagen

January 6, 1998

## Abstract

In this note we present a set of optimisations for an intermediate language of a Standard ML compiler. Most of the optimisations presented are off-the-shelf optimisations, including dead code elimination, constant folding, recursive function specialization, in-lining and value propagation. All optimisations are presented in a typed setting.

## 1 Introduction

Typed intermediate languages of optimising compilers are becoming increasingly recognised, mainly for two reasons. First, several type based optimisations and analyses have been suggested that cannot be done in an untyped setting. Such analyses and optimisations include various sorts of boxing analyses [Ler92, HJ94], intensional polymorphism [HM95] and region inference [TT94, BTV96]. Second, types can be used to provide certain safety guarantees for a program. In particular, by propagating types all the way to the target language, it becomes possible to type check programs just before execution. This has the advantage that typed executables may be shipped across the internet and if the executable type checks, it can be trusted.

In this note we describe optimisations for an intermediate language in the ML Kit with Regions compiler [TBE<sup>+</sup>97] (from hereon just the Kit) which is an optimising compiler for Standard ML [MTHM97]. We show that many of the important optimisations that are possible in an untyped setting, are also possible in a typed setting.

Optimisations performed in an intermediate language of an optimising compiler must satisfy at least two conditions. First, the optimisations must preserve types and semantics. Second, the optimisations may not lead to slower programs or programs that take up more space. In particular, it is important that all optimisations are “safe for space complexity.” That is, no optimisation may destroy space complexity properties of the program. Since space consumptions depend

---

\*Address: Department of Computer Science (DIKU), University of Copenhagen, Universitetsparken 1, DK-2100 Copenhagen, Denmark; email: mael@diku.dk.

on the underlying implementation technique, different restrictions apply for different underlying implementation techniques. In compilers building on garbage collection techniques one must be careful that no expression that may capture data in a closure is in-lined under a lambda-binding. For instance in-lining a selection from a large tuple in the body of a function may cause dead data to be captured in the closure [App92]. This restriction must also be enforced in a compiler building on region inference. Moreover, region inference and region representation analyses lays further restrictions on what optimisations may be performed. Yet, still quite a large set of optimisations are possible. We will not in this note address this issue further.

The optimisations that we present are off-the-shelf optimisations, including dead code elimination, constant folding, recursive function specialization, in-lining and value propagation. Some optimisations trigger other optimisations and vice versa. In the Kit these optimisations are naturally grouped together in what is called a contract phase. By keeping track of usage counts the contract phase may be implemented as a quasi-one-pass algorithm [AJ97]. In the following we focus on a small example language and present the optimisations for this language.

## 2 The Language and its Semantics

The language that we consider is a typed lambda language including a let-construct for polymorphism and a letrec-construct for expressing recursion. Further, the language includes constructs for records and sums.

We assume  $\text{TyVar}$  be a denumerably infinite set of *type variables*, ranged over by  $\alpha$ . Types and type schemes conform to the following syntax.

$$\begin{aligned} \tau & ::= \alpha \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 \times \tau_2 \mid \text{bool} \\ \sigma & ::= \forall \vec{\alpha}. \tau \end{aligned}$$

Type schemes are considered equal up-to renaming of bound type variables. A *substitution*  $S$  is a finite map from type variables to types. When  $A$  is any object and  $S$  is a substitution we write  $S(A)$  to mean simultaneous capture free substitution of  $S$  on  $A$ .

For any type scheme  $\sigma = \forall \alpha_1 \dots \alpha_n. \tau$  and type  $\tau'$ , we say that  $\tau'$  is an *instance of*  $\sigma$  (via  $S$ ), written  $\sigma \geq \tau'$ , if there exists a substitution  $S = \{\alpha_1 \mapsto \tau_1, \dots, \alpha_n \mapsto \tau_n\}$  such that  $S(\tau) = \tau'$ . The *instance list of*  $S$ , written  $il(S)$ , is the list  $[\tau_1, \dots, \tau_n]$  and we shall refer to lists of the above form as instance lists and use  $il$  to range over them. When  $\vec{\alpha} = \alpha_1 \dots \alpha_n$ ,  $n \geq 0$  is a list of type variables and  $il = [\tau_1, \dots, \tau_n]$  is an instance list we write  $\{\vec{il}/\vec{\alpha}\}$  to mean the substitution  $\{\alpha_1 \mapsto \tau_1, \dots, \alpha_n \mapsto \tau_n\}$ . Further, when  $A$  is any object we denote by  $\text{ftv}(A)$  the set of type variables occurring free in  $A$ . We consider a type  $\tau$  to be the type scheme  $\forall().\tau$ , hence the set of types is a subset of the set of type schemes.

