

Algebraic Effects and Effect Handlers for Idioms and Arrows

Sam Lindley

The University of Edinburgh

Abstract

Plotkin and Power’s algebraic effects combined with Plotkin and Pretnar’s effect handlers provide a foundation for modular programming with effects. We present a generalisation of algebraic effects and effect handlers to support other kinds of effectful computations corresponding to McBride and Paterson’s idioms and Hughes’ arrows.

Categories and Subject Descriptors D.1.1 [Applicative (Functional) Programming]; D.3.3 [Language Constructs and Features]; F.3.2 [Semantics of Programming Languages]: Operational semantics

Keywords algebraic effects; effect handlers; idioms; arrows; monads; applicative functors; call-by-push-value

1. Introduction

In previous work [8] we advocated Plotkin and Power’s algebraic effects [21, 22] and effect handlers [23] as a foundation for modular programming with effects. We introduced a statically typed effect handler calculus λ_{eff} along with a sound, terminating, small-step operational semantics, and used it as the basis for practical implementations of handlers in Haskell, ML, and Racket.

Our calculus λ_{eff} (and standard algebraic effects and handlers) provide a means for programming with monadic effects [18], supporting generic effectful computations that can be handled in multiple ways. In this work, we adapt λ_{eff} to accommodate other kinds of effectful computations corresponding to McBride and Paterson’s idioms (also known as applicative functors) [17] and Hughes’ arrows [7]. The resulting calculus, λ_{flow} , extends λ_{eff} with *flow effects*, which explicitly track dependencies between the result of an effectful operation and subsequent effectful computation, allowing us to encode idiom and arrow computations. Crucially, λ_{flow} supports effect handlers for idiom and arrow computations.

A key reason why arrow programs and idiom programs are of interest to functional programmers is that they admit a wider range of implementations than corresponding monadic programs. Every monad is an arrow and every monad is also an idiom. But there exist arrows that are not monads and idioms that are not monads. For instance, non-monadic arrows are often used in functional reactive programming, and non-monadic idioms in parser combinators,

in both cases providing more space and time efficient implementations than monadic alternatives.

The λ_{flow} calculus combines the benefits of λ_{eff} (modular support for handling multiple effects) and our earlier work on the arrow calculus [13, 14] (providing a uniform foundation for programming with idioms, arrows, and monads with a single effect). It allows us to use the same syntax for idiom, arrow, and monad computations. Thus, we can write generic effectful combinators in λ_{flow} that can be used in idiom, arrow, and monad computations.

Just as λ_{eff} gave us a firm basis for building practical implementations of standard effect handlers, we hope that λ_{flow} can provide a firm basis for building practical implementations of idiom and arrow handlers, and in particular support generic effectful programs, combining idiom, arrow, and monad computations.

Our main contributions are as follows:

- We introduce flow effects as a means for distinguishing abstract idiom, arrow, and monad computations.
- We provide a uniform foundation for programming with idioms, arrows, and monads with multiple effects, generalising both λ_{eff} and the arrow calculus.
- An immediate consequence of our formulation is that the inclusions between abstract idiom and abstract arrow computations, and abstract arrow and abstraction monad computations of Lindley et al. [14] are strict. Abstract monad programs are strictly more expressive than abstract arrow programs, which are in turn strictly more expressive than abstract idiom programs.
- We introduce idiom and arrow handlers, as variations on standard monadic effect handlers.
- We illustrate the use of an extension of λ_{flow} for writing generic parser combinators.

The remainder of the paper is structured as follows. Section 2 introduces the key ideas of our approach by first characterising abstract effectful computations as computation trees, and then defining flow effects in terms of computation trees. Section 3 describes λ_{eff} , first focusing on abstract computations, and then on effect handlers. Section 4 presents flow effects and Core λ_{flow} , the fragment of λ_{flow} for expressing abstract idiom, arrow, and monad computations. Section 5 presents handlers in λ_{flow} for idiom, arrow, and monad computations. Section 6 presents a parser combinator implementation using an extension of λ_{flow} . Section 7 discusses related work. Section 8 discusses future work.

2. Effects as Computation Trees

2.1 What is an Effectful Computation?

Plotkin and Power [21, 22] introduced algebraic effects for modelling the semantics of effectful computations. They gave an abstract categorical treatment. We will be much more concrete, and after our initial example consider only free algebras.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WGP '14, August 31 2014, Gothenburg, Sweden.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3042-8/14/08...\$15.00.

<http://dx.doi.org/10.1145/2633628.2633636>

An algebraic effect is given by a signature of operations along with a set of equations on those operations. For example, we might define an algebraic effect for read only boolean state with the following signature:

$$\{\text{get} : 1 \rightarrow \text{Bool}\}$$


and equations:

$$\begin{aligned} \text{get}_{()}(\text{get}_{()}(p, q), r) &= \text{get}_{()}(p, r) = \text{get}_{()}(p, \text{get}_{()}(q, r)) \\ \text{get}_{()}(p, p) &= p \end{aligned}$$


When specifying equations in the algebraic approach, operations are typically written in a continuation passing style (CPS), exposing their algebraic structure. Thus $\text{get}_{()}(p, q)$ corresponds to a term that a functional programmer would typically write in direct style as **let** $x \leftarrow \text{get}()$ **in if** x **then** p **else** q . This continuation passing style corresponds directly to a view of algebraic computations as (possibly infinitely branching) trees, such that:

- nodes are labelled with operations;
- there is an edge labelled with each possible return value in the domain of the operation associated with a parent node;
- each such edge is connected to the corresponding continuation of the computation;
- leaves are labelled with final return values; and
- trees are quotiented modulo the equations.

