

Do Be Do Be Do

Sam Lindley

The University of Edinburgh

Conor McBride

University of Strathclyde

Abstract

We explore the design of Frank, a strict functional programming language designed from the ground up around a novel variant of Plotkin and Pretnar’s *effect handler* abstraction. Inspired by Levy’s call-by-push-value calculus, Frank makes an explicit distinction between computation and value types. Frank combines the advantages of call-by-push-value, call-by-value, and effect typing.

1. Introduction

Plotkin and Power’s algebraic effects [19–22], in conjunction with Plotkin and Pretnar’s handlers for algebraic effects [23, 24], provide a compelling foundation for effectful programming. By separating effect signatures from their implementation, algebraic effects provide a high degree of modularity, allowing programmers to express effectful programs independently of the concrete interpretation of their effects. A handler is an interpretation of the effects of an algebraic computation.

Effect handlers in Frank generalise functions. An effect handler acts as an interpreter for a specified set of commands whose signatures are statically tracked by the type system. A function is the special case of an effect handler whose command set is empty.

The contributions of this paper are:

- Frank, a strict functional programming language featuring a bidirectional effect type system, effect polymorphism, and effect handlers.
- A novel approach to effect polymorphism which avoids all mention of effect variables, crucially relying on the observation that one must always instantiate the effects of a function being applied with the current *ambient* effects.
- The combination of pattern matching and effect handlers in such a way that incomplete or ambiguous pattern matching can be realised as concrete effects that can be handled however the programmer chooses using further effect handlers.
- Multi-handlers as both an abstraction for handling multiple computations over different effect sets simultaneously and a characterisation of effect-handlers as generalised functions.
- A description of pattern matching compilation from Frank into a core language, *Core Frank*.

- A translation from Core Frank into F_{eff} , a polymorphic variant of the λ_{eff} calculus of Kammar et al [8], which in turn extends Levy’s call-by-push-value calculus [11].
- A straightforward small-step operational semantics for F_{eff} , yielding, in combination with the translation to F_{eff} , a type soundness result for Frank.

A number of other languages and libraries are built around effect handlers and algebraic effects.

Bauer and Pretnar’s Eff [2]. A significant difference between Frank and the original version of Eff [2] is that the latter provides no support for effect typing. Recently Bauer and Pretnar have designed an effect type system for Eff [3]. Their implementation [25] supports Hindley-Milner type inference, and the type system incorporates effect sub-typing. In contrast, Frank uses bidirectional type inference, and avoids sub-typing altogether.

Handlers in action [8]. In previous work with Kammar and Oury [8], the first author designed and experimented with a number of effect handler libraries for languages ranging from Racket, to SML, to Haskell. Apart from the Haskell library, these other libraries have no effect typing support. The Haskell library takes advantage of type classes to simulate an effect type system not entirely dissimilar to that of Frank. As Haskell is lazy, the Haskell library cannot be used to write direct-style effectful programs – one must instead adopt a monadic style. Furthermore, although there are a number of ways of almost simulating effect type systems in Haskell, none is without its flaws. Kiselyov et al [9] have designed another Haskell library for effect handlers, making a different collection of design choices.

Brady [4] has designed a library and DSL for programming with effects in his dependently typed Idris language. Like the Haskell libraries, Brady’s library currently requires the programmer to write effectful code in a monadic style.

In Plotkin and Powers setting, one defines algebraic effects with respect to an equational theory. In all of the above implementations, and in Frank, the equational theory is taken to be the free theory, in which there are no equations.

The second author has been plotting Frank since 2007 [14]. He has implemented a prototype of a previous version of Frank [16]. The design described in the current paper has much in common with that implementation, but there are some syntactic and semantic differences. The most important change in the current design is the introduction of multi-handlers as a generalisation of both functions and handlers.

The rest of the paper is structured as follows. Section 2 introduces Frank by example. Section 3 presents a type system for Frank. Section 4 describes how to elaborate multi-handlers and pattern matching into Core Frank, a language of plain call-by-value functions, explicit case analysis and unary handler constructs. Section 5 gives a call-by-value embedding of Core Frank into F_{eff} , a variant of Kammar et al’s λ_{eff} calculus with shallow handlers, explicit polymorphism and general recursion. Section 6 gives a se-

mantics for F_{eff} , which when composed with pattern matching and the call-by-value embedding, yields a semantics for Frank. Section 7 outlines related work and Section 8 discusses future work.

2. Introducing Frank

Frank is a functional programming language with effects and handlers in the style of Eff controlled by a type system inspired by Levy’s call-by-push-value [11]. Doing and Being are clearly separated, and managed by distinguished notions of computation and value types.

Concrete values live in inductive datatypes.

```
data List X
  = nil
  | cons X (List X)
```

```
data Zero =
```

```
data Unit = unit
```

We can write perfectly ordinary functional programs, with (compulsory) type signatures.

```
append : List X -> List X -> List X
append nil      ys = ys
append (x :: xs) ys = cons x (append xs ys)
```

Higher-order functions are passed suspended computations. Braces are “suspenders”.

```
map : {X -> Y} -> List X -> List Y
map f nil      = nil
map f (cons x xs) = cons (f x) (map f xs)
```

A value type V is a datatype $D \ V_1 \ \dots \ V_n$, a suspended computation type $\{C\}$, or a type variable X . A computation type can be a function type $V \rightarrow C$ or an effect-annotated value type $[S_1 \ \dots \ S_n]V$, where an empty bracket may be omitted.

Effects are collections of signatures, which describe a choice of commands. Here are some simple signatures.

```
sig Send X
  = send : Unit(X)
```

```
sig Receive X
  = receive : X
```

```
sig Abort
  = aborting : Zero
```

The `send` command takes an argument of type X and returns a value of type `Unit`. The `receive` command returns a value of type X . The `abort` command returns an element of the empty type, `Zero`.

Frank is a call-by-value language, but it naturally distinguishes $[]V$, the type of suspended pure computations which deliver a V , from V itself. We can thus define a kind of “semicolon” just as the function which ignores its first argument.

```
semi : X Y -> Y
semi x y = y
```

Frank has effect polymorphism, enough to allow higher-order functions to pass effect permissions to their parameters. The following uses `map` to send a list of things, one at a time.

```
sends : List X -> [Send X]Unit
sends xs = semi (map send xs) unit
```

The reason this type checks at all is because `map` is implicitly polymorphic in its effects.

The following does not typecheck, because the `Send` effect is not permitted in the return type of `bad`.

```
bad : List X -> Unit
bad xs = semi (map send xs) unit
```

Writing control operators is not too tricky.

```
bind : X -> {X -> Y} -> Y
bind x f = f x
```

We can use `bind` to define a polymorphic `abort` function.

```
abort : [Abort]X
abort = bind aborting! {}
```

The term `{}` denotes a suspended computation containing an empty collection of pattern matching clauses covering the `Zero` return type of `aborting`.