## 2.1 Typed Expressions

In the following we use  $f$ ,  $x$  and  $y$  to range over a denumerably infinite set  $\text{Lvar}$  of *lambda variables*. The grammar for typed expressions is given below.

$$\begin{array}{l}
 e ::= \lambda x : \tau. e \mid e_1 e_2 \mid (e_1, e_2) \mid \pi_i e \mid x_{il} \\
 \mid \text{let } x : \sigma = e_1 \text{ in } e_2 \mid \text{true} \mid \text{false} \\
 \mid \text{if } e \text{ then } e_1 \text{ else } e_2 \\
 \mid \text{letrec } f : \sigma \ x = e_1 \text{ in } e_2
 \end{array}$$

We sometimes abbreviate  $x_{il}$  with  $x$ , when  $il = []$ .

A *type environment*,  $\Gamma$ , is a mapping from lambda variables to type schemes. The *static semantics* for the language relates types to expressions, under some assumptions, and describes what expressions are well-typed. The static semantics is presented as a set of inference rules allowing inferences among sentences of the form  $\Gamma \vdash e : \tau$ , where  $\Gamma$  is a type environment,  $e$  an expression and  $\tau$  a type.

**Expressions**

$\boxed{\Gamma \vdash e : \tau}$

$$\frac{\Gamma + \{x \mapsto \tau\} \vdash e : \tau'}{\Gamma \vdash \lambda x : \tau. e : \tau \rightarrow \tau'} \quad (1) \qquad \frac{\Gamma \vdash e_1 : \tau_1 \rightarrow \tau_2 \quad \Gamma \vdash e_2 : \tau_1}{\Gamma \vdash e_1 e_2 : \tau_2} \quad (2)$$

$$\frac{\Gamma \vdash e_1 : \tau_1 \quad \Gamma \vdash e_2 : \tau_2}{\Gamma \vdash (e_1, e_2) : \tau_1 \times \tau_2} \quad (3) \qquad \frac{i \in \{1, 2\} \quad \Gamma \vdash e : \tau_1 \times \tau_2}{\Gamma \vdash \pi_i e : \tau_i} \quad (4)$$

$$\frac{\Gamma(x) \geq \tau \text{ via } S}{\Gamma \vdash x_{il(S)} : \tau} \quad (5) \qquad \frac{\Gamma \vdash e_1 : \tau \quad \vec{\alpha} = \text{ftv}(\tau) \setminus \text{ftv}(\Gamma)}{\Gamma + \{x \mapsto \forall \vec{\alpha}. \tau\} \vdash e_2 : \tau'} \quad (6)$$

$$\frac{\Gamma \vdash e : \text{bool} \quad \Gamma \vdash e_1 : \tau \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash \text{if } e \text{ then } e_1 \text{ else } e_2 : \tau} \quad (7) \qquad \frac{}{\Gamma \vdash \text{false} : \text{bool}} \quad (8)$$

$$\frac{\tau = \tau' \rightarrow \tau'' \quad \vec{\alpha} = \text{ftv}(\tau) \setminus \text{ftv}(\Gamma) \quad \Gamma + \{f \mapsto \tau\} \vdash \lambda x : \tau'. e_1 : \tau \quad \Gamma + \{f \mapsto \forall \vec{\alpha}. \tau\} \vdash e_2 : \tau'''}{\Gamma \vdash \text{letrec } f : \forall \vec{\alpha}. \tau \ x = e_1 \text{ in } e_2 : \tau'''} \quad (9) \qquad \frac{}{\Gamma \vdash \text{true} : \text{bool}} \quad (10)$$

When  $e$  is any expression we write  $\text{flv}(e)$  to mean the set of free lambda variables in  $e$ . Further, when  $e$  and  $e'$  are expressions and  $x$  is a lambda variable, we

write  $e\{e'/x\}$  to mean capture free substitution of  $e'$  for  $x$  in  $e$ . Expressions are considered equal up-to renaming of bound lambda variables and type variables.

