

Witnessing Purity, Constancy and Mutability

Ben Lippmeier

Department of Computer Science
Australian National University
Ben.Lippmeier@anu.edu.au

Abstract. Restricting destructive update to values of a distinguished reference type prevents functions from being polymorphic in the mutability of their arguments. This restriction makes it easier to reason about program behaviour during transformation, but the lack of polymorphism reduces the expressiveness of the language. We present a System-F style core language that uses dependently kinded proof witnesses to encode information about the mutability of values and the purity of computations. We support mixed strict and lazy evaluation, and use our type system to ensure that only computations without visible side effects are suspended.

1 Introduction

Suppose we are writing a library that provides a useful data structure such as linked lists. A Haskell-style definition for the list type would be:

```
data List a = Nil | Cons a (List a)
```

The core language of compilers such as GHC is based around System-F [15]. Here is the translation of the standard *map* function to this representation, complete with type abstractions and applications:

```
map :: ∀a b. (a → b) → List a → List b
map = Λa. Λb. λ(f : a → b). λ(list : List a).
  case list of
    Nil      → Nil b
    Cons x xs → Cons b (f x) (map a b f xs)
```

Say we went on to define some other useful list functions, and then decided that we need one to destructively insert a new element into the middle of a list. In Haskell, side effects are carefully controlled and we would need to introduce a monad such as ST or IO [8] to encapsulate the effects due to the update. Destructive update is also limited to distinguished types such as *STRef* and *IORef*. We cannot use our previous list type, so will instead change it to use an *IORef*.

```
data List a = Nil | Cons a (IORef (List a))
```

Unfortunately, as we have changed the structure of our original data type, we can no longer use the previous definition of *map*, or any other functions we defined earlier. We must go back and refactor each of these function definitions to use the new type. We must insert calls to *readIORef* and use monadic sequencing combinators instead of vanilla *let* and *where*-expressions. However, doing so introduces explicit data dependencies into the core program. This in turn reduces the compiler’s ability to perform optimisations such as deforestation and the full laziness transform [6], which require functions to be written in the “pure”, non-monadic style. It appears that we need *two* versions of our list structure and its associated functions, an immutable version that can be optimised, and a mutable one that can be updated.

Variations of this problem are also present in ML and O’Caml. In ML, mutability is restricted to *ref* and *array* types [11]. In O’Caml, record types can have mutable fields, but variant types cannot [9]. Similarly to Haskell, in these languages we are forced to insert explicit reference types into the definitions of mutable data structures, which makes them incompatible with the standard immutable ones. This paper shows how to avoid this problem:

- We present a System-F style core language that uses region and effect typing to guide program optimisation. Optimisations that depend on purity can be performed on the the pure fragments of the program.
- We use region variables and dependently kinded witnesses to encode mutability polymorphism. This allows arbitrary data structures to be mutable without changing the structure of their value types.
- We use call-by-value evaluation as default, but support lazy evaluation via a primitive *suspend* operator. We use witnesses of purity to ensure that only pure function applications can be suspended.

Our goals are similar to those of Benton and Kennedy [3], but as in [15] we use a System-F based core language instead of a monadic one. Type inference and translation from source to core is discussed in [10].

2 Regions, Effects and Mutability Constraints

In Haskell and ML, references and arrays are distinguished values, and are the only ones capable of being destructively updated. This means that the structure of mutable data is necessarily different from the structure of constant data, which makes it difficult to write polymorphic functions that act on both. For example, if we use *IORef Int* as the type of a mutable integer and *Int* as the type of a constant integer, then we would need *readIORef* to access the first, but not the second. On the other hand, if we were to treat all data as mutable, then every function would exhibit a side effect. This would prevent us from using code-motion style optimisations that depend on purity.

Instead, we give integers the type *Int r*, where *r* is a region variable, and constrain *r* to be mutable or constant as needed. Our use of region variables is similar to that by Talpin and Jouvelot [16], where the variable *r* is a name for

a set of locations in the store where a run-time object may lie. We do not use regions for controlling allocation as per [17], due to the difficulty of statically determining when objects referenced by suspended computations can be safely deallocated. We define region variables to have kind $\%$, and use this symbol because pictorially it is two circles separated by a line, a mnemonic for “this, or that”. The kind of value types is $*$, so the *Int* type constructor has kind $Int :: \% \rightarrow *$. The type of a literal integer such as ‘5’ is:

$$5 :: \forall(r : \%). Int\ r$$

In our System-F style language, type application corresponds to instantiation, and ‘5’ is the name of a function that allocates a new integer object into a given region. Note that unlike [16] we do not use allocation effects. This prevents us from optimising away some forms of duplicated computation, such as described in §7 of [4], but also simplifies our type system. For the rest of this paper we will elide explicit kind annotations on binders when they are clear from context.

2.1 Updating Integers

To update an integer we use the *updateInt* function which has type:

$$updateInt :: \forall r_1\ r_2. Mutable\ r_1 \Rightarrow Int\ r_1 \rightarrow Int\ r_2 \xrightarrow{Read\ r_2 \vee Write\ r_1} ()$$

This function reads the value of its second integer argument, and uses this to overwrite the value of the first. As in [16] we annotate function types with their latent effects. We organise effects as a lattice and collect atomic effects with the \vee operator. We use \perp as the effect of a pure function, and unannotated function arrows are taken to have this effect. We also use a set-like subtraction operator where the effect $\sigma \setminus \sigma'$ contains the atomic effects that appear in σ but not σ' . We use $!$ as the kind of effects, so *Read* has kind $Read :: \% \rightarrow !$. The symbol $!$ is a mnemonic for “something’s happening!”.

Returning to the type of *updateInt*, *Mutable* r_1 is a *region constraint* that ensures that only mutable integers may be updated. When we call this function we must pass a *witness* to the fact that this constraint is satisfied, a point we will discuss further in §3.

When the number of atomic effects becomes large, using the above syntax for effects becomes cumbersome. Due to this we sometimes write effect terms after the body of the type instead:

$$updateInt :: \forall r_1\ r_2. Mutable\ r_1 \Rightarrow Int\ r_1 \rightarrow Int\ r_2 \xrightarrow{e_1} () \\ \triangleright e_1 = Read\ r_2 \vee Write\ r_1$$

The symbol \triangleright is pronounced “with”. Note that the effect variable e_1 is not quantified. It has been introduced for convenience only and is not a parameter of the type.