Here is a computation which receives and concatenates lists until one is empty.

```
catter : [Receive (List X)]List X
catter = bind receive!
        { nil -> nil
        | xs -> append xs catter!
        }
```

The command `receive` is a suspended computation of type $\{[Receive (List X)]List X\}$ that delivers a list when forced. The notation `receive!` forces a thunk.

If f is a suspended computation of function type, then $f \ x$ is syntactic sugar for $f! \ x$. Thus we need only explicitly force computations that take no arguments.

Effects are handled by special functions called *effect handlers*. Effect handlers in Frank can take multiple computations as arguments, hence they are *multi-handlers*. In fact standard functions like the ones we have seen so far are just special cases of multi-handlers in which the all the arguments are pure.

A multi-handler has type $R_1 \rightarrow \dots \rightarrow R_m \rightarrow R$ where each R_i and R is an effect-annotated value type. For instance, we can write a pipe multi-handler which handles `send` commands from one computation by matching them against corresponding `receive` commands from another.

```
pipe : [Send X]Unit -> [Receive X]Y -> [Abort]Y
pipe _      y      = y
pipe unit  _      = abort!
pipe (send x ? s) (receive ? r) =
  pipe (s unit) (r x)
```

The type signature conveys several different things. The `pipe` handler must handle all `Send X` commands in its first argument and all `Receive X` commands in its second argument. The first argument returns values of type `Unit` and the second argument returns values of type Y . The handler itself is allowed to perform `Abort` commands and returns a final value of type Y .

Here are some things to send.

```
hello : List Char
hello =
  cons 'h' (cons 'e' (cons 'l'
    (cons 'l' (cons 'o' nil))))
```

```
space : List Char
space = cons ' ' nil
```

```
world : List Char
world =
  cons 'w' (cons 'o' (cons 'r'
    (cons 'l' (cons 'd' nil))))
```

Types	
(values)	$U, V ::= D \bar{U} \mid \{C\} \mid X$
(computations)	$C ::= R \mid R \rightarrow C$
(returners)	$R ::= [\Sigma]V$
(quantifiers)	$Z ::= X \mid \varepsilon$
(polytypes)	$P ::= \forall \bar{Z}. Q$
(thunks)	$Q ::= \{C\}$
(signatures)	$S \bar{X} ::= \cdot \mid c : U(\bar{V}), S \bar{X}$
(effects)	$\Sigma ::= \emptyset \mid \Sigma, S \bar{V} \mid \varepsilon$
(type environments)	$\Gamma ::= \cdot \mid \Gamma, x : V \mid f : P$
Terms	
(inferable values)	$u ::= x \mid f \mid c \mid d$
(checkable values)	$v ::= u \mid k \bar{v} \mid \{e\}$
(inferable computations)	$d ::= u! \mid d e$
(checkable computations)	$e ::= v \mid \bar{r} \mapsto e \mid () \mid e \mid e$
(value patterns)	$p ::= x \mid k \bar{p}$
(computation patterns)	$r ::= p \mid c \bar{p} ? g \mid x!$

Figure 1. Frank Syntax

Here is a computation which sends them.

```
sender : [Send (List Char)]Unit
sender =
  sends (cons (hello
              (cons space (cons world nil))))
```

Here is a main function, which plugs sender and catter together and sends their output to the console.

```
main : [Abort, Console](List Unit)
main = map ouch (pipe sender catter)
```

where the Console operations are handled specially at the top-level according to the following signature.

```
sig Console
= inch : Char(Unit)
  | ouch : Unit(Char)
```

The type system does two separate things:

- It ensures that value types coincide.
- It ensures that effects required are included in effects enabled.

The fun of Frank is that one can say what it is to *be* a computation without saying what it is to *do* it. Doing and being are separately negotiable, and readily interleaved in different ways. Or as Frank Sinatra put it,

do be do be do

3. Type System

In this section we give a formal presentation of Frank's type system. The syntax of Frank types and terms is given in Figure 1. The types are divided into value types and computation types in a similar fashion to Levy's call-by-push-value calculus [11]. Value types are datatypes ($D \bar{U}$), suspended C computations ($\{C\}$), otherwise known as *thunks*, or type variables (X).

Computation types are constructed from *returners*. A returner ($[\Sigma]V$) represents the type of a computation that returns values of type V while performing effects in Σ . In general, computation types represent multi-handlers. The type $[\Sigma_1]V_1 \rightarrow \dots \rightarrow$

$[\Sigma_n]V_n \rightarrow [\Sigma]V$, is the type of an n -handler. For each argument type $[\Sigma_i]V_i$ the multi-handler must handle effects in Σ_i on that argument. Such an n -handler handles all of its arguments simultaneously. As a result of handling its arguments it returns a value of type V and may perform effects in Σ . We often write $[\Sigma]V \rightarrow C$ as an abbreviation for $[\Sigma_1]V_1 \rightarrow \dots \rightarrow [\Sigma_n]V_n \rightarrow C$.

Polytypes are restricted to thunks. Effect polymorphism is restricted to a single effect variable ε , which in practice Frank programmers need never write.

Datatypes and effect signatures are declared at the top-level. An effect signature $S \bar{X}$ consists of a collection of command declarations of the form $c : U(\bar{V})$, denoting that command c takes arguments of types \bar{V} and returns a value of type U . The types \bar{V} and U may all depend on \bar{X} . Each command may appear only once in a signature, and each command may appear in only one signature.

An effect set is a sequence of signatures initiated either with the empty effect \emptyset (yielding a *closed effect set*) or the only effect variable ε (yielding an *open effect set*). Order is important, as repeats are permitted, in which case the right-most signature overrides all others with the same name.

Shadowing arises naturally in two ways. First, given an open effect set Σ we may substitute an arbitrary effect set for ε . Second, we define a notion of effect extension: $\Sigma \oplus \Sigma'$ is the *extension* of effect set Σ with closed effect set Σ' , formally:

$$\begin{aligned} \Sigma \oplus \emptyset &= \Sigma \\ \Sigma \oplus (\Sigma', S \bar{V}) &= (\Sigma \oplus \Sigma'), S \bar{V} \end{aligned}$$

Type environments distinguish monomorphic and polymorphic variables.

Just as with the types, Frank terms are separated into value terms and computation terms. With future extensions, such as dependent types, in mind, Frank adopts a bidirectional typing discipline [18]. Thus terms are further sub-divided into those whose type is inferable, and those that may be checked against a type.

Frank is less strict about the separation between value and computation terms than call-by-push-value is. For instance, inferable computations can sometimes be treated as inferable values. This is a deliberate design decision, with the aim of making Frank convenient to program with.