The *restriction* of an environment  $\Gamma$  to a set of lambda variables  $A \subseteq \text{Dom}(\Gamma)$ , written  $\Gamma \downarrow A$ , is the environment with domain  $A$  and values  $(\Gamma \downarrow A)(x) = \Gamma(x)$ . Further, we say that an environment  $\Gamma$  *enriches* another environment  $\Gamma'$ , written  $\Gamma \sqsupseteq \Gamma'$ , if  $\text{Dom}(\Gamma) \supseteq \text{Dom}(\Gamma')$  and  $\Gamma(x) = \Gamma'(x)$  for all  $x \in \text{Dom}(\Gamma')$ .

The following restriction lemma can be proved by induction over the structure of expressions.

**Lemma 2.1 (Elaboration closed under restriction)** *For all environments  $\Gamma$  and  $\Gamma'$ , expressions  $e$  and types  $\tau$ , if  $\Gamma \vdash e : \tau$  and  $\Gamma' \sqsupseteq (\Gamma \downarrow \text{fv}(e))$  then  $\Gamma' \vdash e : \tau$ .*

Further, the following substitution lemma can be proved by induction over the structure of expressions.

**Lemma 2.2 (Elaboration closed under substitution)** *For all environments  $\Gamma$ , expressions  $e$ , types  $\tau$  and substitutions  $S$ , if  $\Gamma \vdash e : \tau$  then  $S(\Gamma) \vdash S(e) : S(\tau)$ .*

## 2.2 Untyped Expressions

To obtain an untyped expression from a typed expression we define an erasure operation  $er$ . A couple of the defining equations are given below.

$$\begin{aligned} er(\lambda x : \tau. e) &= \lambda x. er(e) \\ er(e_1 e_2) &= er(e_1) er(e_2) \\ &\vdots \end{aligned}$$

The *dynamic semantics* of the language relates untyped expressions to so called values, under some assumptions associating values to variables. Thus, a *dynamic environment*,  $\mathcal{E}$ , maps lambda variables to values, which again conform to the grammar:

$$v ::= \text{clos}(\lambda x. e, \mathcal{E}) \mid \text{true} \mid \text{false} \mid (v_1, v_2)$$

The rules of the dynamic semantics allows inferences among sentences of the form  $\mathcal{E} \vdash e \Downarrow v$ , where  $\mathcal{E}$  is a dynamic environment,  $e$  is an untyped expression and  $v$  is a value. To give meaning to recursive functions we allow creation of non-well-founded objects; i.e. closures,  $cl_\infty$ , with the property

$$cl_\infty = \text{clos}(\lambda x. e, \mathcal{E} + \{f \mapsto cl_\infty\})$$

where  $x$  and  $f$  are lambda variables,  $e$  is an expression and  $\mathcal{E}$  is a dynamic environment [MT91].

## Expressions

$$\boxed{\mathcal{E} \vdash e \Downarrow v}$$

$$\frac{\mathcal{E}(x) = v}{\mathcal{E} \vdash x \Downarrow v} \quad (11)$$

$$\frac{}{\mathcal{E} \vdash \lambda x.e \Downarrow \text{clos}(\lambda x.e, \mathcal{E})} \quad (12)$$

$$\frac{\mathcal{E} \vdash e_1 \Downarrow \text{clos}(\lambda x.e, \mathcal{E}_0) \quad \mathcal{E} \vdash e_2 \Downarrow v \quad \mathcal{E}_0 + \{x \mapsto v\} \vdash e \Downarrow v'}{\mathcal{E} \vdash e_1 e_2 \Downarrow v'} \quad (13)$$

$$\frac{\mathcal{E} \vdash e_1 \Downarrow v_1 \quad \mathcal{E} \vdash e_2 \Downarrow v_2}{\mathcal{E} \vdash (e_1, e_2) \Downarrow (v_1, v_2)} \quad (14)$$

$$\frac{i \in \{1, 2\} \quad \mathcal{E} \vdash e \Downarrow (v_1, v_2)}{\mathcal{E} \vdash \pi_i e \Downarrow v_i} \quad (15)$$

$$\frac{\mathcal{E} \vdash e \Downarrow \text{true} \quad \mathcal{E} \vdash e_1 \Downarrow v}{\mathcal{E} \vdash \text{if } e \text{ then } e_1 \text{ else } e_2 \Downarrow v} \quad (16)$$

$$\frac{\mathcal{E} \vdash e \Downarrow \text{false} \quad \mathcal{E} \vdash e_2 \Downarrow v}{\mathcal{E} \vdash \text{if } e \text{ then } e_1 \text{ else } e_2 \Downarrow v} \quad (17)$$