For example, the CPS term $\text{get}_{()}(\text{get}_{()}(1, 2), 3)$ is given by the computation tree:



which by the first equation is equivalent to the CPS term $\text{get}_{()}(1, 3)$, whose computation tree is:



Following our previous work [8] on handlers for algebraic effects, we consider only algebraic effects for which there are no equations. We call these *abstract effects* and algebraic computations over them *abstract computations*. Thus an abstract computation is a plain unquotiented tree.

Abstract effects are closely related to *monads*, which Moggi successfully advocated [18] as a tool for modelling the semantics of effectful computation. Indeed, an abstract effect over an effect signature Σ is exactly the free monad construction over the functor generated by Σ [25].

2.2 Idioms are Oblivious, Arrows are Meticulous, Monads are Promiscuous

In previous work [14], we analysed the relative expressiveness of abstract idiom, arrow, and monad computations, proving that abstract idiom computations are less expressive than abstract arrow computations, which in turn are less expressive than abstract monad computations. We also gave an informal characterisation of the differences in terms of *control flow* and *data flow*, which we will

	control flow	data flow
idiom	static	static
arrow	static	dynamic
monad	dynamic	dynamic

Table 1. Flow for Idioms, Arrows, and Monads


now develop into a crisp characterisation in terms of constraints on abstract computation trees.

We say that data flow is dynamic if the value passed to an operation can depend on the results of prior operations. We say that control flow is dynamic if which operation to invoke (or whether to invoke an operation at all) can depend on the results of prior operations. The nature of data flow and control flow for idioms, arrows, and monads is given in Table 1. Idioms are entirely static. Arrows have static control flow but dynamic data flow. Monads are entirely dynamic.


Now we consider how this characterisation relates to abstract computation trees. First let us define some effect signatures:

$$\text{GB} = \{\text{get} : 1 \rightarrow \text{Bool} \quad \text{beep} : 1 \rightarrow 1\} \quad \text{State } S = \{\text{get} : 1 \rightarrow S \quad \text{put} : S \rightarrow 1\}$$

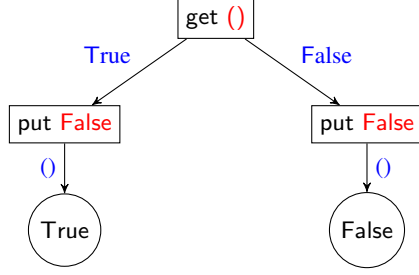
Monad trees are just abstract computation trees, as monads impose no restrictions on dependencies. For example, the CPS term, $\text{noisy} = \text{get}(\text{True}, \text{beep}(\text{False}))$, over effect signature GB, represents monad tree:



Arrow trees are those monad trees for which control flow is static. Concretely, this means that the tree must be completely balanced, and at each level of the tree each node must be labelled with the same operation (though the parameters to the operations may differ). For example, the CPS term, $\text{flip} = \text{get}(\text{put}_{\text{False}}(\text{True}), \text{put}_{\text{True}}(\text{False}))$, over effect signature State Bool, represents the arrow tree:



Idiom trees are those arrow trees for which not only is the control flow static, but so is the data flow. Concretely, this means that at each level of the tree the parameters of the operations on each node are identical. For example, the CPS term, $\text{reset} = \text{get}(\text{put}_{\text{False}}(\text{True}), \text{put}_{\text{False}}(\text{False}))$, over effect signature State Bool, represents the idiom tree:



Freedom Just as abstract monadic effects correspond exactly with the free monad construction over an effect signature, so abstract arrow effects correspond exactly with the free arrow construction over an effect signature, and abstract idiom effects correspond exactly with the free idiom construction over an effect signature.

Strange Computations An obvious omission from Table 1 is the case where control flow is dynamic, but data flow is static. We are not aware of a corresponding structure in the literature, and it seems rather strange to have dynamic control flow without dynamic data flow as well, so we will refer to such computations as *strange*. We can give a characterisation of strange computations in terms of abstract computation trees. *Strange trees* are those abstract computation trees for which control flow is dynamic, but data flow is static. Concretely, this means that the tree must be completely balanced, and at each level of the tree the parameters of the operations on each node must be identical.

Memory Flow From the structure of abstract computation trees, it is apparent that there is one other place that results may flow to: the leaves of the tree. We refer to this kind of flow as *memory*. We say that memory flow is dynamic if the values at the leaves depend on the results of prior computations and static if they do not. Idiom, arrow, and monad computations all have dynamic memory flow. One might conceive of explicitly distinguishing effectful computations that have static memory flow. The leaves of the corresponding computation trees are all the same; hence they must always have the same return value. This notion does not seem particularly useful, as we can always achieve the same behaviour via computations with unit return type paired up with the constant return value.

2.3 Flow Effects

In order to distinguish static and dynamic control and data flow, we will introduce special effect annotations which we call *flow effects*. The two flow effects are *c*, for dynamic control flow, and *d*, for dynamic data flow. We explain how they fit into λ_{flow} in Section 4.

3. An Effect Calculus

As a starting point for a calculus of effects and handlers for idioms, arrows, and monads, we take the λ_{eff} -calculus [8], which provides an effect type system and operational semantics for standard monadic algebraic effects and effect handlers, providing a foundation for generic effectful programming. In this section, we recapitulate the details of λ_{eff} .

3.1 Abstract Effects

In this subsection, we introduce Core λ_{eff} , the fragment of λ_{eff} that describes abstract effects in isolation from their handlers. This sub-language allows us to write abstract computations over an effect signature. We deviate slightly from our previous presentation [8]. Apart from superficial differences in lexical syntax, we omit computation products and unit, as they are orthogonal to the current work, and we adopt direct-style (as opposed to CPS) operations as primitive.