2.2 Updating Algebraic Data


Along with primitive types such as *Int*, the definition of an algebraic data type can also contain region variables. For example, we define our lists as follows:

data $List\ r\ a = Nil \mid Cons\ a\ (List\ r\ a)$

This definition is similar to the one from §1 except that we have also applied the $List$ constructor to a region variable. This variable identifies the region that contains the list cells, and can be constrained to be constant or mutable as needed. The definition also introduces data constructors that have the following types:

$$\begin{aligned} Nil &:: \forall r\ a. List\ r\ a \\ Cons &:: \forall r\ a. a \rightarrow List\ r\ a \rightarrow List\ r\ a \end{aligned}$$

In the type of Nil , the fact that r is quantified indicates that this constructor allocates a new Nil object. Freshly allocated objects do not alias with existing objects, so they can be taken to be in any region. On the other hand, in the type of $Cons$, the type of the second argument and return value share the same region variable r , which means the new cons-cell is allocated into the same region as the existing cells. For example, evaluation of the following expression produces the store objects shown below.

$$\begin{aligned} list &:: List\ r_5\ (Int\ r_6) \\ list &= Cons\ r_5\ (Int\ r_6)\ (2\ r_6)\ (Cons\ r_5\ (Int\ r_6)\ (3\ r_6)\ (Nil\ r_5)) \end{aligned}$$


As the list cells and integer elements are in different regions, we can give them differing mutabilities. If the type of $list$ was constrained as follows, then we would be free to update the integer elements, but not the spine.

$$list :: Const\ r_5 \Rightarrow Mutable\ r_6 \Rightarrow List\ r_5\ (Int\ r_6)$$

The definition of an algebraic type also introduces a set of update operators, one for each updatable component of the corresponding value. For our list type, as we could usefully update the head and tail pointers in a cons-cell, we get the following operators:

$$\begin{aligned} update_{Cons,0} &:: \forall r\ a. Mutable\ r \Rightarrow List\ r\ a \rightarrow a \xrightarrow{Write\ r} () \\ update_{Cons,1} &:: \forall r\ a. Mutable\ r \Rightarrow List\ r\ a \rightarrow List\ r\ a \xrightarrow{Write\ r} () \end{aligned}$$

These operators both take a list and a new value. If the list contains an outer cons-cell, then the appropriate pointer in that cell is updated to point to the new value. If the list is not a cons, then a run-time error is raised.

3 Witnesses and Witness Construction

The novel aspect of our core language is that it uses dependently kinded witnesses to manage information about purity, constancy and mutability. A witness is a special type that can occur in the term being evaluated, and its occurrence guarantees a particular property of the program. The System-Fc [15] language uses a similar mechanism to manage information about non-syntactic type equality. Dependent kinds were introduced by the Edinburgh Logical Framework (LF) [1] which uses them to encode logical rules.

Note that although our formal operational semantics manipulates witnesses during reduction, in practice they are only used to reason about the program during compilation, and are not needed at runtime. Our compiler erases witnesses before code generation, along with all other type information.

3.1 Region Handles

We write witnesses with an underline, and the first we discuss are the region allocation witnesses $\underline{\rho}_n$. These are also called *region handles* and are introduced into the program with the **letregion** r **in** t expression. Reduction of this expression allocates a fresh handle $\underline{\rho}$ and substitutes it for all occurrences of the variable r in t . To avoid problems with variable capture we require all bound variables r in the initial program to be distinct. Although region handles are not needed at runtime, we can imagine them to be operational descriptions of physical regions of the store, perhaps incorporating a base address and a range. For example, the following program adds two to its argument, while storing an intermediate value in a region named r_3 .

$$\begin{aligned} \text{addTwo} &:: \forall r_1 r_2. \text{Int } r_1 \xrightarrow{\text{Read } r_1} \text{Int } r_2 \\ \text{addTwo} &= \lambda r_1 r_2. \lambda(x : \text{Int } r_1). \\ &\quad \mathbf{letregion } r_3 \mathbf{ in } \text{succ } r_3 r_2 (\text{succ } r_1 r_3 x) \end{aligned}$$

This program makes use of the primitive *succ* function that reads its integer argument and produces a new value into a given region:

$$\text{succ} :: \forall r_1 r_2. \text{Int } r_1 \xrightarrow{\text{Read } r_1} \text{Int } r_2$$

Note the phase distinction between region variables r_n and region handles $\underline{\rho}_n$. Region handles are bound by region variables. As no regions exist in the store before execution, region handles may not occur in the initial program. Also, although the outer call to *succ* reads a value in r_3 , this effect is not observable by calling functions, so is masked and not included in the type signature of *addTwo*. This is similar to the system of [17].

3.2 Witnesses of Constancy and Mutability

The constancy or mutability of values in a particular region is represented by the witnesses const $\underline{\rho}$ and mutable $\underline{\rho}$. Once again, these witnesses may not occur in the initial program. Instead, they are created with the *MkConst* and *MkMutable* witness type constructors which have the following kinds:

$$\begin{aligned} \mathit{MkConst} &:: \Pi(r : \%). \mathit{Const} \ r \\ \mathit{MkMutable} &:: \Pi(r : \%). \mathit{Mutable} \ r \end{aligned}$$

Both constructors take a region handle and produce the appropriate witness. To ensure that both $\underline{\text{const } \rho_n}$ and $\underline{\text{mutable } \rho_n}$ for the same ρ_n cannot be created by a given program, we require the mutability of a region to be set at the point it is introduced. We introduce new regions with **letregion**, so extend this construct with an optional witness binding that specifies the desired mutability. If a function accesses values in a given region, and does not possess either a witness of constancy or mutability for that region, then it cannot assume either. For example, the following function computes the length of a list by destructively incrementing a local accumulator, then copying out the final value.

$$\begin{aligned} \mathit{length} &:: \forall a \ r_1 \ r_2. \mathit{List} \ r_1 \ a \xrightarrow{\text{Read } r_1} \mathit{Int} \ r_2 \\ \mathit{length} &= \Lambda a \ r_1 \ r_2. \lambda(\mathit{list} : \mathit{List} \ r_1 \ a). \\ &\quad \mathbf{letregion} \ r_3 \ \mathbf{with} \ \{w = \mathit{MkMutable} \ r_3\} \ \mathbf{in} \\ &\quad \mathbf{let} \ (\mathit{acc} : \mathit{Int} \ r_3) = 0 \ r_3 \\ &\quad \quad (\mathit{length}' : \dots) \\ &\quad \quad = \lambda(xx : \mathit{List} \ r_1 \ a). \\ &\quad \quad \quad \mathbf{case} \ xx \ \mathbf{of} \\ &\quad \quad \quad \mathit{Nil} &\rightarrow \mathit{copyInt} \ r_3 \ r_2 \ \mathit{acc} \\ &\quad \quad \quad \mathit{Cons} \ _ \ xs &\rightarrow \mathbf{let} \ (_ : ()) = \mathit{incInt} \ r_3 \ w \ \mathit{acc} \\ &\quad \quad \quad &\quad \mathbf{in} \ \mathit{length}' \ xs \\ &\quad \mathbf{in} \ \mathit{length}' \ \mathit{list} \end{aligned}$$

where

$$\begin{aligned} \mathit{copyInt} &:: \forall r_1 \ r_2. \mathit{Int} \ r_1 \xrightarrow{\text{Read } r_1} \mathit{Int} \ r_2 \\ \mathit{incInt} &:: \forall r_1. \mathit{Mutable} \ r_1 \Rightarrow \mathit{Int} \ r_1 \xrightarrow{\text{Read } r_1 \vee \text{Write } r_1} () \end{aligned}$$

The set after the **with** keyword binds an optional witness type variable. If the region variable bound by the **letregion** is r , then the right of the witness binding must be either $\mathit{MkConst} \ r$ or $\mathit{MkMutable} \ r$. The type constructors $\mathit{MkConst}$ and $\mathit{MkMutable}$ may not occur elsewhere in the program. The length function above makes use of incInt which requires its integer argument to be in a mutable region, and we satisfy this constraint by passing it our witness to the fact.