$$\frac{\mathcal{E} \vdash e_1 \Downarrow v_1 \quad \mathcal{E} + \{x \mapsto v_1\} \vdash e_2 \Downarrow v_2}{\mathcal{E} \vdash \text{let } x = e_1 \text{ in } e_2 \Downarrow v_2} \quad (18)$$

$$\frac{}{\mathcal{E} \vdash \text{false} \Downarrow \text{false}} \quad (19)$$

$$\frac{cl_\infty = \text{clos}(\lambda x.e_1, \mathcal{E} + \{f \mapsto cl_\infty\}) \quad \mathcal{E} + \{f \mapsto cl_\infty\} \vdash e_2 \Downarrow v_2}{\mathcal{E} \vdash \text{letrec } f \ x = e_1 \text{ in } e_2 \Downarrow v_2} \quad (20)$$

$$\frac{}{\mathcal{E} \vdash \text{true} \Downarrow \text{true}} \quad (21)$$

### 2.3 Expression Contexts

We define three notions of expression contexts; local expression contexts, allowing one to single out a local expression (not going under a lambda or a letrec construct); global expression contexts, allowing one to single out any sub-expression; and multi expression contexts, allowing one to single out multiple occurrences of an expression. A *local expression context*,  $L$ , takes the following form.

$$L ::= [\cdot] \mid L e \mid e L \mid (L, e) \mid (e, L) \mid \pi_i L$$

$$\begin{array}{l} | \text{let } x : \sigma = L \text{ in } e \mid \text{let } x : \sigma = e \text{ in } L \\ | \text{letrec } f : \sigma \ x = e \text{ in } L \\ | \text{if } L \text{ then } e_1 \text{ else } e_2 \\ | \text{if } e \text{ then } L \text{ else } e_2 \\ | \text{if } e \text{ then } e_1 \text{ else } L \end{array}$$

Further, *global expression contexts*,  $C$ , takes the following form.

$$C ::= L \mid L[\lambda x : \tau. C] \mid L[\mathbf{letrec} f : \sigma x = C \mathbf{in} e]$$

Finally, *multi expression contexts*,  $M$ , takes the following form.

$$\begin{aligned} M ::= & [\cdot] \mid e \mid \lambda x : \tau. M \mid M_1 M_2 \mid (M_1, M_2) \mid \pi_i M \\ & \mid \mathbf{let} x : \sigma = M_1 \mathbf{in} M_2 \mid \mathbf{if} M \mathbf{then} M_1 \mathbf{else} M_2 \\ & \mid \mathbf{letrec} f : \sigma x = M_1 \mathbf{in} M_2 \end{aligned}$$

When  $\mathcal{C}$  is either a local expression context, a global expression context or a multi expression context, and  $e$  is an expression, the effect of *filling*  $\mathcal{C}$  with  $e$ , written  $\mathcal{C}[e]$  is to replace all appearances of the hole in  $\mathcal{C}$  with  $e$ . We write  $\text{flv}(\mathcal{C})$  to mean the set of free lambda variables in  $\mathcal{C}$ .

## 2.4 Small, Calling and Safe Expressions

A *small* expression is an expression with at most  $k_{\text{small}}$  nodes in the abstract syntax tree for the expression. The constant controls what non-recursive functions are in-lined and what recursive functions are specialized. It is a trade-off between code-size and speed. In the Kit a value of 20 is used.

An expression  $e$  is considered a *calling expression* if it can be written,  $L[e_1 e_2]$ , where  $L$  is any local expression context and  $e_1$  and  $e_2$  are any expressions.

The set of terminating expressions that cannot perform any side-effects on the store and cannot raise any exceptions are candidates to dead code elimination and other kinds of optimising reductions. As a simple approximation we consider an expression to be safe if it is not a calling expression and if it cannot raise any exception or update the store. In the small example language considered here an expression is *safe* if it is not a calling expression.

## 3 Simple Optimising Reductions

In the following we use the term reduction to refer to optimisation reductions, relating expressions. We use  $\longrightarrow$  to range over reductions, and we say that  $e$  *reduces* to  $e'$  if  $e \longrightarrow e'$ . We now define what it means for a reduction to be type preserving.

**Definition 3.1** A reduction,  $\longrightarrow$ , is *type preserving*, if for all environments  $\Gamma$ , expressions  $e$  and  $e'$ , and types  $\tau$ , such that  $\Gamma \vdash e : \tau$  and  $e \longrightarrow e'$  then  $\Gamma \vdash e' : \tau$ .  $\square$

The following lemma can be proved by induction over the structure of global expression contexts.