The syntax of types is as follows:

(values)	$A, B ::= 1 \mid A_1 \times A_2$ $\quad \quad \quad \mid 0 \mid A_1 + A_2$ $\quad \quad \quad \mid \{C\}_E$
(computations)	$C ::= [A] \mid A \rightarrow C$
(effect signatures)	$E ::= \{\text{op} : A \rightarrow B\} \uplus E \mid \emptyset$
(environments)	$\Gamma ::= x_1 : A_1, \dots, x_n : A_n$

Following Levy’s call-by-push-value [10], types are partitioned into value types (A, B) and computation types (C). The primary benefit we gain from a call-by-push-value approach is that it makes an explicit distinction between supplying an argument to a function and forcing a suspended computation. Effects are associated only with the latter action.

Value types (V, W) comprise unit (1), product ($A_1 \times A_2$), empty (0), sum ($A_1 + A_2$), and thunk types ($\{C\}_E$). In $\{C\}_E$, the computation type C is allowed to perform operations in the *effect signature* E . Computation types (C) comprise returners ($[A]$), which yields values of type A , and function types $A \rightarrow C$, from an argument of type A to a computation of type C . An effect signature comprises a collection of operation signatures $\{\text{op}_i : A_i \rightarrow B_i\}_i$. Operations in a signature must have distinct names, but order is not important. Type environments (Γ) are standard.

Intuitively, one can think of a returner type $[A]$ as being inhabited by computation trees. The structure of such trees can be quite rich. For a start, given a node for operation $\text{op} : A \rightarrow B$, it has n children, where n is the number of inhabitants of B . If we simply add a base type of natural numbers, then this means that our trees can be infinitely wide. More interestingly, an operation may take a suspended returner computation as a parameter, which yields a kind of “hyper tree” structure in which internal nodes may themselves contain further trees. Amongst other things, operations with computation parameters are useful for implementing a binary choice between two computations (see Section 6) and certain forms of aspect-oriented programming patterns [6, 12].

The syntax of terms is as follows:

(values)	$V, W ::= x \mid () \mid (V_1, V_2) \mid \text{inj}_i V \mid \{M\}$
(computations)	$M, N ::= \text{split}(V, x_1.x_2.M) \mid \text{case}_0(V)$ $\quad \quad \quad \mid \text{case}(V, x_1.M_1, x_2.M_2) \mid V!$ $\quad \quad \quad \mid \text{return } V \mid \text{let } x \leftarrow M \text{ in } N$ $\quad \quad \quad \mid \lambda x.M \mid M V$ $\quad \quad \quad \mid \text{op } V$

As with types, the terms are partitioned into value terms and computation terms. Value terms comprise variables (x), unit ($()$), pairs ((V_1, V_2)), injections ($\text{inj}_i V$), and thunks ($\{M\}$). Value terms are inert, in that all computation takes place in computation terms. Thus, all of the value constructs apart from variables are introduction forms. Computation terms comprise elimination forms for pairs ($\text{split}(V, x_1.x_2.M)$), the empty type ($\text{case}_0(V)$), sum types ($\text{case}(V, x_1.M_1, x_2.M_2)$), and thunks ($V!$), introduction ($\text{return } V$) and elimination ($\text{let } x \leftarrow M \text{ in } N$) forms for returners, introduction ($\lambda x.M$) and elimination ($M V$) forms for functions, and operation applications ($\text{op } V$).

The typing rules for Core λ_{eff} are given in Figure 1. The value judgement $\Gamma \vdash V : A$ asserts that value term V has type A in type environment Γ . The computation judgement $\Gamma \vdash_E M : C$ asserts that computation term M has type C with effects E in type environment Γ . The small-step operational semantics for Core λ_{eff} is given in Figure 2.

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

$$\begin{array}{c} \text{VAR} \\ \frac{(x : A) \in \Gamma}{\Gamma \vdash x : A} \end{array} \quad \begin{array}{c} \text{UNIT} \\ \frac{}{\Gamma \vdash () : 1} \end{array} \quad \begin{array}{c} \text{PAIR} \\ \frac{\Gamma \vdash V_1 : A_1 \quad \Gamma \vdash V_2 : A_2}{\Gamma \vdash (V_1, V_2) : A_1 \times A_2} \end{array} \quad \begin{array}{c} \text{INJ}_i \\ \frac{\Gamma \vdash V : A_i}{\Gamma \vdash \text{inj}_i V : A_1 + A_2} \end{array} \quad \begin{array}{c} \text{THUNK} \\ \frac{\Gamma \vdash_E M : C}{\Gamma \vdash \{M\} : \{C\}_E} \end{array}$$

$$\boxed{\Gamma \vdash_E M : C}$$

$$\begin{array}{c} \text{SPLIT} \\ \frac{\Gamma \vdash V : A_1 \times A_2 \quad \Gamma, x_1 : A_1, x_2 : A_2 \vdash_E M : C}{\Gamma \vdash_E \text{split}(V, x_1.x_2.M) : C} \end{array} \quad \begin{array}{c} \text{CASEZERO} \\ \frac{}{\Gamma \vdash_E \text{case}_0(V) : C} \end{array} \quad \begin{array}{c} \text{CASE} \\ \frac{\Gamma \vdash V : A_1 + A_2 \quad \Gamma, x_1 : A_1 \vdash_E M_1 : C \quad \Gamma, x_2 : A_2 \vdash_E M_2 : C}{\Gamma \vdash_E \text{case}(V, x_1.M_1, x_2.M_2) : C} \end{array} \quad \begin{array}{c} \text{RETURN} \\ \frac{\Gamma \vdash V : A}{\Gamma \vdash_E \text{return } V : [A]} \end{array}$$