3.3 Laziness and Witnesses of Purity

Although we use call-by-value evaluation as the default, we can suspend the evaluation of an arbitrary function application with the *suspend* operator:

$$\mathit{suspend} :: \forall a \ b \ e. \mathit{Pure} \ e \Rightarrow (a \xrightarrow{e} b) \rightarrow a \rightarrow b$$

$\mathit{suspend}$ takes a parameter function of type $a \xrightarrow{e} b$, its argument of type a , and defers the application by building a thunk at runtime. When the value of the thunk is demanded, the contained function will be applied to its argument, yielding a result of type b . As per [7], values are demanded when they are used as

the function in an application, or are inspected by a case-expression or primitive operator such as *update*. The constraint *Pure e* indicates that we must also provide a witness that the application to be suspended is observably pure. Witnesses of purity are written $\underline{\text{pure } \sigma}$ where σ is some effect. They can be created with the *MkPurify* witness type constructor. For example, the following function computes the successor of its argument, but only when the result is demanded:

$$\begin{aligned} \text{succL} &:: \forall r_1 r_2. \text{Const } r_1 \Rightarrow \text{Int } r_1 \xrightarrow{\text{Read } r_1} \text{Int } r_2 \\ \text{succL} &= \Lambda r_1 r_2 (w : \text{Const } r_1). \lambda(x : \text{Int } r_1). \\ &\quad \text{suspend } (\text{Int } r_1) (\text{Int } r_2) (\text{Read } r_1) (\text{MkPurify } r_1 w) \\ &\quad (\text{succ } r_1 r_2) x \end{aligned}$$

MkPurify takes a witness that a particular region is constant, and produces a witness proving that a read from that region is pure. It has the following kind:

$$\text{MkPurify} :: \Pi(r : \%). \text{Const } r \rightarrow \text{Pure } (\text{Read } r)$$

Reads of constant regions are pure because it does not matter when the read takes place, the same value will be returned each time. Note that in our system there are several ways of writing the effect of a pure function. As mentioned in §2.1 the effect term \perp is manifestly pure. However, we can also treat any other effect as pure if we can produce a witness of the appropriate kind. For example, *Read r₅* is pure if we can produce a witness of kind *Pure (Read r₅)*.

3.4 Witness Joining and Explicit Effect Masking

Purity constraints extend naturally to higher order functions. Here is the definition of a lazy map function, *mapL*, which constructs the first list element when called, but only constructs subsequent elements when they are demanded:

$$\begin{aligned} \text{mapL} &:: \forall a b r_1 r_2 e. \\ &\quad \text{Const } r_1 \Rightarrow \text{Pure } e \Rightarrow (a \xrightarrow{e} b) \rightarrow \text{List } r_1 a \rightarrow \text{List } r_2 b \\ \text{mapL} &= \Lambda a b r_1 r_2 e (w_1 : \text{Const } r_1) (w_2 : \text{Pure } e). \\ &\quad \lambda(f : a \xrightarrow{e} b) (\text{list} : \text{List } r_1 a). \\ &\quad \mathbf{mask} (\text{MkPureJoin } (\text{Read } r_1) e (\text{MkPurify } r_1 w_1) w_2) \mathbf{in} \\ &\quad \mathbf{case list of} \\ &\quad \quad \text{Nil} \quad \quad \rightarrow \text{Nil } r_2 \\ &\quad \quad \text{Cons } x \text{ } xs \rightarrow \text{Cons } r_2 b (f x) \\ &\quad \quad \quad (\text{suspend } (\text{List } r_1 a) (\text{List } r_2 b) \perp \text{MkPure} \\ &\quad \quad \quad (\text{mapL } a b r_1 r_2 e w_1 w_2 f) xs) \end{aligned}$$

The inner case-expression in this function has the effect *Read r₁ ∨ e*. The first part is due to inspecting the list constructors, and the second is due to the application of the argument function *f* to the list element *x*. However, as the recursive call to *mapL* is suspended, *mapL* itself must be pure. One way to satisfy this constraint would be to pass a witness showing that *Read r₁ ∨ e* is

pure directly to *suspend*. This works, but leaves *mapL* with a type that contains this (provably pure) effect term. Instead, we have chosen to explicitly mask this effect in the body of *mapL*. This gives *mapL* a manifestly pure type, and allows us to pass a trivial witness to *suspend* to show that the recursive call is pure.

The masking is achieved with the **mask** δ **in** t expression, which contains a witness of purity δ and a body t . The type and value of this expression is the same as for t , but its effect is the effect of t minus the terms which δ proves are pure. In our *mapL* example we prove that $\text{Read } r_1 \vee e$ is pure by combining two other witnesses, w_1 which proves that the list cells are in a constant region, and w_2 which proves that the argument function itself is pure. They are combined with the *MkPureJoin* witness type constructor which has the following kind:

$$\text{MkPureJoin} :: \Pi(e_1 : !). \Pi(e_2 : !). \text{Pure } e_1 \rightarrow \text{Pure } e_2 \rightarrow \text{Pure } (e_1 \vee e_2)$$

The masking expression in our example also uses *MkPure*, which introduces a witness that the effect \perp is pure. Note that our type for *mapL* now contains exactly the constraints that are *implicit* in a lazy language such as Haskell. In Haskell, all algebraic data is constant, and all functions are pure. In our language, we can suspend function applications as desired, but doing so requires the functions and data involved to satisfy the usual constraints of lazy evaluation.

4 Language

We are now in a position to formally define our core language and its typing rules. The structure of the language is given in Fig. 1. Most has been described previously, so we only discuss the aspects not covered so far. Firstly, we use \diamond as the result kind of witness kind constructors, so a constructor such as *Mutable* has kind $\text{Mutable} :: \% \rightarrow \diamond$. This says that a witness of kind *Mutable* r guarantees a property of a region, where \diamond refers to the guarantee.