**Lemma 3.2** *Assume  $\longrightarrow$  is type preserving. For all environments  $\Gamma$ , expressions  $e$  and  $e'$ , global expression contexts  $C$  and types  $\tau$ , if  $\Gamma \vdash C[e] : \tau$  and  $e \longrightarrow e'$  then  $\Gamma \vdash C[e'] : \tau$ .*

We now present a set of simple optimising reductions. By use of Lemma 2.1 and Lemma 2.2 each of the reductions can be shown to preserve types.

### 3.1 Reduction of Projections from Explicit Records

Folding of conditional expressions on constants is implemented by value propagation. Reduction of projections from explicit records is implemented by the following rules.

$$\pi_1 (e_1, e_2) \xrightarrow[\text{proj}]{e_2 \text{ safe}} e_1 \quad (22)$$

$$\pi_2 (e_1, e_2) \xrightarrow[\text{proj}]{e_1 \text{ safe}} e_2 \quad (23)$$

### 3.2 Dead Code Elimination

Dead code elimination is implemented by the following rules.

$$\text{let } x : \sigma = e_1 \text{ in } e_2 \xrightarrow[\text{dce}]{\substack{e_1 \text{ safe} \\ x \notin \text{flv}(e_2)}} e_2 \quad (24)$$

$$\text{letrec } f : \sigma \ x = e_1 \text{ in } e_2 \xrightarrow[\text{dce}]{f \notin \text{flv}(e_2)} e_2 \quad (25)$$

### 3.3 Let-reductions

Simple let-constructs are reduced and explicit applications of lambda-constructs are reduced to let-constructs.

$$\text{let } x : \forall \vec{\alpha}. \tau = e_1 \text{ in } x_{il} \xrightarrow{\text{let}} e_1 \{il/\vec{\alpha}\} \quad (26)$$

$$(\lambda x : \tau. e_2) e_1 \xrightarrow{\lambda} \text{let } x : \tau = e_1 \text{ in } e_2 \quad (27)$$

### 3.4 Letrec-reductions

Several reductions are performed on letrec-constructs. Non-recursive functions bound by letrec-constructs are reduced to let-bindings of lambda-constructs as follows.

$$\begin{aligned} \text{letrec } f : \forall \vec{\alpha}. \tau_1 \rightarrow \tau_2 \ x = e_1 \text{ in } e_2 & \xrightarrow[\text{rec1}]{f \notin \text{flv}(e_1)} \\ \text{let } f : \forall \vec{\alpha}. \tau_1 \rightarrow \tau_2 = \lambda x : \tau_1. e_1 \text{ in } e_2 & \end{aligned} \quad (28)$$

The following reduction rule opens for other reductions.

$$\begin{aligned} (\text{letrec } f : \sigma \ x = e_1 \text{ in } f_{il}) e_2 & \xrightarrow{\text{rec2}} \\ \text{letrec } f : \sigma \ x = e_1 \text{ in } (f_{il} e_2) & \end{aligned} \quad (29)$$

## 4 In-lining Non-recursive Functions

In-lining ( $\beta$ -reduction) is an important optimisation for functional programs. Non-recursive functions referenced only once may always be in-lined and small functions referenced more than once may also be in-lined. The Kit implements the following in-lining strategies.

$$\text{let } f : \forall \vec{\alpha}. \tau = \lambda x : \tau'. e \text{ in } C[fi] \xrightarrow[\text{inl1}]{f \notin \text{fv}(C)} C[(\lambda x : \tau'. e)\{i/\vec{\alpha}\}] \quad (30)$$

$$\text{let } f : \forall \vec{\alpha}. \tau = \lambda x : \tau'. e \text{ in } e' \xrightarrow[\text{inl2}]{e \text{ small}} e' \{((\lambda x : \tau'. e)\{i/\vec{\alpha}\})/f_i\} \quad (31)$$

## 5 Specializing Recursive Functions

Specialization of recursive functions is an important optimisation in a compiler for a functional language [SW95]. In particular specializations of small functions as *fold* and *map* with respect to their first arguments lead to important speedups without drastically increasing the code size.

As an example consider the following Standard ML program.

```
let fun map f [] = []
    | map f (x::xs) = f x :: map f xs
in map (fn x => x+1) [1,2]
end
```

By specializing recursive functions this program is transformed into the following optimised program.

```
let fun map [] = []
    | map (x::xs) = x+1 :: map xs
in map [1,2]
end
```

The function `map` is now first-order and further, the successor function is in-lined into the body of the `map` function. This optimisation has an important effect on performance.