The typing rules for Frank are given in Figure 2. The judgement $\Gamma [\Sigma] \vdash u \text{ infers } V$ says that given type environment Γ , ambient effects Σ , and inferable value term u , then we can *infer* that the type of u is V . The judgement $\Gamma [\Sigma] \vdash V \text{ checks } v$ says that given type environment Γ , ambient effects Σ , value type V , and checkable value term v , then we can *check* that V is the type of v . The judgement $\Gamma [\Sigma] \vdash d \text{ infers } C$ says that given type environment Γ , ambient effects Σ , and inferable computation term d , then we can *infer* that the type of d is C . The judgement $\Gamma \vdash C \text{ checks } e$ says that given type environment Γ , computation type C , and checkable computation term e , then we can *check* that C is the type of e .

The types of monomorphic variables (x) are simply looked up in the type environment. The types of polymorphic variables (f) are looked up and instantiated. In practice this means applying a simple unification-based algorithm. The type of a command (c) is looked up from the ambient effects. An inferable computation (d) is also an inferable value, providing it is a returner whose effects agree with the ambient effects.

Any inferable value (u) is also checkable against its inferred type. Datatype ($k \bar{v}$) and thunk ($\{e\}$) terms are checkable by checking their components. The side condition on the thunk introduction rule requires that the pattern matching clauses of $e : C$ cover C .

A thunk u can be forced ($u!$) if its inferred type agrees with the ambient effects. To infer the type of a handler application $d e$, we first infer the type of d , and then check that the argument matches

$$\boxed{\Gamma [\Sigma] \vdash u \text{ infers } V}$$

$$\frac{x : V \in \Gamma}{\Gamma [\Sigma] \vdash x \text{ infers } V} \quad \frac{f : \forall \bar{Z}. Q \in \Gamma \quad \text{dom}(\theta) = \{\bar{Z}\} \quad \varepsilon \in \text{dom}(\theta) \implies \theta(\varepsilon) = \Sigma}{\Gamma [\Sigma] \vdash f \text{ infers } \theta(Q)} \quad \frac{c : U(\bar{V}) \in \Sigma}{\Gamma [\Sigma] \vdash c \text{ infers } \{\overline{[]V} \rightarrow [\Sigma]U\}}$$

$$\frac{\Gamma [\Sigma] \vdash d \text{ infers } [\Sigma]V}{\Gamma [\Sigma] \vdash d \text{ infers } V}$$

$$\boxed{\Gamma [\Sigma] \vdash V \text{ checks } v}$$

$$\frac{\Gamma [\Sigma] \vdash u \text{ infers } U \quad U = V}{\Gamma [\Sigma] \vdash V \text{ checks } u} \quad \frac{(\Gamma [\Sigma] \vdash V_i \text{ checks } v_i)_i \quad k \bar{V} \in D \bar{U}}{\Gamma [\Sigma] \vdash D \bar{U} \text{ checks } k \bar{v}} \quad \frac{\Gamma \vdash C \text{ checks } e \quad e \text{ covers } C}{\Gamma [\Sigma] \vdash \{C\} \text{ checks } \{e\}}$$

$$\boxed{\Gamma [\Sigma] \vdash d \text{ infers } C}$$

$$\frac{\Gamma [\Sigma] \vdash u \text{ infers } \{C\} \quad C \text{ does } \Sigma}{\Gamma [\Sigma] \vdash u! \text{ infers } C} \quad \frac{\Gamma [\Sigma] \vdash d \text{ infers } [\Sigma']V \rightarrow C \quad \Gamma \vdash [\Sigma \oplus \Sigma']V \text{ checks } e}{\Gamma [\Sigma] \vdash d e \text{ infers } C}$$

$$\frac{(\Gamma, f : \forall \bar{Z}. \bar{Q} \vdash e_i \text{ checks } Q_i)_i \quad \Gamma, f : \forall \bar{Z}. \bar{Q} [\Sigma] \vdash d \text{ infers } C}{\Gamma [\Sigma] \vdash \text{letrec } f : \forall \bar{Z}. \bar{Q} = e \text{ in } d \text{ infers } C}$$

$$\boxed{\Gamma \vdash C \text{ checks } e}$$

$$\frac{\Gamma [\Sigma] \vdash V \text{ checks } v}{\Gamma \vdash [\Sigma]V \text{ checks } v} \quad \frac{\overline{\Sigma'} \text{ closed} \quad (([\Sigma'_i]V_i \text{ matches } r_i \rightarrow \Delta_i)_i \quad \Gamma, \bar{\Delta} \vdash C \text{ checks } e)}{\Gamma \vdash \overline{[\Sigma']V} \rightarrow C \text{ checks } \bar{r} \mapsto e} \quad \frac{V \text{ uninhabited}}{\Gamma \vdash [\Sigma']V \rightarrow C \text{ checks } ()}$$

$$\frac{\Gamma \vdash R \rightarrow C \text{ checks } e \quad \Gamma \vdash R \rightarrow C \text{ checks } e'}{\Gamma \vdash R \rightarrow C \text{ checks } e \mid e'}$$

$$\boxed{C \text{ does } \Sigma}$$

$$\frac{}{\overline{[\Sigma]V \text{ does } \Sigma}} \quad \frac{C \text{ does } \Sigma}{R \rightarrow C \text{ does } \Sigma}$$

$$\boxed{V \text{ matches } p \rightarrow \Delta}$$

$$\frac{}{\overline{V \text{ matches } x \rightarrow x : V}} \quad \frac{k \bar{V} \in D \bar{U} \quad (V_i \text{ matches } p_i \rightarrow \Delta_i)_i}{D \bar{U} \text{ matches } k \bar{p} \rightarrow \bar{\Delta}}$$

$$\boxed{R \text{ matches } r \rightarrow \Delta}$$

$$\frac{V \text{ matches } p \rightarrow \Delta}{\overline{[\Sigma]V \text{ matches } p \rightarrow \Delta}} \quad \frac{c : U(\bar{V}) \in \Sigma \quad (V_i \text{ matches } p_i \rightarrow \Delta_i)_i}{\overline{[\Sigma]V \text{ matches } c \bar{p} ? g \rightarrow \bar{\Delta}, g : \{\{\emptyset\}U \rightarrow [\Sigma]V\}}} \quad \frac{}{\overline{[\Sigma]V \text{ matches } x! \rightarrow x : \{[\Sigma]V\}}}$$

Figure 2. Frank Typing Rules

the inferred argument type. Note that the effect set on an argument must be closed, and thus we may extend the ambient effect set with the argument effect set when checking the argument. The term $\text{letrec } \overline{f : P = e} \text{ in } d$ binds the mutually recursive polymorphic functions $\overline{f : P}$ in d .