We use τ_i as binders for value types, σ_i as binders for effect types, and δ_i as binders for type expressions that construct witness types. Δ_i refers to constructed witnesses of the form $\underline{\rho}$, $\underline{\text{const } \rho}$, $\underline{\text{mutable } \rho}$ or $\underline{\text{pure } \sigma}$. φ_i can refer to any type expression.

The values in our term language are identified with v . Weak values, v° , consist of the values as well as suspended function applications $\text{suspend } \overline{\varphi} v_1^\circ v_2^\circ$. A suspension is only forced when its (strong) value is demanded by using it as the function in an application, the discriminant of a case expression, or as an argument to a primitive operator such as *update*. Store locations l_i are discussed in §4.2. The other aspects of our term language are standard. Recursion can be introduced via **fix** in the usual way, but we omit it to save space. To simplify the presentation we require the alternatives in a case-expression to be exhaustive.

4.1 Typing Rules

In Fig. 2 the judgement form $\Gamma \vdash_{\kappa} \kappa :: \kappa'$ reads: with type environment Γ , kind κ has kind κ' . We could have added a super-kind stratum containing \diamond , but

Symbol Classes
 $a, r, e, w \rightarrow$ (type variable) $T \rightarrow$ (type constructor)
 $x \rightarrow$ (value variable) $K \rightarrow$ (data constructor)

Kinds
 $\kappa ::= \kappa \varphi \mid \Pi(a : \kappa_1). \kappa_2$ (kinds)
 $\mid * \mid \% \mid ! \mid \diamond$ (base kinds)
 $\mid \text{Const} \mid \text{Mutable} \mid \text{Pure}$ (kind constrs)

Types
 $\varphi, \tau, \sigma, \delta, \Delta$
 $::= a \mid \forall(a : \kappa). \tau \mid \varphi_1 \varphi_2 \mid (\rightarrow) \mid () \mid T$ (types)
 $\mid \sigma_1 \vee \sigma_2 \mid \perp \mid \text{Read} \mid \text{Write}$ (effects)
 $\mid \text{MkConst} \mid \text{MkMutable} \mid \text{MkPure} \mid \text{MkPurify} \mid \text{MkPureJoin}$ (witness constrs)
 $\mid \underline{\rho} \mid \text{const } \underline{\rho} \mid \text{mutable } \underline{\rho} \mid \text{pure } \underline{\sigma}$ (witnesses)

Terms
 $t ::= v \mid t \varphi \mid t_1 t_2 \mid \text{letregion } r \text{ with } \overline{\{w = \delta\}} \text{ in } t \mid K \overline{\varphi} \overline{t}$
 $\mid \text{case } t \text{ of } \overline{K \overline{x} : \overline{\tau}} \rightarrow t' \mid \text{update}_{K,i} \overline{\varphi} t_1 t_2 \mid \text{suspend } \overline{\varphi} t_1 t_2$
 $\mid \text{mask } \delta \text{ in } t$
 $v^\circ, u^\circ ::= v \mid \text{suspend } \overline{\varphi} v_1^\circ v_2^\circ$ (weak values)
 $v, u ::= x \mid l \mid () \mid \Lambda(a : \kappa). t \mid \lambda(x : \tau). t$ (values)

Derived Forms
 $\kappa_1 \rightarrow \kappa_2 \stackrel{\text{def}}{=} \Pi(_ : \kappa_1). \kappa_2$ $\text{let } (x : \tau) = t_1 \text{ in } t_2 \stackrel{\text{def}}{=} (\lambda(x : \tau). t_2) t_1$
 $\kappa \Rightarrow \tau \stackrel{\text{def}}{=} \forall(_ : \kappa). \tau$ $\text{letregion } r \text{ in } t \stackrel{\text{def}}{=} \text{letregion } r \text{ with } \emptyset \text{ in } t$
Store Typing
 $\Sigma ::= l : \tau \mid \underline{\rho} \mid \text{const } \underline{\rho} \mid \text{mutable } \underline{\rho}$
Type Environment
 $\Gamma ::= a : \kappa \mid x : \tau$
Fig. 1. Core Language

$$\boxed{\Gamma \vdash_{\mathbb{K}} \kappa :: \kappa'}$$

$$\frac{\kappa \in \{*, \%, !, \diamond\}}{\Gamma \vdash_{\mathbb{K}} \kappa :: \kappa} \text{ (KsRefl)} \quad \frac{\Gamma \vdash_{\mathbb{K}} \kappa_1 :: \kappa_{11} \rightarrow \kappa_{12} \quad \Gamma \vdash_{\mathbb{K}} \varphi :: \kappa_{11}}{\Gamma \vdash_{\mathbb{K}} \kappa_1 \varphi :: \kappa_{12}} \text{ (KsApp)}$$

$$\Gamma \vdash_{\mathbb{K}} \text{Const} :: \% \rightarrow \diamond \quad \Gamma \vdash_{\mathbb{K}} \text{Mutable} :: \% \rightarrow \diamond \quad \Gamma \vdash_{\mathbb{K}} \text{Pure} :: ! \rightarrow \diamond$$

Fig. 2. Kinds of Kinds

inspired by [12] we cap the hierarchy in this way to reduce the volume of typing rules.

In Fig. 3 the judgement form $\Gamma \mid \Sigma \vdash_{\mathbb{T}} \varphi :: \kappa$ reads: with type environment Γ and store typing Σ , type φ has kind κ . We discuss store typings in §4.3.