$$\begin{array}{c} \text{LET} \\ \frac{\Gamma \vdash_E M : [A] \quad \Gamma, x : A \vdash_E N : C}{\Gamma \vdash_E \text{let } x \leftarrow M \text{ in } N : C} \end{array} \quad \begin{array}{c} \text{ABS} \\ \frac{}{\Gamma \vdash_E \lambda x.M : A \rightarrow C} \end{array} \quad \begin{array}{c} \text{APP} \\ \frac{\Gamma \vdash_E M : A \rightarrow C \quad \Gamma \vdash V : A}{\Gamma \vdash_E M V : C} \end{array} \quad \begin{array}{c} \text{FORCE} \\ \frac{}{\Gamma \vdash_E V! : C} \end{array} \quad \begin{array}{c} \text{OP} \\ \frac{(\text{op} : A \rightarrow B) \in E \quad \Gamma \vdash V : A}{\Gamma \vdash_E \text{op } V : [B]} \end{array}$$

Figure 1. Typing Rules for Core λ_{eff}

$$\begin{array}{l} (\beta. \times) \quad \text{split}((V_1, V_2), x_1.x_2.M_1) \longrightarrow M[V_1/x_1, V_2/x_2] \\ (\beta. +) \quad \text{case}(\text{inj}_i V, x_1.M_1, x_2.M_2) \longrightarrow M_i[V/x_i] \\ (\beta. \{ \}) \quad \{M\}! \longrightarrow M \end{array}$$

$$\begin{array}{l} (\beta. []) \quad \text{let } x \leftarrow \text{return } V \text{ in } M \longrightarrow M[V/x] \\ (\beta. \rightarrow) \quad (\lambda x.M) V \longrightarrow M[V/x] \end{array}$$

$$\frac{M \longrightarrow N}{E[M] \longrightarrow E[N]}$$

$$E ::= [] \mid E V \mid \text{let } x \leftarrow E \text{ in } N$$

Figure 2. Operational Semantics for Core λ_{eff}

$$\begin{array}{l} \text{noisy} : \{[\text{Bool}]\}_{\text{GB}} \\ \text{noisy} = \{\text{let } x \leftarrow \text{get}() \text{ in} \\ \quad \text{if } x \text{ then return } x \text{ else (beep()); return } x\} \end{array}$$

$$\begin{array}{l} \text{flip} : \{[\text{Bool}]\}_{\text{State Bool}} \\ \text{flip} = \{\text{let } x \leftarrow \text{get}() \text{ in} \\ \quad \text{let } y \leftarrow \neg x \text{ in} \\ \quad \text{put}(y); \text{return } x\} \end{array}$$

$$\begin{array}{l} \text{reset} : \{[\text{Bool}]\}_{\text{State Bool}} \\ \text{reset} = \{\text{let } x \leftarrow \text{get}() \text{ in} \\ \quad \text{put}(\text{False}); \text{return } x\} \end{array}$$

Figure 3. Core λ_{eff} examples

Syntactic Sugar We will use the following syntactic sugar:

$$\begin{array}{l} M; N \equiv \text{let } x \leftarrow M \text{ in } N, \quad x \text{ fresh} \\ \text{Bool} \equiv 1 + 1 \\ \text{True} \equiv \text{inj}_1 () \\ \text{False} \equiv \text{inj}_2 () \end{array}$$

$$\begin{array}{l} \text{if } V \text{ then } M \text{ else } N \equiv \text{case}(V, x.M, y.N), \quad x, y \text{ fresh} \\ \neg V \equiv \text{if } V \text{ then return True else return False} \end{array}$$

Examples Figure 3 presents the four example abstract computations from Section 2 as λ_{eff} terms.

3.2 Effect Handlers

Core λ_{eff} allows us to write abstract computations over arbitrary effect signatures. An effect handler provides an *interpretation* of an abstract computation.

We extend the grammar for Core λ_{eff} as follows.

Handler types

$$R ::= A \xrightarrow{E} \xrightarrow{E'} C$$

Handlers

$$H ::= \{\text{return } x \mapsto M\} \mid H \uplus \{\text{op } p \ k \mapsto N\}$$

Handling

$$M ::= \dots \mid \text{handle } M \text{ with } H$$

A handler type $A \xrightarrow{E} \xrightarrow{E'} C$ interprets a returner computation of type $[A]$ with effects E as a computation of type C with effects

E' . A handler is defined by a return clause $\text{return } x \mapsto M$ and a collection of operation clauses $\{\text{op}_i \ p \ k \mapsto N_i\}_i$. The return clause defines how to handle final return values. The returned value is bound to the variable x in M . The operation clauses define how to handle each operation. The operation parameter is bound to p , and the continuation is bound to k in N . Providing direct access to the whole continuation allows it to be used non-linearly, which is important for implementing effects such as exceptions, and non-deterministic choice, for instance.

For any handler:

$$H = \{\text{return } x \mapsto M\} \uplus \{\text{op}_i \ p \ k \mapsto N_i\}_i$$

we define the action of H on return values as follows:

$$H(\text{return}, V) = M[V/x]$$

and the action of H on operations handled by H as follows:

$$H(\text{op}_i, V, W) = N_i[V/p, W/k]$$

The typing rules for handlers are given in Figure 4. The operational semantics for handlers is given in Figure 5.