Any term v that type checks with ambient effects Σ at value type V also type checks at computation type $[\Sigma]V$. A multi-handler of type $\overline{[\Sigma']V} \rightarrow C$ is built by composing clauses of the form $\bar{r} \mapsto e$, where \bar{r} is a sequence of *computation patterns* whose variables are bound in e . The $e \mid e'$ construct composes the clauses of e with those of e' . A multi-handler is defined by a collection of clauses that

provides complete coverage of the input types. The $()$ construct allows empty sets of clauses to be constructed in the event that the value type of a function argument is uninhabited.

We preclude composition of non-clauses (i.e. returners) by constraining $e \mid e'$ to have function type. It is perfectly legitimate to compose clauses with different numbers of arguments, though as a preliminary part of pattern matching compilation, we first transform the program to ensure that all composed clauses do have the same number of arguments.

Value patterns are standard, consisting of variable patterns (x) and datatype constructor patterns ($k \bar{p}$).

Computation patterns are more interesting. Any value pattern p is also a computation pattern (used to match against the value returned by the returner computation). A request pattern $c \bar{p} ? g$ matches against a computation of the form $c! \bar{v}$ if the values v match against p . Furthermore, it also binds g to the continuation of the computation delimited by the nearest enclosing multi-handler. This is where the real power of multi-handlers arises. A thunk pattern $x!$ matches any computation reifying it as a thunk bound to x .

Sequencing Computations We write $\text{let } x = e \text{ in } e'$ as syntactic sugar for $\text{bind } e \{x \mapsto e'\}$, where bind is defined in Section 2. More verbosely:

$$\begin{aligned} \text{let } x = e \text{ in } e' &\equiv \\ \text{letrec } (\text{bind} : \forall \varepsilon X Y. X \rightarrow \{X \rightarrow [\varepsilon]Y\} \rightarrow [\varepsilon]Y) = \\ &x f \mapsto f x \text{ in } \text{bind}! e \{x \mapsto e'\} \end{aligned}$$

3.1 Effect Polymorphism with an Invisible Effect Variable

We now give a brief description of how syntactic sugar allows Frank programmers to omit effect variables completely. Consider the type of map in Section 2:

$$\{X \rightarrow Y\} \rightarrow \text{List } X \rightarrow \text{List } Y$$

Apart from perhaps the curly braces, this looks pretty much the same as the type a functional programmer might expect to write in a language without support for effect typing.

In fact, this type desugars into the rather more verbose:

$$\{[\emptyset]X \rightarrow [\varepsilon]Y\} \rightarrow [\emptyset](\text{List } X) \rightarrow [\varepsilon](\text{List } Y)$$

Let us distinguish between returners $[\Sigma]V$ in argument position, whose effects Σ must be closed, and those in tail position, whose effects need not be closed. We call the former *ports*, and the latter *pegs*.

Observe that ports are closed, so it is never necessary to write the \emptyset in their effects. Pegs, on the other hand, may be open or closed. We adopt the convention that ε may be omitted from the start of a peg's effects. Thus, if we know it is a peg, then V means $[\]V$, which means $[\varepsilon]V$, and $[\text{Abort}]V$ means $[\varepsilon, \text{Abort}]V$.

We now summarise the syntactic sugar. For ports and pegs:

$$V \equiv [\]V$$

For ports:

$$[\overline{S \bar{V}}]U \equiv [\emptyset, \overline{S \bar{V}}]U$$

For pegs:

$$[\overline{S \bar{V}}]U \equiv [\varepsilon, \overline{S \bar{V}}]U$$

With this syntactic sugar in place, we can now avoid writing the effect variable ε in Frank programs, ever. In addition, we need never write \emptyset in port effect sets. It is sometimes necessary to explicitly write \emptyset in peg effect sets. In particular, a pure top-level program returning values of type V has type $[\emptyset]V$.

4. Pattern Matching Compilation

We take a fairly standard approach to compiling away pattern matching. As we may match simultaneously against multiple side-effecting computations, we must be somewhat careful about order. Optionally, we can expose incomplete or ambiguous pattern matching as concrete effects.

The target language of pattern matching compilation, Core Frank, replaces multi-handlers with a combination of call-by-value functions, case statements, and unary effect handlers.

The syntax of Core Frank is given in Figure 3. The Core Frank typing rules are given in Figure 4.

Multi-handlers in Frank become curried functions over suspended computations in Core Frank. Shallow pattern matching on

Types

(values)	$U, V ::= D \bar{U} \mid \{C\} \mid X$
(computations)	$C ::= R \mid V \rightarrow C$
(returners)	$R ::= [\Sigma]V$
(quantifiers)	$Z ::= X \mid \varepsilon$
(polytypes)	$P ::= \forall \bar{Z}. Q$
(thunks)	$Q ::= \{C\}$
(signatures)	$S \bar{X} ::= \cdot \mid c : U(\bar{V}), S \bar{X}$
(effects)	$\Sigma ::= \emptyset \mid \Sigma, S \bar{V} \mid \varepsilon$
(type environments)	$\Gamma ::= \cdot \mid \Gamma, x : V \mid f : P$

Terms

(inferable values)	$u ::= x \mid f \mid c \mid d$
(checkable values)	$v ::= u \mid k \bar{v} \mid \{e\}$
(inferable computations)	$d ::= u! \mid d v$
	$\mid \text{letrec } \overline{f : P = e} \text{ in } d$
(checkable computations)	$e ::= v \mid \lambda x. e$
	$\mid \text{case } u \text{ of } (k \bar{x}_k \mapsto e_k)_{\bar{k}}$
	$\mid \text{handle } d \text{ with } (c \bar{x}_c ? g_c \mapsto e_c)_{\bar{c}}$
	$\mid x \mapsto e$

Figure 3. Core Frank Syntax

a single request becomes unary effect handling. Shallow pattern matching on a datatype value becomes case analysis. Nested pattern matching on multiple computations is realised as a pattern matching tree constructed from handlers and case statements.

We may adapt standard algorithms for pattern matching compilation apply (e.g. [1] or [13]). Rather than committing to a particular one, we outline how a pattern matching compiler fits into our setting, what input it takes, and what kind of output it must produce.

Given a Frank expression $\{e\}$ such that $\Gamma \vdash R_1 \dots R_n \rightarrow R$ **checks** e we compile it to an equivalent Core Frank expression $\{\langle e \rangle\}$. First we expand all of the clauses in e to yield an n column pattern matrix. For instance, suppose the arguments have types $[\text{Send Char}, \text{Abort}] \text{Unit}$ and $[\text{Receive Char}]$, and we have the following clauses:

$(\text{send } x ? s)$	$(\text{receive } ? r)$	$\mapsto e_1$
$(\text{send } x ? s)$	z	$\mapsto e_2$
$(\text{abort } ? s)$	e_3	
unit	e_4	

then this becomes:

$(\text{send } x ? s)$	$(\text{receive } ? r)$	$\mapsto e_1$
$(\text{send } x ? s)$	z	$\mapsto e_2$
$(\text{abort } ? s)$	$y!$	$\mapsto \text{bind } y! e_3$
unit	$y!$	$\mapsto \text{bind } y! e_4$

where bind is defined in Section 2. The missing patterns have been inserted as thunk patterns, which match any computation. Invoking bind allows us to forward the computation bound by the thunk pattern to the existing continuation. Next we generate a vector of fresh variables, one for each argument.

$$x_0 x_1$$

The goal of pattern matching compilation is to generate a pattern matching tree that matches the variable vector against all of the patterns in the pattern matrix in the correct order.