$$\boxed{\Gamma \mid \Sigma \vdash_{\mathsf{T}} \varphi :: \kappa}$$

$$\begin{array}{c}
\Gamma, a : \kappa \mid \Sigma \vdash_{\mathsf{T}} a :: \kappa \text{ (KiVar)} \\
\Gamma \mid \Sigma, \underline{\rho} \vdash_{\mathsf{T}} \underline{\rho} :: \% \text{ (KiHandle)} \quad \Gamma \mid \Sigma \vdash_{\mathsf{T}} \perp :: ! \text{ (KiBot)} \\
\frac{\Gamma \vdash_{\kappa} \kappa_1 :: \kappa'_1 \quad \Gamma, a : \kappa_1 \mid \Sigma \vdash_{\mathsf{T}} \tau :: \kappa_2 \quad a \notin \text{fv}(\Gamma)}{\Gamma \mid \Sigma \vdash_{\mathsf{T}} \forall(a : \kappa_1). \tau :: \kappa_2} \text{ (KiAll)} \quad \frac{\Gamma \mid \Sigma \vdash_{\mathsf{T}} \sigma_1 :: ! \quad \Gamma \mid \Sigma \vdash_{\mathsf{T}} \sigma_2 :: !}{\Gamma \mid \Sigma \vdash_{\mathsf{T}} \sigma_1 \vee \sigma_2 :: !} \text{ (KiJoin)} \\
\frac{\Gamma \mid \Sigma \vdash_{\mathsf{T}} \varphi_1 :: \Pi(a : \kappa_1). \kappa_2 \quad \Gamma \mid \Sigma \vdash_{\mathsf{T}} \varphi_2 :: \kappa_1}{\Gamma \mid \Sigma \vdash_{\mathsf{T}} \varphi_1 \varphi_2 :: \kappa_2[\varphi_2/a]} \text{ (KiApp)} \\
\Gamma \mid \Sigma, \underline{\text{const } \rho} \vdash_{\mathsf{T}} \underline{\text{const } \rho} :: \text{Const } \underline{\rho} \text{ (KiConst)} \\
\Gamma \mid \Sigma, \underline{\text{mutable } \rho} \vdash_{\mathsf{T}} \underline{\text{mutable } \rho} :: \text{Mutable } \underline{\rho} \text{ (KiMutable)} \\
\Gamma \mid \Sigma \vdash_{\mathsf{T}} \underline{\text{pure } \perp} :: \text{Pure } \perp \text{ (KiPure)} \\
\Gamma \mid \Sigma, \underline{\text{const } \rho} \vdash_{\mathsf{T}} \underline{\text{pure (Read } \rho)} :: \text{Pure (Read } \underline{\rho})} \text{ (KiPurify)} \\
\frac{\Gamma \mid \Sigma \vdash_{\mathsf{T}} \underline{\text{pure } \sigma_1} :: \text{Pure } \sigma_1 \quad \Gamma \mid \Sigma \vdash_{\mathsf{T}} \underline{\text{pure } \sigma_2} :: \text{Pure } \sigma_2}{\Gamma \mid \Sigma \vdash_{\mathsf{T}} \underline{\text{pure } (\sigma_1 \vee \sigma_2)} :: \text{Pure } (\sigma_1 \vee \sigma_2)} \text{ (KiPureJoin)} \\
\Gamma \mid \Sigma \vdash_{\mathsf{T}} (\rightarrow) \quad :: * \rightarrow * \rightarrow ! \rightarrow * \quad \Gamma \mid \Sigma \vdash_{\mathsf{T}} () \quad :: * \\
\Gamma \mid \Sigma \vdash_{\mathsf{T}} \text{Bool} \quad :: \% \rightarrow * \quad \Gamma \mid \Sigma \vdash_{\mathsf{T}} \text{Read} \quad :: \% \rightarrow ! \\
\Gamma \mid \Sigma \vdash_{\mathsf{T}} \text{MkConst} \quad :: \Pi(r : \%). \text{Const } r \quad \Gamma \mid \Sigma \vdash_{\mathsf{T}} \text{Write} \quad :: \% \rightarrow ! \\
\Gamma \mid \Sigma \vdash_{\mathsf{T}} \text{MkMutable} \quad :: \Pi(r : \%). \text{Mutable } r \quad \Gamma \mid \Sigma \vdash_{\mathsf{T}} \text{MkPure} \quad :: \text{Pure } \perp \\
\Gamma \mid \Sigma \vdash_{\mathsf{T}} \text{MkPurify} \quad :: \Pi(r : \%). \text{Const } r \rightarrow \text{Pure (Read } r) \\
\Gamma \mid \Sigma \vdash_{\mathsf{T}} \text{MkPureJoin} \quad :: \Pi(e_1 : !). \Pi(e_2 : !). \text{Pure } e_1 \rightarrow \text{Pure } e_2 \rightarrow \text{Pure } (e_1 \vee e_2)
\end{array}$$

Fig. 3. Kinds of Types

In Fig. 4 the judgement form $\Gamma \mid \Sigma \vdash t :: \tau; \sigma$ reads: with type environment Γ and store typing Σ , term t has type τ and effect σ . In `TyLetRegion` the premise “ $\overline{\delta}_i$ well formed” refers to the requirement discussed in §3.2 that the witness introduced by a **letregion** must concern the bound variable r . In `TyUpdate` and `TyAlt`, the meta-function `ctorTypes(T)` returns a set containing the types of the data constructors associated with type constructor T .

4.2 Dynamic Semantics

During evaluation, all updatable data is held in the store (also known as the heap), which is defined in Fig. 5. The store contains bindings that map abstract store locations to store objects. Each store object consists of a constructor tag