Small recursive functions may be specialized according to the following rule.

$$\text{letrec } f : \forall \vec{\alpha}. \tau_1 \rightarrow (\tau_2 \rightarrow \tau_3) \ x = \lambda y : \tau_2. M[f \ x] \text{ in } C[fi \ e] \xrightarrow[\text{spec1}]{M[f] \text{ small}} \quad (32)$$

$$\text{letrec } f : \forall \vec{\alpha}. \tau_1 \rightarrow (\tau_2 \rightarrow \tau_3) \ x = \lambda y : \tau_2. M[f \ x] \text{ in } C \left[ \begin{array}{l} \text{let } x : \tau_1 \{i/\vec{\alpha}\} = e \text{ in} \\ \text{letrec } f : (\tau_2 \rightarrow \tau_3) \{i/\vec{\alpha}\} \ y = \\ M[f] \{i/\vec{\alpha}\} \\ \text{in } f \end{array} \right]$$

Note that the original binding of the recursive function is not removed by the reduction.

Even large recursive functions may be specialized. This is captured by the following rule.

$$\begin{array}{l}
\mathbf{letrec} \ f : \forall \vec{\alpha}. \tau_1 \rightarrow (\tau_2 \rightarrow \tau_3) \ x = \\
\quad \lambda y : \tau_2. M[f \ x] \\
\mathbf{in} \ C[f_{il} \ e]
\end{array}
\quad
\begin{array}{l}
\begin{array}{l}
f \notin \text{flv}(M) \\
f \notin \text{flv}(C) \cup \text{flv}(e) \\
\longrightarrow_{\text{spec2}}
\end{array}
\end{array}
\quad (33)$$

$$C \left[ \begin{array}{l}
\mathbf{let} \ x : \tau_1 \{il/\vec{\alpha}\} = e \ \mathbf{in} \\
\mathbf{letrec} \ f : (\tau_2 \rightarrow \tau_3) \{il/\vec{\alpha}\} \ y = \\
\quad M[f] \{il/\vec{\alpha}\} \\
\mathbf{in} \ f
\end{array} \right]$$

In this case the original binding of the recursive function is removed by the reduction.

## 6 Value Propagation

By propagating information throughout a program about to what kinds of values a variable is bound, it is possible to eliminate many unnecessary fetches from records and unnecessary conditional checks. The set of abstract values `AbsVal` is defined by the following syntax. We use  $V$  to range over `AbsVal`.

$$V ::= \top \mid x_{il} \mid (V_1, V_2) \mid \mathbf{true} \mid \mathbf{false}$$

We say that an abstract value  $V$  is *simple* if  $V$  is either a constant `true` or `false`, or a variable  $x_{il}$ . When  $V_1$  and  $V_2$  are abstract values we define the *least upper bound* of  $V_1$  and  $V_2$ , written  $V_1 \sqcap V_2$ , recursively as follows.

$$V \sqcap V' = \begin{cases} V & \text{if } V = V' \text{ and } V \text{ simple} \\ (V_1 \sqcap V'_1, V_2 \sqcap V'_2) & \text{if } V = (V_1, V_2) \text{ and } V' = (V'_1, V'_2) \\ \top & \text{otherwise} \end{cases}$$

Further, when  $V$  is an abstract value and  $x$  is a lambda variable, we define the *exclusion* of  $x$  from  $V$ , written  $V \setminus x$ , recursively as follows.

$$V \setminus x = \begin{cases} \top & \text{if } V = x_{il} \\ (V_1 \setminus x, V_2 \setminus x) & \text{if } V = (V_1, V_2) \\ V & \text{otherwise} \end{cases}$$

A propagation environment  $\Phi$  is a finite mapping from lambda variables to pairs of a list of type variables and an abstract value.

$$\Phi \in \text{PropEnv} = \text{Lvar} \xrightarrow{\text{fin}} \text{TyVar}^{(k)} \times \text{AbsVal}$$

Below we state a propagation function  $\mathcal{P} : \text{Lexp} \rightarrow \text{PropEnv} \rightarrow \text{Lexp} \times \text{AbsVal}$ . Given an expression and a propagation environment the propagation function  $\mathcal{P}$  computes an optimised expression and an abstract value for the expression.