C does Σ

$$\frac{}{\overline{[\Sigma]V \text{ does } \Sigma}} \qquad \frac{C \text{ does } \Sigma}{\overline{V \rightarrow C \text{ does } \Sigma}}$$

$\Gamma[\Sigma] \vdash u$ infers V

$$\frac{x : V \in \Gamma}{\Gamma[\Sigma] \vdash x \text{ infers } V} \qquad \frac{f : \forall \bar{Z}. Q \in \Gamma \quad \text{dom}(\theta) = \{\bar{Z}\} \quad \varepsilon \in \text{dom}(\theta) \implies \theta(\varepsilon) = \Sigma}{\Gamma[\Sigma] \vdash f \text{ infers } \theta(Q)} \qquad \frac{c : U(\bar{V}) \in \Sigma}{\Gamma[\Sigma] \vdash c \text{ infers } \{\bar{V} \rightarrow [\Sigma]U\}}$$

$$\frac{\Gamma[\Sigma] \vdash d \text{ infers } [\Sigma]V}{\Gamma[\Sigma] \vdash d \text{ infers } V}$$

$\Gamma[\Sigma] \vdash V$ checks v

$$\frac{\Gamma[\Sigma] \vdash u \text{ infers } U \quad U = V}{\Gamma[\Sigma] \vdash V \text{ checks } u} \qquad \frac{(\Gamma[\Sigma] \vdash V_i \text{ checks } v_i)_i \quad k \bar{V} \in D \bar{U}}{\Gamma[\Sigma] \vdash D \bar{U} \text{ checks } k \bar{v}} \qquad \frac{\Gamma \vdash C \text{ checks } e}{\Gamma[\Sigma] \vdash \{C\} \text{ checks } \{e\}}$$

$\Gamma[\Sigma] \vdash d$ infers C

$$\frac{\Gamma[\Sigma] \vdash u \text{ infers } \{C\} \quad C \text{ does } \Sigma}{\Gamma[\Sigma] \vdash u! \text{ infers } C} \qquad \frac{\Gamma[\Sigma] \vdash d \text{ infers } [\Sigma']V \rightarrow C \quad \Gamma[\Sigma \oplus \Sigma'] \vdash V \text{ checks } v}{\Gamma[\Sigma] \vdash d v \text{ infers } C}$$

$$\frac{(\Gamma, f : \forall \bar{Z}. Q \vdash e_i \text{ checks } Q_i)_i \quad \overline{\Gamma, f : \forall \bar{Z}. Q [\Sigma] \vdash d \text{ infers } C}}{\Gamma[\Sigma] \text{ letrec } f : \forall \bar{Z}. Q = e \text{ in } d \text{ infers } C}$$

$\Gamma \vdash C$ checks e

$$\frac{\Gamma[\Sigma] \vdash V \text{ checks } v}{\Gamma \vdash [\Sigma]V \text{ checks } v} \qquad \frac{\Gamma, x : V \vdash C \text{ checks } e}{\Gamma \vdash V \rightarrow C \text{ checks } \lambda x. e}$$

$$\frac{\Gamma[\Sigma] \vdash u \text{ infers } D \bar{U} \quad C \text{ does } \Sigma \quad (\Gamma, \bar{x}_k : \bar{V} \vdash C \text{ checks } e_k)_{k \bar{V} \in D \bar{U}}}{\Gamma \vdash C \text{ checks case } u \text{ of } (k \bar{x}_k \mapsto e_k)_k}$$

$$\frac{\Sigma' \text{ closed} \quad C \text{ does } \Sigma \quad \Gamma[\Sigma \oplus \Sigma'] \vdash d \text{ infers } [\Sigma \oplus \Sigma']V \quad (\Gamma, \bar{x}_c : \bar{V}, g_c : \{U \rightarrow [\Sigma \oplus \Sigma']V\} \vdash C \text{ checks } e_c)_{c:U(\bar{V}) \in \Sigma'} \quad \Gamma, x : V \vdash C \text{ checks } e}{\Gamma \vdash C \text{ checks handle } d \text{ with } (c \bar{x}_c ? g_c \mapsto e_c)_c \mid x \mapsto e}$$

Figure 4. Core Frank Typing Rules

In Frank, pattern matching trees M are built up from leaves, case analysis, and handlers.

$$M ::= \bar{e} \mid \text{case } x \text{ of } (k \bar{x}_k \mapsto M_k)_{k \in D} \mid \text{handle } x! \text{ with } (c \bar{x}_c ? g_c \mapsto M_c)_{c \in \Sigma} \mid x \mapsto M$$

The leaves consist of a sequence of checkable computation expressions. Each element corresponds to one way of matching all of the patterns. If there exists a leaf with no elements, then the pattern matching is incomplete; if there exists a leaf with multiple elements, then the pattern matching is ambiguous. Our default strategy (as indicated by the think introduction rule) is to class incomplete pattern matching as a type error, and to keep only the first element in the case of ambiguous pattern matching. Our example generates

the following pattern matching tree:

$$M = \text{handle } x_0! \text{ with} \\ \text{send } x ? s \mapsto \\ \text{handle } x_1! \text{ with} \\ \text{receive } ? r \mapsto e_1 \\ z \mapsto e_2 \\ \text{abort } ? s \mapsto \\ \text{handle } x_1! \text{ with} \\ \text{receive } ? r \mapsto \text{bind } (r! \text{ receive!}) e_3 \\ z \mapsto \text{bind } z e_3 \\ y \mapsto \\ \text{case } y \text{ of} \\ \text{unit} \mapsto \text{handle } x_1! \text{ with} \\ \text{receive } ? r \mapsto \text{bind } (r! \text{ receive!}) e_4 \\ z \mapsto \text{bind } z e_4$$

Each think pattern has been expanded out to explicitly list all of the cases according to its type. We obtain the corresponding Core

Frank code by abstracting over the fresh variables.

$$\lambda x_0 x_1. M$$

Some pattern matching operations reorder columns as an optimisation. Column reordering is not in general a valid optimisation in Frank. This is because commands in the ambient effects, but not in the argument effects, are implicitly forwarded, and the order in which they are forwarded is left-to-right. (The forwarding behaviour is made precise in the Section 6.)