$$\boxed{\Gamma \mid \Sigma \vdash t :: \tau; \sigma}$$

$$\begin{array}{c}
 \Gamma, x: \tau \mid \Sigma \vdash x :: \tau; \perp \text{ (TyVar)} \quad \Gamma \mid \Sigma, l: \tau \vdash l :: \tau; \perp \text{ (TyLoc)} \\
 \Gamma, K: \tau \mid \Sigma \vdash K :: \tau; \perp \text{ (TyCtor)} \quad \Gamma \mid \Sigma \vdash () :: (); \perp \text{ (TyUnit)} \\
 \\
 \frac{\Gamma, a: \kappa \mid \Sigma \vdash t_2 :: \tau_2; \sigma_2}{\Gamma \mid \Sigma \vdash \Lambda(a: \kappa). t_2 :: \forall(a: \kappa). \tau_2; \sigma_2} \text{ (TyAbsT)} \\
 \\
 \frac{\Gamma \mid \Sigma \vdash t_1 :: \forall(a: \kappa_{11}). \varphi_{12}; \sigma_1 \quad \Gamma \mid \Sigma \vdash_{\text{T}} \varphi_2 :: \kappa_{11}}{\Gamma \mid \Sigma \vdash t_1 \varphi_2 :: \varphi_{12}[\varphi_2/a]; \sigma_1[\varphi_2/a]} \text{ (TyAppT)} \\
 \\
 \frac{\Gamma, x: \tau_1 \mid \Sigma \vdash t :: \tau_2; \sigma}{\Gamma \mid \Sigma \vdash \lambda(x: \tau_1). t :: \tau_1 \xrightarrow{\sigma} \tau_2; \perp} \text{ (TyAbs)} \\
 \\
 \frac{\Gamma \mid \Sigma \vdash t_2 :: \tau_{11}; \sigma_2 \quad \Gamma \mid \Sigma \vdash t_1 :: \tau_{11} \xrightarrow{\sigma} \tau_{12}; \sigma_1}{\Gamma \mid \Sigma \vdash t_1 t_2 :: \tau_{12}; \sigma_1 \vee \sigma_2 \vee \sigma} \text{ (TyApp)} \\
 \\
 \frac{\overline{\delta_i} \text{ well formed} \quad r \notin fv(\tau) \quad \Gamma \vdash_{\text{K}} \kappa_i :: \diamond \quad \Gamma, r: \%, \overline{w_i: \kappa_i} \mid \Sigma \vdash t :: \tau; \sigma \quad \Gamma \mid \Sigma \vdash_{\text{T}} \delta_i :: \kappa_i}{\Gamma \mid \Sigma \vdash \mathbf{letregion } r \text{ with } \{\overline{w_i = \delta_i}\} \text{ in } t :: \tau; \sigma \setminus (\text{Read } r \vee \text{Write } r)} \text{ (TyLetRegion)} \\
 \\
 \frac{\Gamma \mid \Sigma \vdash t :: T \varphi \overline{\varphi'}; \sigma \quad \overline{\Gamma \mid \Sigma \vdash p_i \rightarrow t_i :: T \varphi \overline{\varphi'} \rightarrow \tau; \sigma'_i}^n}{\Gamma \mid \Sigma \vdash \mathbf{case } t \text{ of } \overline{p \rightarrow t} :: \tau; \sigma \vee \text{Read } \varphi \vee \sigma'_0 \vee \sigma'_1 \dots \vee \sigma'_n} \text{ (TyCase)} \\
 \\
 \frac{\Gamma \mid \Sigma \vdash_{\text{T}} \delta :: \text{Mutable } \varphi \quad \Gamma \mid \Sigma \vdash t' :: \tau_i[\varphi/r][\overline{\varphi'/a}]; \sigma' \quad \Gamma \mid \Sigma \vdash t :: T \varphi \overline{\varphi'}; \sigma \quad K :: \forall(r: \%)(\overline{a: \kappa}). \overline{\tau} \rightarrow T r \overline{a} \in \text{ctorTypes}(T)}{\Gamma \mid \Sigma \vdash \mathbf{update}_{\text{K}, i} \varphi \overline{\varphi'} \delta t t' :: (); \sigma \vee \sigma' \vee \text{Write } \varphi} \text{ (TyUpdate)} \\
 \\
 \frac{\Gamma \mid \Sigma \vdash_{\text{T}} \delta :: \text{Pure } \sigma \quad \Gamma \mid \Sigma \vdash t_1 :: \tau_{11} \xrightarrow{\sigma} \tau_{12}; \sigma_1 \quad \Gamma \mid \Sigma \vdash t_2 :: \tau_{11}; \sigma_2}{\Gamma \mid \Sigma \vdash \mathbf{suspend } \tau_{11} \tau_{12} \sigma \delta t_1 t_2 :: \tau_{12}; \sigma_1 \vee \sigma_2} \text{ (TySuspend)} \\
 \\
 \frac{\Gamma \mid \Sigma \vdash t :: \tau; \sigma \quad \Gamma \mid \Sigma \vdash_{\text{T}} \delta :: \text{Pure } \sigma'}{\Gamma \mid \Sigma \vdash \mathbf{mask } \delta \text{ in } t :: \tau; \sigma \setminus \sigma'} \text{ (TyMaskPure)} \\
 \\
 \boxed{\Gamma \mid \Sigma \vdash p \rightarrow t :: \tau \rightarrow \tau'; \sigma} \\
 \\
 \frac{K :: \forall(r: \%)(\overline{a: \kappa}). \overline{\tau} \rightarrow T r \overline{a} \in \text{ctorTypes}(T) \quad \Gamma, x: \theta(\tau) \mid \Sigma \vdash t :: \tau'; \sigma \quad \theta = [\varphi/r \overline{\varphi'/a}]}{\Gamma \mid \Sigma \vdash K \overline{x} \rightarrow t :: T \varphi \overline{\varphi'} \rightarrow \tau'; \sigma} \text{ (TyAlt)}
 \end{array}$$

Fig. 4. Types of Terms

l	\rightarrow	(store location)	
ρ	\rightarrow	(region handle)	
o	$::=$	$\rho \mid \text{const } \rho \mid \text{mutable } \rho$	(property)
π	$::=$	$l \mid () \mid \Lambda(a : \kappa).t \mid \lambda(x : \tau).t \mid \text{suspend } \overline{\varphi} \pi \pi'$	(store value)
μ	$::=$	$C_K \overline{\pi}$	(store object)
H	$:$	$\{ l \xrightarrow{\rho} \mu \} + \{ o \}$	(store)

Fig. 5. Stores and Store Objects

C_K and a list of store values $\overline{\pi}$, where each value can be a location, unit value, abstraction or suspension. Each binding is annotated with a region handle ρ that specifies the region that the binding belongs to. Note that store *objects* can be usefully updated, but store *values* can not.

The store also contains *properties* that specify how bindings in the various regions may be used. The properties are ρ , (const ρ) and (mutable ρ). The last two indicate whether a binding in that region may be treated as constant, or updated. When used as a property, a region handle ρ indicates that the corresponding region has been created and is ready to have bindings allocated into it. Note that the region handles of store bindings and properties are not underlined because those occurrences are not used as types.

In Fig. 6 the judgement form $H; \delta \rightsquigarrow \delta'$ reads: with store H , witness δ produces witness δ' . Operationally, properties can be imagined as protection flags on regions of the store — much like the read, write and execute bits in a hardware page table. The witness constructors *MkConst* and *MkMutable* test for these properties, producing a type-level artefact showing that the property was set. If we try to evaluate either constructor when the desired property is *not* set, then the evaluation becomes stuck.

In Fig. 7 the judgement form $H; t \longrightarrow H'; t'$ reads: in heap H term t reduces to a new heap H' and term t' . In *EvLetRegion* the *propOf* meta-function maps a witness to its associated store property. Also, note that the premise of *EvLetRegion* is always true, and produces the required witnesses and properties from the given witness constructions $\overline{\delta}_i$.

4.3 Soundness

In the typing rules we use a store typing Σ that models the state of the heap as the program evaluates. The store typing contains the type of each store location, along with witnesses to the current set of store properties. We say the store typing *models* the store, and write $\Sigma \models H$, when all members of the store typing correspond to members of the store. Conversely, we say the store is *well typed*, and write $\Sigma \vdash H$ when it contains all the bindings and properties predicted by the store typing. Both the store and store typing grow as the program evaluates, and neither bindings, properties or witnesses are removed once added.