Reifying Handlers In essence, the type $A \xrightarrow{E} \xrightarrow{E'} C$ behaves like a suspended function of type $\{\{[A]\}_E \rightarrow C\}_{E'}$. Indeed we can reify a handler H as a value as follows:

$$\frac{\Gamma \vdash H : A \xrightarrow{E} \xrightarrow{E'} C}{\Gamma \vdash \{\lambda x. \text{handle } x! \text{ with } H\} : \{\{[A]\}_E \rightarrow C\}_{E'}}$$

$\Gamma \vdash_E M : C$	$\Gamma \vdash H : A \Rightarrow^{E'} C$
<p style="text-align: center;">HANDLE</p> $\frac{\Gamma \vdash_E M : [A] \quad \Gamma \vdash H : A \Rightarrow^{E'} C}{\Gamma \vdash_{E'} \mathbf{handle} M \mathbf{with} H : C}$	<p style="text-align: center;">HANDLER</p> $\frac{E = \{\text{op}_i : A_i \rightarrow B_i\}_i \quad H = \{\mathbf{return} x \mapsto M\} \uplus \{\text{op}_i p k \mapsto N_i\}_i \quad [\Gamma, p : A_i, k : \{B_i \rightarrow C\}_{E'} \vdash_{E'} N_i : C]_i \quad \Gamma, x : A \vdash_{E'} M : C}{\Gamma \vdash H : A \Rightarrow^{E'} C}$

Figure 4. Typing Rules for Handlers

$(\mathit{handle}.\square) \mathbf{handle} (\mathbf{return} V) \mathbf{with} H \longrightarrow H(\mathbf{return}, V)$ $(\mathit{handle}.\text{op}) \mathbf{handle} D[\text{op} V] \mathbf{with} H \longrightarrow H(\text{op}, V, \{\lambda z. \mathbf{handle} D[\mathbf{return} z] \mathbf{with} H\})$
<p>(delimited computation contexts) $D ::= [] \mid D V \mid \mathbf{let} x \leftarrow D \mathbf{in} N$</p> <p>(evaluation contexts) $E ::= [] \mid E V \mid \mathbf{let} x \leftarrow E \mathbf{in} N \mid \mathbf{handle} E \mathbf{with} H$</p>

Figure 5. Operational Semantics for Handlers

Examples As a basic example, consider a handler for stateful computations:

```
evalState : A State S =>^0 (S -> [A])
evalState = return x -> \s. return x
           get p k -> \s. k s s
           put p k -> \s. k () p
```

This handler interprets a stateful computation of type $[A]$ over state type S as a pure function of type $S \rightarrow [A]$. The state is threaded through the computation such that it can be read and written using `get` and `put`, but is discarded when the computation returns. Now if we instantiate S to `Bool`, then we can apply `evalState` to flip:

handle flip! with evalState

yielding a function of type $\text{Bool} \rightarrow [\text{Bool}]$ that flips its argument.

An important property of abstract computations is that they are generic in the sense that they do not commit to a particular interpretation of the operations. For instance, we can define alternative handlers for state. A mild variation on the `evalState` handler is a handler that interprets `get` and `put` in the same way, but returns the final state along with the final return value of the computation.

```
runState : A State S =>^0 (S -> [A x S])
runState = return x -> \s. return (x, s)
           get p k -> \s. k s s
           put p k -> \s. k () p
```

A slightly more interesting variation is the following handler, which logs each put operation, returning the list of all values written to the state cell in a list alongside the final return value.

```
logState : A State S =>^0 (S -> [A x List S])
logState = return x -> \s. return (x, Nil)
           get p k -> \s. k s s
           put p k -> \s. let z -< k () p in
                        split(z, x.ss.(x, Cons s ss))
```

For this example, we have assumed a standard extension of λ_{eff} with list types `List X` and corresponding list constructors `Nil` : `List X` and `Cons` : `X -> List X -> List X`.

3.3 Effect Forwarding

In our prior work [8], we considered a number of practical extensions and variations on handlers, including shallow handlers, parameterised handlers, open handlers, effect forwarding, and effect polymorphism. Perhaps the most important extension is effect forwarding in conjunction with open handlers.

The idea is that an open handler handles all of the operations explicitly specified by its type, but it also supports other operations, which are forwarded to be handled by an outer handler. A key advantage of open handlers is that they compose. We might, for instance, define an open handler for state and an open handler for exceptions. We can then handle a computation that uses state, exceptions, and possibly other effects as well by first handling it with the state handler, and then handling the resulting computation with the exception handler, or vice-versa.

It is straightforward to adapt λ_{eff} to support open handlers. The typing rule for open handlers is:

<p style="text-align: center;">OPENHANDLER</p> $\frac{E = E' \oplus \{\text{op}_i : A_i \rightarrow B_i\}_i \quad H = \{\mathbf{return} x \mapsto M\} \uplus \{\text{op}_i p k \mapsto N_i\}_i \quad [\Gamma, p : A_i, k : \{B_i \rightarrow C\}_{E'} \vdash_{E'} N_i : C]_i \quad \Gamma, x : A \vdash_{E'} M : C}{\Gamma \vdash H : A \Rightarrow^{E'} C}$
--

The only difference from the vanilla `HANDLER` rule is that the input effects are $E' \oplus E$ instead of just E , where $E' \oplus E$ is the extension of E' by E (where any clashes are resolved in favour of E).