Of course, because Core Frank takes values as arguments whereas Frank takes computations, each argument must be wrapped in a thunk constructor. The type translation is given simply by the homomorphic extension of the following equation on function types:

$$([\Sigma]V) \rightarrow \langle C \rangle = \{([\Sigma])\langle V \rangle\} \rightarrow \langle C \rangle$$

A correct pattern matching translation $(-)$ from Frank to Core Frank should be type preserving.

- If $\Gamma [\Sigma] \vdash u$ infers V then $(\Gamma) [\langle \Sigma \rangle] \vdash (u)$ infers $\langle V \rangle$.
- If $\Gamma [\Sigma] \vdash V$ checks v then $(\Gamma) [\langle \Sigma \rangle] \vdash \langle V \rangle$ checks $\langle v \rangle$.
- If $\Gamma [\Sigma] \vdash d$ infers C then $(\Gamma) [\langle \Sigma \rangle] \vdash \langle d \rangle$ infers $\langle C \rangle$.

4.1 Incomplete and Ambiguous Pattern Matching as Effects

As an extension to Frank, we might allow incomplete and ambiguous pattern matching. The former may be permitted if the ambient effects contain the *Abort* signature, in which case incomplete patterns are translated into the *abort* : *Zero* command, which can then be handled however the programmer wishes. Similarly, we can define a *choice* : $X(X, X)$ command, in order to allow ambiguous pattern matches to be handled by the programmer.

5. Explicit Control Flow and Polymorphism

In order to make control flow explicit, and to ease the definition of an operational semantics, we translate Core Frank into a call-by-push-value calculus F_{eff} based on Kammar et al's λ_{eff} [8].

At the same time, we make polymorphism explicit. The syntax of F_{eff} is given in Figure 5. Broadly, F_{eff} types are similar to Frank types, with an explicit division between value type and computation types. A superficial difference is that rather than annotating value returning computations with effects, we shift such labels to the thunk containing the computation. The former design seems more convenient to program with, which is why we adopt it in the source language. The latter design leads to a more uniform presentation of the typing rules, and matches the design of λ_{eff} . Another minor difference is that in F_{eff} commands must always be fully applied, which leads to a cleaner operational semantics. Note that *let* is actually redundant in F_{eff} , just as it is in Frank and Core Frank. We include *let* as a special construct because doing so makes the semantics cleaner.

The typing rules for F_{eff} are given in Figure 6. Typing in F_{eff} is unidirectional.

Call-by-push-value calculi such as λ_{eff} and F_{eff} make a strict separation between values and computations, not dissimilar from CPS or A-normal form representations, in which all reduction takes place at the level of computations. In such a setting it would seem most natural to add type abstractions to computations rather than values. However, this does not give us what we need in the presence of effects, as we need to be able to quantify over effects. Our solution is to build the universal quantifier into the thunk type (the introduction rule for thunks is the place where ambient effects are reified in a type) and build type application into forcing.

Types

(values)	$U, V ::= D \bar{U} \mid \forall \bar{Z}. [\Sigma] \{C\} \mid X$
(computations)	$C ::= \langle V \rangle \mid V \rightarrow C$
(quantifiers)	$Z ::= X \mid \varepsilon$
(arguments)	$T ::= V \mid \Sigma$
(polytypes)	$P ::= \forall \bar{Z}. Q$
(monothunks)	$Q ::= [\Sigma] \{C\}$
(signatures)	$S ::= \cdot \mid c : U(\bar{V}), S$
(effects)	$\Sigma ::= \emptyset \mid \Sigma, S \bar{V} \mid \varepsilon$
(type environments)	$\Gamma ::= \cdot \mid \Gamma, x : V$

Terms

(values)	$u, v ::= x \mid k \bar{v} \mid \Lambda \bar{Z}. \{d\}$
(computations)	$d, e ::= \text{case } u \text{ of } (k \bar{x}_k \mapsto e_k)_k \mid (u \bar{T})!$ $\mid \lambda x. e \mid d v \mid \text{ret } v \mid \text{let } x = e \text{ in } e' \mid c \bar{v}$ $\mid \text{handle } d \text{ with } (c \bar{x}_c ? g \mapsto e_c)_c \mid x \mapsto e$ $\mid \text{letrec } f : P = e \text{ in } e'$

Figure 5. F_{eff} Syntax

The type translation from Core Frank to F_{eff} is given by the homomorphic extension of the following equations:

$$\begin{aligned} \llbracket \{C\} \rrbracket &= \llbracket [\Sigma] \{C\} \rrbracket, \quad \text{where } C \text{ does } \Sigma \\ \llbracket [\Sigma] V \rightarrow C \rrbracket &= \llbracket [\Sigma] \langle V \rangle \rrbracket \rightarrow \llbracket C \rrbracket \\ \llbracket [\Sigma] V \rrbracket &= \langle \llbracket V \rrbracket \rangle \end{aligned}$$

In order to simplify the definition of the term translation, as a preprocessing step we annotate instances of polymorphic variables f with the type arguments they are instantiated with and commands with their arities. The term translation is a call by value embedding of Core Frank into F_{eff} :

$$\begin{aligned} \llbracket x \rrbracket &= \text{ret } x \\ \llbracket f \bar{T} \rrbracket &= \text{ret } \{f \langle \bar{T} \rangle\} \\ \llbracket c^n \rrbracket &= \text{ret } \{\lambda x_1 \dots x_n. c x_1 \dots x_n\} \\ \llbracket k \bar{v} \rrbracket &= \text{let } \bar{x} = \langle \bar{v} \rangle \text{ in ret } (k \bar{x}) \\ \llbracket \{e\} \rrbracket &= \text{ret } \{\langle e \rangle\} \\ \llbracket d v \rrbracket &= \text{let } x = \langle v \rangle \text{ in } \llbracket d \rrbracket x \\ \llbracket u! \rrbracket &= \text{let } x = \langle u \rangle \text{ in } x! \\ \llbracket \text{letrec } f : \forall \bar{Z}. Q = e \text{ in } d \rrbracket &= \text{letrec } f : \langle \forall \bar{Z}. Q \rangle = \langle e \rangle \text{ in } \llbracket d \rrbracket \end{aligned}$$

$$\begin{aligned} \llbracket \lambda x. e \rrbracket &= \lambda x. \llbracket e \rrbracket \\ \llbracket \text{case } u \text{ of } (k \bar{x}_k \mapsto e_k)_k \rrbracket &= \\ &\text{let } x = \langle u \rangle \text{ in case } x \text{ of } (k \bar{x}_k \mapsto \langle e_k \rangle)_k \\ \llbracket \text{handle } d \text{ with } (c \bar{x}_c ? g_c \mapsto e_c)_c \mid x \mapsto e \rrbracket &= \\ &\text{handle } \llbracket d \rrbracket \text{ with } (c \bar{x}_c ? g_c \mapsto \langle e_c \rangle)_c \mid x \mapsto \langle e \rangle \end{aligned}$$