Store bindings can be modified by the *update* operator, but the typing rules for update ensure that bindings retain the types predicted by the store typing. The rule *TyLoc* of Fig. 4 and *KiHandle*, *KiConst*, *KiMutable* and *KiPurify* of

$$\begin{array}{c}
 \frac{H ; \delta \rightsquigarrow \delta'}{H ; E_w[\delta] \rightsquigarrow E_w[\delta']} \qquad E_w ::= [] \mid \text{MkPurify } \rho \ E_w \\
 \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \mid \text{MkPureJoin } \sigma_1 \ \sigma_2 \ E_w \ E_w \\
 \\
 H[\text{const } \rho] ; \text{MkConst } \underline{\rho} \rightsquigarrow \underline{\text{const } \rho} \qquad \text{(EwConst)} \\
 H[\text{mutable } \rho] ; \text{MkMutable } \underline{\rho} \rightsquigarrow \underline{\text{mutable } \rho} \qquad \text{(EwMutable)} \\
 H ; \text{MkPure} \rightsquigarrow \underline{\text{pure } \perp} \qquad \text{(EwPure)} \\
 H ; \text{MkPurify } \underline{\rho} \ \text{const } \underline{\rho} \rightsquigarrow \underline{\text{pure (Read } \rho)} \qquad \text{(EwPurify)} \\
 H ; \text{MkPureJoin } \sigma_1 \ \sigma_2 \ \underline{\text{pure } \sigma_1} \ \underline{\text{pure } \sigma_2} \rightsquigarrow \underline{\text{pure } (\sigma_1 \vee \sigma_2)} \qquad \text{(EwPureJoin)}
 \end{array}$$

Fig. 6. Witness Construction

$$\begin{array}{c}
 \frac{e \longrightarrow e'}{E_v[e] \longrightarrow E_v[e']} \qquad E_v ::= [] \mid E_v \ \varphi \mid E_v \ t_2 \mid v \ E_v \mid \mathbf{case} \ E_v \ \mathbf{of} \ \overline{alt} \\
 \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \mid K \ \overline{\varphi} \ E_v \ t_1 \ \dots \mid K \ \overline{\varphi} \ v_0 \ E_v \ \dots \mid \dots \\
 \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \mid \text{update}_{K,i} \ \overline{\varphi} \ E_v \ t_2 \mid \text{update}_{K,i} \ \overline{\varphi} \ l \ E_v \\
 \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \mid \text{suspend} \ \overline{\varphi} \ E_v \ t_2 \mid \text{suspend} \ \overline{\varphi} \ v \ E_v \\
 \\
 H ; (\Lambda(a :: \kappa). t) \ \varphi \longrightarrow H ; t[\varphi/a] \qquad \text{(EvTAppAbs)} \\
 H ; (\lambda(x :: \tau). t) \ v^\circ \longrightarrow H ; t[v^\circ/x] \qquad \text{(EvAppAbs)} \\
 \\
 \frac{H, \overline{\text{propOf}(\Delta_i)} ; \overline{\delta_i[\rho/r]} \rightsquigarrow \overline{\Delta_i} \quad \rho \ \text{fresh}}{H ; \mathbf{letregion} \ r \ \mathbf{with} \ \{\overline{w_i} = \overline{\delta_i}\} \ \mathbf{in} \ t \longrightarrow H, \rho, \overline{\text{propOf}(\Delta_i)} ; t[\overline{\Delta_i/w_i}][\rho/r]} \qquad \text{(EvLetRegion)} \\
 \\
 H[\rho] ; K \ \underline{\rho} \ \overline{\varphi} \ \overline{v^\circ} \longrightarrow H, l \stackrel{\rho}{\mapsto} C_K \ \overline{v^\circ} ; l \quad l \ \text{fresh} \qquad \text{(EvAlloc)} \\
 H[l \stackrel{\rho}{\mapsto} C_K \ \overline{v^\circ}] ; \mathbf{case} \ l \ \mathbf{of} \ \dots K \ \overline{x} \rightarrow t \dots \longrightarrow H ; t[\overline{v^\circ}/\overline{x}] \qquad \text{(EvCase)} \\
 \\
 \frac{H ; \delta \rightsquigarrow \delta'}{H ; \text{update} \ \overline{\varphi} \ \delta \ t \ t' \longrightarrow H ; \text{update} \ \overline{\varphi} \ \delta' \ t \ t'} \qquad \text{(EvUpdateW)} \\
 \\
 H[\text{mutable } \rho], l \stackrel{\rho}{\mapsto} C_K \ \overline{v^\circ} ; \text{update}_{K,i} \ \overline{\varphi} \ \underline{\text{mutable } \rho} \ l \ u^\circ \qquad \text{(EvUpdate)} \\
 \longrightarrow H, l \stackrel{\rho}{\mapsto} C_K \ v_0..u_i^\circ..v_n ; () \\
 \\
 H[\text{mutable } \rho], l \stackrel{\rho}{\mapsto} C_K \ \overline{v^\circ} ; \text{update}_{K',i} \ \overline{\varphi} \ \underline{\text{mutable } \rho} \ l \ u^\circ \qquad \text{(EvFail)} \\
 \longrightarrow H ; \text{fail} \qquad K \neq K' \\
 \\
 \frac{H ; \delta \rightsquigarrow \delta'}{H ; \text{suspend} \ \overline{\varphi} \ \delta \ t \ t' \longrightarrow H ; \text{suspend} \ \overline{\varphi} \ \delta' \ t \ t'} \qquad \text{(EvSuspendW)} \\
 \\
 H ; \text{suspend} \ \tau \ \tau' \ \sigma \ \underline{\text{pure } \sigma} \ (\lambda(x : \tau). t) \ v^\circ \longrightarrow H ; t[v^\circ/x] \qquad \text{(EvSuspend)} \\
 \\
 H ; \mathbf{mask} \ \delta \ \mathbf{in} \ t \longrightarrow H ; t \qquad \text{(EvMask)}
 \end{array}$$

Fig. 7. Term Evaluation

Fig. 3 ensure that if a location or witness occurs in the term, then it also occurs in the store typing. Provided the store typing models the store, this also means that the corresponding binding or property is present in the store. From the evaluation rules in Fig. 7, the only term that adds properties to the store is **letregion**, and when it does, it also introduces the corresponding witnesses into the expression. The well-formedness restriction on **letregion** guarantees that a witnesses of mutability and constancy for the same region cannot be created. This ensures that if we have, say, the witness $\text{const } \rho$ in the term, then there is *not* a (mutable ρ) property in the store. This means that bindings in those regions can never be updated, and it is safe to suspend function applications that read them.

Our progress and preservation theorems are stated below. We do not prove these here, but [10] contains a proof for a similar system. The system in this paper supports full algebraic data types, whereas the one in [10] is limited to booleans. Also, here we include a $r \notin \text{fv}(\tau)$ premise in the TyLetRegion rule, which makes the presentation easier. See [10] for a discussion of this point.

Progress. If $\emptyset \mid \Sigma \vdash t :: \tau; \sigma$ and $\Sigma \models H$ and $\Sigma \vdash H$ and $\text{nofab}(t)$ then either $t \in \text{Value}$ or for some H', t' we have $(H; t \longrightarrow H'; t' \text{ and } \text{nofab}(t'))$ or $H; t \longrightarrow H'; \text{fail}$).