As far as the semantics goes, we simply extend the action of a handler to forward operations with no operation clause:

$$H(\text{op}, V, W) = \mathbf{let} x \leftarrow \text{op} V \mathbf{in} W! x, \quad \text{op} \neq \text{op}_i \text{ for any } i$$

4. Flow Effects

In this section we introduce `Core λ_{flow}` , a variation of `Core λ_{eff}` that supports abstract idiom, arrow, and monad computations using a uniform syntax, and thus supporting generic effectful programming over all three kinds of effectful computation. The grammar of types is as follows:

(values)	$A, B ::= 1 \mid A_1 \times A_2 \mid 0 \mid A_1 + A_2 \mid \{C\}_E$
(computations)	$C ::= [A] \mid A \rightarrow C$
(effect signatures)	$E ::= \{\text{op} : A \rightarrow B\} \uplus E \mid \{f\} \uplus E \mid \emptyset$
(flow effects)	$f ::= c \mid d$
(environments)	$\Gamma ::= \Gamma, x : A \mid \Gamma, x^* : A \mid \cdot$

The differences are highlighted in grey. As well as the usual operations, effect signatures can also include flow effects `c` and `d`, which denote dynamic control and data flow. In λ_{flow} , type environments distinguish between *active* ($x^* : A$) and *inactive* ($x : A$) variables. Only active variables can be directly used. The flow effects allow inactive variables to be activated appropriately in order to realise dynamic control and data flow.

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

$$\frac{\text{VAR}^*}{(x^* : A) \in \Gamma} \Gamma \vdash x : A$$

$$\frac{\text{UNIT}}{\Gamma \vdash () : 1}$$

$$\frac{\text{PAIR}}{\Gamma \vdash V_1 : A_1 \quad \Gamma \vdash V_2 : A_2} \Gamma \vdash (V_1, V_2) : A_1 \times A_2$$

$$\frac{\text{INJ}_i}{\Gamma \vdash V : A_i} \Gamma \vdash \text{inj}_i V : A_1 + A_2$$

$$\frac{\text{THUNK}}{\Gamma \vdash_E M : C} \Gamma \vdash \{M\} : \{C\}_E$$

$$\boxed{\Gamma \vdash_E M : C}$$

$$\frac{\text{SPLIT}^*}{\Gamma \vdash_E \text{split}(V, x_1.x_2.M) : C} \Gamma^* \vdash V : A_1 \times A_2 \quad \Gamma, x_1 : A_1, x_2 : A_2 \vdash_E M : C$$

$$\frac{\text{CASEZERO}^*}{\Gamma \vdash_E \text{case}_0(V) : C} \Gamma^* \vdash V : 0$$

$$\frac{\text{CASE}}{\Gamma \vdash_E \text{case}(V, x_1.M_1, x_2.M_2) : C} \Gamma \vdash V : A_1 + A_2 \quad \Gamma, x_1 : A_1 \vdash_E M_1 : C \quad \Gamma, x_2 : A_2 \vdash_E M_2 : C$$

$$\frac{\text{RETURN}^*}{\Gamma \vdash_E \text{return } V : [A]} \Gamma^* \vdash V : A$$

$$\frac{\text{LET}}{\Gamma \vdash_E \text{let } x \leftarrow M \text{ in } N : C} \Gamma \vdash_E M : [A] \quad \Gamma, x : A \vdash_E N : C$$

$$\frac{\text{ABS}}{\Gamma \vdash_E \lambda x.M : A \rightarrow C} \Gamma, x : A \vdash_E M : C$$

$$\frac{\text{APP}^*}{\Gamma \vdash_E M V : C} \Gamma \vdash_E M : A \rightarrow C \quad \Gamma^* \vdash V : A$$

$$\frac{\text{FORCE}}{\Gamma \vdash_E V! : C} \Gamma \vdash V : \{C\}_E$$

$$\frac{\text{OP}}{\Gamma \vdash_E \text{op } V : [B]} (\text{op} : A \rightarrow B) \in E \quad \Gamma \vdash V : A$$

$$\frac{\text{OP}^*}{\Gamma \vdash_E \text{op } V : [B]} d \in E \quad (\text{op} : A \rightarrow B) \in E \quad \Gamma^* \vdash V : A$$

$$\frac{\text{FORCE}^*}{\Gamma \vdash_E V! : C} c, d \in E \quad \Gamma^* \vdash V : \{C\}_E$$

$$\frac{\text{RETURN}^*}{\Gamma \vdash_E \text{return } M : [A]} E' \subseteq \{c, d\} \quad \Gamma^* \vdash_{E'} M : [A]$$

Figure 6. Typing Rules for Core λ_{flow}

We define two meta operations on type environments. The first, *activation* (Γ^*), activates all of the variables in Γ .

$$\begin{aligned} \cdot^* &= \cdot \\ (\Gamma, x : A)^* &= \Gamma^*, x^* : A \\ (\Gamma, x^* : A)^* &= \Gamma^*, x^* : A \end{aligned}$$

The second, *flushing* (Γ^\dagger), removes all of the inactive variables from Γ .

$$\begin{aligned} \cdot^\dagger &= \cdot \\ (\Gamma, x : A)^\dagger &= \Gamma^\dagger \\ (\Gamma, x^* : A)^\dagger &= \Gamma^\dagger, x^* : A \end{aligned}$$

The typing rules for Core λ_{flow} are given in Figure 6. The differences from Core λ_{eff} are again highlighted in grey. The variable rule restricts access to active variables.

$$\frac{\text{VAR}^*}{(x^* : A) \in \Gamma} \Gamma \vdash x : A$$

The application rule, activates all variables in the argument value.

$$\frac{\text{APP}^*}{\Gamma \vdash_E M V : C} \Gamma \vdash_E M : A \rightarrow C \quad \Gamma^* \vdash V : A$$

This is always sound because β -reduction will always bind the argument value to an inactive variable. The elimination rules for products and empty sums are amended in order to activate the type environment in the eliminated value.

$$\frac{\text{SPLIT}^*}{\Gamma \vdash_E \text{split}(V, x_1.x_2.M) : C} \Gamma^* \vdash V : A_1 \times A_2 \quad \Gamma, x_1 : A_1, x_2 : A_2 \vdash_E M : C$$

$$\frac{\text{CASEZERO}^*}{\Gamma \vdash_E \text{case}_0(V) : C} \Gamma^* \vdash V : 0$$