PROPOSITION 1. *The translation $\llbracket - \rrbracket$ from Core Frank to F_{eff} is type preserving.*

- If $\Gamma [\Sigma] \vdash u$ infers V then $\llbracket \Gamma \rrbracket [\langle \Sigma \rangle] \vdash \langle u \rangle : \langle \llbracket V \rrbracket \rangle$.
- If $\Gamma [\Sigma] \vdash V$ checks v then $\llbracket \Gamma \rrbracket [\langle \Sigma \rangle] \vdash \langle v \rangle : \langle \llbracket V \rrbracket \rangle$.
- If $\Gamma [\Sigma] \vdash d$ infers C then $\llbracket \Gamma \rrbracket [\langle \Sigma \rangle] \vdash \langle d \rangle : \langle \llbracket C \rrbracket \rangle$.
- If $\Gamma \vdash c$ checks e and C does Σ then $\llbracket \Gamma \rrbracket [\langle \Sigma \rangle] \vdash \langle e \rangle : \langle \llbracket C \rrbracket \rangle$.

Kammar et al [8] classify a number of different varieties of handler. The handlers in F_{eff} (and indeed Frank) are *shallow* in that, unlike Plotkin and Pretnar's original deep handlers, the handler is not automatically rewrapped around the continuation when handling a command. Roughly, deep handlers perform a fold over computa-

$$\boxed{\Gamma \vdash v : V}$$

$$\frac{x : V \in \Gamma}{\Gamma \vdash x : V}$$

$$\frac{\Gamma \vdash V_i : v_i \quad k \bar{V} \in D \bar{U}}{\Gamma \vdash D \bar{U} : k \bar{v}}$$

$$\frac{\Gamma [\Sigma] \vdash e : C}{\Gamma \vdash \Lambda \bar{Z}. \{e\} : \forall \bar{Z}. [\Sigma] \{C\}}$$

$$\boxed{\Gamma [\Sigma] \vdash e : C}$$

$$\frac{\Gamma \vdash u : D \bar{U} \quad (\Gamma, \bar{x}_k : \bar{V} [\Sigma] \vdash e_k : C)_{k \bar{v} \in D \bar{U}}}{\Gamma [\Sigma] \vdash \text{case } u \text{ of } (k \bar{x}_k \mapsto e_k)_k : C}$$

$$\frac{\Gamma \vdash u : \forall \bar{Z}. [\Sigma] \{C\}}{\Gamma [\Sigma] \vdash (u \bar{T})! : C[\bar{T}/\bar{Z}]}$$

$$\frac{\Gamma, x : V [\Sigma] \vdash e : C}{\Gamma [\Sigma] \vdash \lambda x. e : V \rightarrow C}$$

$$\frac{\Gamma [\Sigma] \vdash d : V \rightarrow C \quad \Gamma \vdash v : V}{\Gamma [\Sigma] \vdash d v : C}$$

$$\frac{\Gamma \vdash v : V}{\Gamma [\Sigma] \vdash \text{ret } v : \langle V \rangle}$$

$$\frac{\Gamma [\Sigma] \vdash d : \langle V \rangle \quad \Gamma, x : V [\Sigma] \vdash e : C}{\Gamma [\Sigma] \vdash \text{let } x = d \text{ in } e : C}$$

$$\frac{c : U(\bar{V}) \in \Sigma \quad \Gamma \vdash v_i : V_i}{\Gamma [\Sigma] \vdash c \bar{v} : \langle U \rangle}$$

$$\frac{\Gamma [\Sigma \oplus \Sigma'] \vdash d : \langle V \rangle \quad (\Gamma, \bar{x}_c : \bar{V}, g : [\Sigma \oplus \Sigma'] \{U \rightarrow \langle V \rangle\} [\Sigma] \vdash e_c : C)_{c \bar{v} \in \Sigma'} \quad \Gamma, x : V [\Sigma] \vdash e : C}{\Gamma [\Sigma] \vdash \text{handle } d \text{ with } (c \bar{x}_c ? g \mapsto e_c)_c \mid x \mapsto e : C}$$

$$\frac{(\Gamma, f : \forall \bar{Z}. Q [\Sigma] \vdash e_i : Q_i)_i \quad \Gamma, f : \forall \bar{Z}. Q [\Sigma] \vdash d : C}{\Gamma [\Sigma] \vdash \text{letrec } f : \forall \bar{Z}. Q = e \text{ in } d : C}$$

Figure 6. F_{eff} Typing rules

tions, whereas shallow handlers perform a case split. Deep handlers are denotationally better behaved than shallow handlers, but shallow handlers sometimes appear more convenient to program with. The handlers in F_{eff} (and indeed Frank) are polymorphic forwarding handlers. The form of polymorphism amounts to a kind of row polymorphism with shadowing, unlike the more conventional kind of row polymorphism suggested by Kammar et al.

6. Small-step semantics

We give a small step operation semantics for F_{eff} inspired by Kammar et al's semantics for λ_{eff} [8]. All of the rules except the ones for handlers are pretty standard β -reductions. The rules for forcing are slightly unusual due to the hard-wiring of polymorphism into thunks.

Returns $\text{ret } v$ are handled by substituting the value into the handler's return clause. Commands are handled by capturing the continuation up to the current handler and dispatching to the appropriate clause for the command. If a clause is defined in the handler, then that clause is selected. If not, then the command is forwarded to be handled by an outer handler, but its continuation is handled by the current handler.

Delimited computation contexts D are used to characterise the continuation up to the current handler. Computation contexts E amount to evaluation contexts for a call-by-push-value setting.

PROPOSITION 2 (Type Soundness).

- If $\Gamma [\Sigma] \vdash d : C$ and $d \longrightarrow e$ then $\Gamma [\Sigma] \vdash e : C$.
- If $\Gamma [\emptyset] \vdash d : \langle V \rangle$ then either there exists v such that $d = \text{ret } v$ or there exists e such that $d \longrightarrow e$.

7. Related Work

We have discussed much of the related work throughout the paper. Here we briefly mention some other related work.

A natural implementation for handlers is to use *free monads* [8]. Swierstra [27] illustrates how to write effectful programs with free

monads in Haskell, taking advantage of type-classes to provide a certain amount of modularity.

Inspired by Bauer and Pretnar's Eff, Visscher has implemented the `effects` library [28]. The key idea is to layer continuation monads in the style of Filinski [5], using Haskell type classes to automatically infer lifting between layers.

Filinski's work on monadic reflection and layered monads is closely related to effect handlers [6]. Monadic reflection supports a similar style of composing effects. The key difference is that monadic reflection interprets monadic computations in terms of other monadic computations, rather than abstracting over and interpreting operations

Languages other than Frank that attempt to elide some effect variables from source code include Links [12] and Daan Leijen's Koka [10]. Neither eliminates effect variables altogether.