Preservation. If $\Gamma \mid \Sigma \vdash t :: \tau; \sigma$ and $H; t \longrightarrow H'; t'$ then for some Σ', σ' we have $\Gamma \mid \Sigma' \vdash t' :: \tau; \sigma'$ and $\Sigma' \supseteq \Sigma$ and $\Sigma' \models H'$ and $\Sigma' \vdash H'$ and $\Gamma \mid \Sigma \vdash \sigma' \sqsubseteq \sigma$.

In the Progress Theorem, “nofab” is short for “no fabricated region witnesses”, and refers to the syntactic constraint that *MkConst* and *MkMutable* may only appear in the witness binding of a **letregion** and not elsewhere in the program. We could perhaps recast these two constructors as separate syntactic forms of **letregion**, and remove the need for nofab, but we have chosen not to do this because we prefer the simpler syntax.

In the Preservation Theorem, note that the latent effect of the term reduces as the program progresses. The \sqsubseteq relationship on effects is defined in the obvious way, apart from the following extra rule:

$$\frac{\Gamma \mid \Sigma \vdash_{\tau} \delta :: \text{Pure } \sigma}{\Gamma \mid \Sigma \vdash \sigma \sqsubseteq \perp} \text{ (SubPurify)}$$

This says that if we can construct a witness that a particular effect is pure, then we can treat it as such. This allows us to erase read effects on constant regions during the proof of Preservation. It is needed to show that forcing a suspension does not have a visible effect, and that we can disregard explicitly masked effect terms when entering into the body of a mask-expression.

5 Related Work

The inspiration for our work has been to build on the monadic intermediate languages of [18], [3] and [13]. Note that for our purposes, the difference between

using effect and monadic typing is largely syntactic. We prefer effect typing because it mirrors our operational intuition more closely, but [19] gives a translation between the two. Our system extends the previous languages with region, effect and mutability polymorphism, which improves the scope of optimisations that can be performed. In [4] and [2], Benton *et al* present similar monadic languages that include region and effect polymorphism, but do not consider mutability polymorphism or lazy evaluation.

The Capability Calculus [5] provides region based memory management, whereby a capability is associated with each region, and an expression can only access a region when it holds its capability. When the region is deallocated, its associated capability is revoked, ensuring soundness. The capabilities of [5] have similarities to the witnesses of our system, but theirs are not reified in the term being evaluated, and we do not allow ours to be revoked.

The BitC [14] language permits any location, whether on the stack, heap or within data structures to be mutated. Its operational semantics includes an explicit stack as well as a heap, and function arguments are implicitly copied onto the stack during application. BitC includes mutability annotations, but does not use region or effect typing.

6 Conclusion and Future Work

We have presented a System-F style intermediate language that supports mutability polymorphism as well as lazy evaluation, and uses dependently kinded witnesses to track the purity of effects and the mutability and constancy of regions. One of the current limitations of our system is that the results of all case alternatives must have the same type. This prevents us from choosing between, say, a mutable and a constant integer. In future work we plan to provide a new region constraint *Blocked* that represents the fact that an object could be in either a mutable or constant region. We would permit such objects to be read, but not updated, and computations that read them could not be suspended. Doing so would likely require introducing a notion of subtyping into the system, so the types of all alternatives could be coerced to a single upper bound.

The system presented in this paper has been implemented in the prototype Disciplined Disciple Compiler (DDC) which can be obtained from the `haskell.org` website.

References

1. Avron, A., Honsell, F., Mason, I.A.: An overview of the Edinburgh Logical Framework. In: Current Trends in Hardware Verification and Automated Theorem Proving, pp. 323–340. Springer, Heidelberg (1989)
2. Benton, N., Buchlovsky, P.: Semantics of an effect analysis for exceptions. In: Proc. of TLDI 2007, pp. 15–26. ACM, New York (2007)
3. Benton, N., Kennedy, A.: Monads, effects and transformations. Electronic Notes in Theoretical Computer Science, pp. 1–18. Elsevier, Amsterdam (1999)

4. Benton, N., Kennedy, A., Beringer, L., Hofmann, M.: Relational semantics for effect-based program transformations with dynamic allocation. In: Proc. of PPDP, pp. 87–96. ACM, New York (2007)
5. Crary, K., Walker, D., Morrisett, G.: Typed memory management in a calculus of capabilities. In: Proc. of POPL, pp. 262–275. ACM, New York (1999)
6. de Medeiros Santos, A.L.: Compilation by Transformation in Non-Strict Functional Languages. PhD thesis, University of Glasgow (1995)
7. Launchbury, J.: A natural semantics for lazy evaluation. In: Proc. of POPL, pp. 144–154. ACM, New York (1993)
8. Launchbury, J., Jones, S.P.: Lazy functional state threads. In: Proc. of PLDI, pp. 24–35. ACM, New York (1994)
9. Leroy, X., Doligez, D., Garrigue, J., Rémy, D., Vouillon, J.: The Objective Caml system, release 3.11, documentation and user’s manual. Technical report, INRIA (2008)
10. Lippmeier, B.: Type Inference and Optimisation for an Impure World. PhD thesis, Australian National University (June 2009) (submitted)
11. MacQueen, D.B.: Standard ML of New Jersey. In: Małuszyński, J., Wirsing, M. (eds.) PLILP 1991. LNCS, vol. 528, pp. 1–13. Springer, Heidelberg (1991)
12. Jones, S.P., Meijer, E.: Henk: a typed intermediate language. In: Proc. of the Workshop on Types in Compilation (1997)
13. Jones, S.P., Shields, M., Launchbury, J., Tolmach, A.: Bridging the gulf: a common intermediate language for ML and Haskell. In: Proc. of POPL, pp. 49–61. ACM, New York (1998)
14. Shapiro, J., Sridhar, S., Doerrie, S.: BitC language specification. Technical report, The EROS Group and Johns Hopkins University (2008)
15. Sulzmann, M., Chakravarty, M.M.T., Jones, S.P., Donnelly, K.: System-F with type equality coercions. In: Proc. of TLDI. ACM, New York (2007)
16. Talpin, J.-P., Jouvelot, P.: The type and effect discipline. In: Proc. of Logic in Computer Science, pp. 162–173. IEEE, Los Alamitos (1992)
17. Tofte, M., Birkedal, L., Elsmann, M., Hallenberg, N., Olesen, T.H., Sestoft, P.: Programming with regions in the MLKit (revised for version 4.3.0). Technical report, IT University of Copenhagen, Denmark (January 2006)
18. Tolmach, A.: Optimizing ML using a hierarchy of monadic types. In: Leroy, X., Ogori, A. (eds.) TIC 1998. LNCS, vol. 1473, pp. 97–115. Springer, Heidelberg (1998)
19. Wadler, P., Thiemann, P.: The marriage of effects and monads. *ACM Trans. Computation and Logic* 4(1), 1–32 (2003)