This is sound because decomposing a product or an empty sum has no impact on control or data flow. The rule for returning a value activates all variables in the value.

$$\frac{\text{RETURN}^*}{\Gamma \vdash_E \text{return } V : [A]} \Gamma^* \vdash V : A$$

Activating the type environment supports dynamic memory flow. In order to support dynamic data flow, we add a variant of the rule for operations that only applies if the d effect is present.

$$\frac{\text{OP}^*}{\Gamma \vdash_E \text{op } V : [B]} d \in E \quad (\text{op} : A \rightarrow B) \in E \quad \Gamma^* \vdash V : [A]$$

As well as the standard **return** V construct, we introduce a special **return** M construct, which yields the value returned by the pure returner computation M .

$$\frac{\text{RETURN}^*}{\Gamma \vdash_E \text{return } M : [A]} E' \subseteq \{c, d\} \quad \Gamma^* \vdash_{E'} M : [A]$$

This rule allows final return values to be computed from any variables in the type environment using an arbitrary pure computation (E' can only include flow effects, so it must be pure). Finally, we include a special rule for forcing thunks, that only applies to thunked computations with the c and d flow effects.

$$\frac{\text{FORCE}^*}{\Gamma \vdash_E V! : C} c, d \in E \quad \Gamma^* \vdash V : \{C\}_E$$

This rule activates all of the variables in a type environment when forcing a thunk.

If $c, d \in E$ then the following derivation applies:

$$\frac{\Gamma^* \vdash_E M : C}{\Gamma \vdash_E \{M\}! : C} \text{THUNK} \quad \frac{\Gamma^* \vdash \{M\} : \{C\}_E}{\Gamma \vdash_E \{M\}! : C} \text{FORCE}^*$$

Hence in the presence of all flow effects we can systematically activate all variables in the type environment and λ_{flow} degenerates into λ_{eff} , which explains why both the data flow and control flow effects appear in the (FORCE^*) rule.

The operational semantics for λ_{flow} is given in Figure 7. The only differences from Core λ_{eff} (highlighted in grey) are the additional evaluation context for computing inside returned values and

$$\begin{array}{l}
(\beta.\times) \quad \mathbf{split}((V_1, V_2), x_1.x_2.M_1) \longrightarrow M[V_1/x_1, V_2/x_2] \\
(\beta.+) \quad \mathbf{case}(\mathbf{inj}_i V, x_1.M_1, x_2.M_2) \longrightarrow M_i[V/x_i] \\
(\beta.\{\}) \quad \{M\}! \longrightarrow M \\
(\beta.[]) \quad \mathbf{let } x \leftarrow \mathbf{return } V \mathbf{ in } M \longrightarrow M[V/x] \\
(\beta.\rightarrow) \quad (\lambda x.M) V \longrightarrow M[V/x] \\
(\mathit{ret.ret}) \quad \mathbf{return } (\mathbf{return } V) \longrightarrow \mathbf{return } V \\
\hline
\frac{M \longrightarrow N}{E[M] \longrightarrow E[N]} \\
E ::= [] \mid E V \mid \mathbf{let } x \leftarrow E \mathbf{ in } N \mid \mathbf{return } E
\end{array}$$

Figure 7. Operational Semantics for Core λ_{flow}

the $(\mathit{ret.ret})$ rule for converting the result of a pure computation into a plain value.

Examples Figure 8 presents the four example abstract computations of Section 2 and Figure 3 as λ_{flow} terms. The terms `noisy` and `flip` differ from the corresponding terms in Figure 3. The term `reset` is identical because it does not involve dynamic data or control flow.

In λ_{flow} we explicitly list the c and d effects in the type of `noisy` and we explicitly list the d effect in the type of `flip`. The λ_{eff} version of `noisy` does not type check in λ_{flow} because x is not active in the conditional test position. Thunking and forcing the conditional activates it. The λ_{eff} version of `flip` does not type check in λ_{flow} because x is not active in the negation. We take advantage of λ_{flow} 's ability to treat pure computations like values in order to activate it.

5. Handling Flow

In this section we present effect handlers for λ_{flow} . As there are two flow effects (c and d), there are four possible kinds of computation we might try to handle: monads ($\{c, d\}$), arrows ($\{d\}$), idioms (\emptyset), and something strange ($\{c\}$). As λ_{flow} does not have adequate support for writing computations of the latter kind, we will only consider handlers for the other three kinds.

As we have the inclusions $\emptyset \subseteq \{d\} \subseteq \{c, d\}$, monad handlers can handle arrow and idiom computations, and arrow handlers can handle idiom computations. However, there exist interpretations of arrow computations that cannot be specified using monad handlers and interpretations of idiom computations that cannot be specified using arrow or monad handlers.

The typing rules for monad, arrow, and idiom handlers are given in Figure 9. The operational semantics is given in Figure 10. We use the syntactic sugar:

$$\lambda(x, y).M \equiv \lambda z.\mathbf{split}(z, x.y.M), \quad z \text{ fresh}$$

We now describe in detail the design of the different kinds of handler. Each kind of handler has the same syntax as standard handlers. The differences are in the typing rules, operational semantics, and the **handle** constructs.