Swamy et al [26] add support for monads in ML, supporting direct-style effectful programming in a strict language. Unlike Frank, their system is based on monad transformers rather than effect handlers.

8. Future Work

Our first priority is to implement a prototype for the current Frank design. Having done that, there is much scope for exploring different implementations of handlers, both for performance and for exploring new abstractions. We would like to combine effect handlers with richer type systems, following the work of McBride [15] and Brady [4]. Handlers provide some of the same functionality as modules and type classes. We would like to formally relate all three abstractions. We are exploring algebraic effects and effect handlers for idiom [17] and arrow [7] computations.

References

- [1] L. Augustsson. Compiling pattern matching. In *FPCA*, pages 368–381, 1985.
- [2] A. Bauer and M. Pretnar. Programming with algebraic effects and handlers. *CoRR*, abs/1203.1539, 2012.

$$\begin{aligned}
& \text{(Delimited computation contexts) } D ::= [\] \mid D v \mid \text{let } x = D \text{ in } e \\
& \text{(Computation contexts) } E ::= [\] \mid E v \mid \text{let } x = E \text{ in } e \mid \text{handle } E \text{ with } H \\
& \text{case } k_j \bar{v} \text{ of } (k_i \bar{x}_i \mapsto e_i)_i \longrightarrow e_j[\bar{v}/\bar{x}_j] \\
& \quad ((\Lambda \bar{Z}. \{e\}) \bar{T})! \longrightarrow e[\bar{T}/\bar{Z}] \\
& \quad (\lambda x. e) v \longrightarrow e[v/x] \\
& \quad \text{let } x = \text{ret } v \text{ in } e \longrightarrow e[v/x] \\
& \text{letrec } f : \forall \bar{Z}. Q = e \text{ in } d \longrightarrow d[\Lambda \bar{Z}. \{e\}[\Lambda \bar{Z}. \{\text{letrec } f : \forall \bar{Z}. Q = e \text{ in } (f \bar{Z})!\}/\bar{f}]/\bar{f}] \\
& \quad \text{handle } (\text{ret } v) \text{ with } H \longrightarrow H(\text{ret}, v) \\
& \quad \text{handle } D[c \bar{v}] \text{ with } H \longrightarrow H(c, \bar{v}, \{\lambda z. D[\text{ret } z]\})
\end{aligned}$$

where the action of $H = (c_i \bar{x}_i ? g \mapsto e_i)_i \mid x \mapsto e$ is given by:

$$\begin{aligned}
H(\text{ret}, v) &= e[v/x] \\
H(c_j, \bar{v}, u) &= e_j[\bar{v}/\bar{x}, u/g] \\
H(c, \bar{v}, u) &= \text{let } z = c \bar{v} \text{ in handle } u! z \text{ with } H, \quad c \neq c_i \text{ for any } i
\end{aligned}$$

$$\frac{d \longrightarrow e}{E[d] \longrightarrow E[e]}$$

Figure 7. Small-step operational semantics for F_{eff}

- [3] A. Bauer and M. Pretnar. An effect system for algebraic effects and handlers. In *CALCO*, volume 8089 of *Lecture Notes in Computer Science*, pages 1–16. Springer, 2013.
- [4] E. Brady. Programming and reasoning with algebraic effects and dependent types. In *ICFP*. ACM, 2013.
- [5] A. Filinski. Representing layered monads. In *POPL*. ACM, 1999.
- [6] A. Filinski. Monads in action. In *POPL*. ACM, 2010.
- [7] J. Hughes. Programming with arrows. In *Advanced Functional Programming*, volume 3622 of *Lecture Notes in Computer Science*, pages 73–129. Springer, 2004.
- [8] O. Kammar, S. Lindley, and N. Oury. Handlers in action. In *ICFP*, pages 145–158. ACM, 2013.
- [9] O. Kiselyov, A. Sabry, and C. Swords. Extensible effects: an alternative to monad transformers. In *Haskell*, pages 59–70. ACM, 2013.
- [10] D. Leijen. Koka: Programming with row-polymorphic effect types. Technical Report MSR-TR-2013-79, Microsoft Research, 2013.
- [11] P. B. Levy. *Call-By-Push-Value: A Functional/Imperative Synthesis*, volume 2 of *Semantics Structures in Computation*. Springer, 2004.
- [12] S. Lindley and J. Cheney. Row-based effect types for database integration. In *TLDI*. ACM, 2012.
- [13] L. Maranget. Compiling pattern matching to good decision trees. In *ML*, pages 35–46. ACM, 2008.
- [14] C. McBride. How might effectful programs look? In *Workshop on Effects and Type Theory*, 2007. <http://cs.ioc.ee/efftt/mcbride-slides.pdf>.
- [15] C. McBride. Kleisli arrows of outrageous fortune, 2011. Accepted for publication. <https://personal.cis.strath.ac.uk/conor.mcbride/Kleisli.pdf>.
- [16] C. McBride. Frank (0.3), 2012. <http://hackage.haskell.org/package/Frank>.
- [17] C. McBride and R. Paterson. Applicative programming with effects. *J. Funct. Program.*, 18(1):1–13, 2008.
- [18] B. C. Pierce and D. N. Turner. Local type inference. *ACM Trans. Program. Lang. Syst.*, 22(1):1–44, 2000.
- [19] G. D. Plotkin and J. Power. Semantics for algebraic operations. *Electr. Notes Theor. Comput. Sci.*, 45, 2001.
- [20] G. D. Plotkin and J. Power. Adequacy for algebraic effects. In *FoSSaCS*. Springer-Verlag, 2001.
- [21] G. D. Plotkin and J. Power. Notions of computation determine monads. In *FoSSaCS*. Springer-Verlag, 2002.
- [22] G. D. Plotkin and J. Power. Algebraic operations and generic effects. *Appl. Categ. Structures*, 11(1):69–94, 2003.
- [23] G. D. Plotkin and M. Pretnar. Handlers of algebraic effects. In *ESOP*. Springer-Verlag, 2009.
- [24] G. D. Plotkin and M. Pretnar. Handling algebraic effects. *Logical Methods in Computer Science*, 9(4), 2013.
- [25] M. Pretnar. Inferring algebraic effects. *CoRR*, abs/1312.2334, 2013.
- [26] N. Swamy, N. Guts, D. Leijen, and M. Hicks. Lightweight monadic programming in ML. In *ICFP*. ACM, 2011.
- [27] W. Swierstra. Data types à la carte. *J. Funct. Program.*, 18(4):423–436, 2008.
- [28] S. Visscher. The effects package (0.2.2), 2012. <http://hackage.haskell.org/package/effects>.