$$\begin{aligned}
\mathcal{P} \llbracket \text{true} \rrbracket \Phi &= (\text{true}, \text{true}) \\
\mathcal{P} \llbracket \text{false} \rrbracket \Phi &= (\text{false}, \text{false}) \\
\mathcal{P} \llbracket \lambda x : \tau. e \rrbracket \Phi &= \text{let } (e', \_ ) = \mathcal{P} \llbracket e \rrbracket (\Phi + \{x \mapsto (\_, \top)\}) \\
&\quad \text{in } (\lambda x : \tau. e', \top) \\
\mathcal{P} \llbracket e_1 e_2 \rrbracket \Phi &= \text{let } (e'_1, \_ ) = \mathcal{P} \llbracket e_1 \rrbracket \Phi \\
&\quad (e'_2, \_ ) = \mathcal{P} \llbracket e_2 \rrbracket \Phi \\
&\quad \text{in } (e'_1 e'_2, \top) \\
\mathcal{P} \llbracket (e_1, e_2) \rrbracket \Phi &= \text{let } (e'_1, V_1) = \mathcal{P} \llbracket e_1 \rrbracket \Phi \\
&\quad (e'_2, V_2) = \mathcal{P} \llbracket e_2 \rrbracket \Phi \\
&\quad \text{in } ((e'_1, e'_2), (V_1, V_2)) \\
\mathcal{P} \llbracket \text{let } x : \forall \vec{\alpha}. \tau = e_1 \text{ in } e_2 \rrbracket \Phi &= \\
&\quad \text{let } (e'_1, V_1) = \mathcal{P} \llbracket e_1 \rrbracket \Phi \\
&\quad (e'_2, V_2) = \mathcal{P} \llbracket e_2 \rrbracket (\Phi + \{x \mapsto (\vec{\alpha}, V_1)\}) \\
&\quad \text{in } (\text{let } x : \forall \vec{\alpha}. \tau = e'_1 \text{ in } e'_2, V_2 \parallel x) \\
\mathcal{P} \llbracket \text{letrec } f : \forall \vec{\alpha}. \tau x = e_1 \text{ in } e_2 \rrbracket \Phi &= \\
&\quad \text{let } (e'_1, \_ ) = \mathcal{P} \llbracket e_1 \rrbracket (\Phi + \{f \mapsto (\_, \top), x \mapsto (\_, \top)\}) \\
&\quad (e'_2, V_2) = \mathcal{P} \llbracket e_2 \rrbracket (\Phi + \{f \mapsto (\vec{\alpha}, \top)\}) \\
&\quad \text{in } (\text{letrec } f : \forall \vec{\alpha}. \tau x = e'_1 \text{ in } e'_2, V_2 \parallel f) \\
\mathcal{P} \llbracket \pi_i e \rrbracket \Phi &= \\
&\quad \begin{cases} (V_i, V_i) & \text{if } V = (V_1, V_2), V_i \text{ simple and } e' \text{ safe} \\ (\pi_i e', V_i) & \text{if } V = (V_1, V_2) \text{ and } (V_i \text{ not simple or } e' \text{ not safe}) \\ (\pi_i e', \top) & \text{otherwise} \end{cases} \\
&\quad \text{where } (e', V) = \mathcal{P} \llbracket e \rrbracket \Phi \\
\mathcal{P} \llbracket \text{if } e \text{ then } e_1 \text{ else } e_2 \rrbracket \Phi &= \\
&\quad \begin{cases} (e'_1, V_1) & \text{if } V' \text{ not simple, } V = \text{true and } e' \text{ safe} \\ (e'_2, V_2) & \text{if } V' \text{ not simple, } V = \text{false and } e' \text{ safe} \\ (V', V') & \text{if } V' \text{ simple and } e' \text{ safe} \\ (\text{if } e' \text{ then } e'_1 \text{ else } e'_2, V') & \text{otherwise} \end{cases} \\
&\quad \text{where } (e', V) = \mathcal{P} \llbracket e \rrbracket \Phi \\
&\quad (\Phi_t, \Phi_f) = \text{if } V = x_\square \text{ then } ( \Phi + \{x \mapsto (\_, \text{true})\}, \\
&\quad \quad \quad \Phi + \{x \mapsto (\_, \text{false})\}) \\
&\quad \quad \quad \text{else } (\Phi, \Phi) \\
&\quad (e'_1, V_1) = \mathcal{P} \llbracket e_1 \rrbracket \Phi_t \\
&\quad (e'_2, V_2) = \mathcal{P} \llbracket e_2 \rrbracket \Phi_f \\
&\quad V' = V_1 \sqcap V_2 \\
\mathcal{P} \llbracket x_{il} \rrbracket \Phi &= \begin{cases} (V, V)\{il/\vec{\alpha}\} & \text{if } V \text{ simple} \\ (x_{il}, x_{il}) & \text{if } V = \top \\ (x_{il}, V)\{il/\vec{\alpha}\} & \text{otherwise} \end{cases} \\
&\quad \text{where } (\vec{\alpha}, V) = \Phi(x)
\end{aligned}$$

## 7 Implementation

All optimisations presented here have been implemented in the ML Kit with Regions compiler. Due to the sensitivity of region inference and region representation analyses w.r.t. changes in the program, the value propagation algorithm used in the Kit is not as aggressive as the algorithm presented here.

A straight-forward implementation of the optimisation rules will quickly show to be at least quadratic in the size of the program. Inspired by [AJ97] the optimisations presented here are implemented in the Kit by two functions, *reduce* and *contract*. The function *contract* maintains an environment and applies the function *reduce* on the way down the syntax tree and the way up the syntax tree. The function *reduce* reduces redexes w.r.t. environment and usage information. It updates usage information when inserting and removing sub-expressions.

Consider the expression `let  $x = e_1$  in  $e_2$` . On the way down, if there is zero uses of  $x$  then eliminate  $e_1$  (if it is safe) and decrement uses in  $e_1$  prior to recurring on  $e_2$ . This may trigger in-lining in  $e_2$  of functions also applied in  $e_1$ , e.g. On the way up, if there is now zero uses of  $x$  then eliminate the binding (if it is safe) and decrement uses in  $e_1$  prior to returning.

## 8 Conclusions

In this note we have presented a set of off-the-shelf optimisations for a typed intermediate language of the Kit which is a Standard ML compiler. It has not been possible to present all optimisations performed in the Kit using the small example language presented in this note. For instance, the intermediate language of the Kit includes a fix-construct to allow for mutually recursive functions. As an optimisation, the Kit minimises each fix-construct by finding strongly connected components of the call-graph associated with each fix-construct.

Further, the Kit implements a few other optimisations that we do not mention here. These are optimisations that are performed only to improve on region inference.

## References

- [AJ97] Andrew W. Appel and Trevor Jim. Shrinking lambda expressions in linear time. In *Journal of Functional Programming*, 1997.
- [App92] Andrew W. Appel. *Compiling With Continuations*. Cambridge University Press, 1992.
- [BTV96] Lars Birkedal, Mads Tofte, and Magnus Vejlstrup. From region inference to von Neumann machines via region representation inference. In *23st ACM Symposium on Principles of Programming Languages*, January 1996.

- [HJ94] Fritz Henglein and Jesper Jørgensen. Formally optimal boxing. In *21st Annual ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages. Portland, Oregon*, pages 213–226, January 1994.
- [HM95] Robert Harper and Greg Morrisett. Compiling polymorphism using intensional type analysis. In *Principles of Programming Languages*, San Francisco, January 1995.
- [Ler92] Xavier Leroy. Unboxed objects and polymorphic typing. In *Principles of Programming Languages*, pages 177–188, 1992.
- [MT91] Robin Milner and Mads Tofte. Co-induction in relational semantics. In *Theoretical Computer Science 87*, 1991.
- [MTHM97] Robin Milner, Mads Tofte, Robert Harper, and David MacQueen. *The Definition of Standard ML (Revised)*. MIT Press, 1997.
- [SW95] Manuel Serrano and Pierre Weis. Bigloo: a portable and optimizing compiler for strict functional languages. In *Second International Symposium on Static Analysis (SAS)*, pages 366–381, September 1995.
- [TBE<sup>+</sup>97] Mads Tofte, Lars Birkedal, Martin Elsman, Niels Hallenberg, Tommy Højfeldt Olesen, Peter Sestoft, and Peter Bertelsen. Programming with regions in the ML Kit. Technical Report DIKU-TR-97/12, Dept. of Computer Science, University of Copenhagen, 1997. (<http://www.diku.dk/research-groups/topps/activities/kit2>).
- [TT94] Mads Tofte and Jean-Pierre Talpin. Implementation of the typed call-by-value  $\lambda$ -calculus using a stack of regions. In *21st ACM Symposium on Principles of Programming Languages*, January 1994.