5.1 Monad Handlers

We already know how to handle arbitrary monadic computations. The typing rules are the same as for standard handlers, except the effects of a handled computation may additionally include arbitrary flow effects. As abstract arrow and idiom computations are just restricted monad computations, it is sound for a monad handler to handle an abstract arrow or idiom computation. The operational semantics is unchanged. We annotate monad handlers and monad handler types with a \top subscript.

$$\begin{array}{l}
\mathit{noisy} : \{[\mathbf{Bool}]\}_{\text{GB} \cup \{c, d\}} \\
\mathit{noisy} = \{\mathbf{let } x \leftarrow \mathbf{get}() \mathbf{ in} \\
\quad \{\mathbf{if } x \mathbf{ then return } x \mathbf{ else (beep()); return } x\}!\} \\
\mathit{flip} : \{[\mathbf{Bool}]\}_{\text{State Bool} \cup \{d\}} \\
\mathit{flip} = \{\mathbf{let } x \leftarrow \mathbf{get}() \mathbf{ in} \\
\quad \mathbf{let } y \leftarrow \mathbf{return } \neg x \mathbf{ in} \\
\quad \mathbf{put}(y); \mathbf{return } x\} \\
\mathit{reset} : \{[\mathbf{Bool}]\}_{\text{State Bool}} \\
\mathit{reset} = \{\mathbf{let } x \leftarrow \mathbf{get}() \mathbf{ in} \\
\quad \mathbf{put}(\mathbf{False}); \mathbf{return } x\}
\end{array}$$

Figure 8. Core λ_{flow} examples

5.2 Arrow Handlers

The challenge of adapting conventional effect handlers to interpret arrow computations is that each operation clause must bind a continuation representing the rest of the computation, but in general this continuation need not inhabit the usual function space. Indeed, a key feature of arrows is that they abstract over computations with input and output, that is continuations, in such a way that the input need not be provided up-front.

For a monad T , a continuation is exactly a *Kleisli arrow* of type $A \rightarrow T B$ (or $A \rightarrow [B]$ in λ_{flow}). But arrows generalise Kleisli arrows, and thus continuations. For instance, for any state type S , a state transformer [19] of type:

$$\{S \rightarrow [A]\}_{\{c, d\}} \rightarrow \{[S \rightarrow [B]]\}_{\{c, d\}}$$

is an arrow with input type A and output type B .

Recall that an effect handler is a compositional interpreter for an abstract computation. An arrow handler provides an interpretation of an arrow computation with an input and an output. Arrow handler syntax is exactly the same as standard handler syntax. Arrow handler types do however differ from standard handler types.

We let X range over type variables, which we will use in order to ensure that arrow handlers are parametric in the input type. This is crucial, as it allows us to manually thread a context through arrow computations.

Arrow handler types have the following shape:

$$R_{\rightsquigarrow} ::= A^E \Rightarrow_{\rightsquigarrow} G$$

where G is a unary type operator mapping types to types. The idea is that this interprets an abstract arrow computation of type $X \rightarrow [A]$ as a computation of type $G X$, where X can be instantiated at any type, and the interpretation is parametric in X . We will sometimes write type operators as type-level lambda functions of the form $\lambda X.C$, where the type variable X is bound in C .

The (ARROWHANDLE) rule describes how to handle an arrow computation with an arrow handler. The type environment Γ is flushed in M meaning that all dynamic input to the computation must be packaged up in z . As the handled computation has an input z , it is written as a lambda $\lambda z.M$. The return type of the conclusion is $G B$.

The (ARROWHANDLER) rule follows a similar structure to the (MONADHANDLER) rule. The key differences arise because arrow handlers handle computations with inputs. Thus the type of x in the return clause is a function from the input type X to A and similarly the type of the parameter p in an operation clause is a function from X to A_i . The computation type of the continuation k is $G(X \times B_i)$ instead of $B_i \rightarrow C$. The idea is that G models the type of an arrow computation: the argument to G is the input type, and the result

$$\boxed{\Gamma \vdash_E M : C}$$

$$\dots$$

$$\frac{\text{MONADHANDLE} \quad \Gamma \vdash_E M : [A] \quad \Gamma \vdash H : A \xrightarrow{E} \xrightarrow{\top}^{E'} C}{\Gamma \vdash_{E'} \mathbf{handle}_{\top} M \mathbf{with} H : C} \quad \frac{\text{ARROWHANDLE} \quad \Gamma^\dagger, z : B \vdash_E M : [A] \quad \Gamma \vdash H : A \xrightarrow{E} \xrightarrow{\infty}^{E'} G}{\Gamma \vdash_{E'} \mathbf{handle}_{\infty} (\lambda z.M) \mathbf{with} H : G B}$$

$$\frac{\text{IDIOMHANDLE} \quad \Gamma^\dagger, z : B \vdash_E M : [A] \quad \Gamma \vdash H : A \xrightarrow{E} \xrightarrow{\top}^{E'} G}{\Gamma \vdash_{E'} \mathbf{handle}_! (\lambda z.M) \mathbf{with} H : G B}$$

$$\boxed{\Gamma \vdash H : A \xrightarrow{E} \xrightarrow{\top}^{E'} C}$$

$$\text{MONADHANDLER}$$

$$\frac{E = \{\text{op}_i : A_i \rightarrow B_i\}_i \cup \{f_j\}_j \quad H = \{\mathbf{return} x \mapsto M\} \uplus \{\text{op}_i p k \mapsto N_i\}_i \quad [\Gamma, p : A_i, k : \{B_i \rightarrow C\}_{E'} \vdash_{E'} N_i : C\}_i \quad \Gamma, x : A \vdash_{E'} M : C}{\Gamma \vdash H : A \xrightarrow{E} \xrightarrow{\top}^{E'} C}$$

$$\boxed{\Gamma \vdash H : A \xrightarrow{E} \xrightarrow{\infty}^{E'} G}$$

$$\text{ARROWHANDLER}$$

$$\frac{X \text{ fresh} \quad E = \{\text{op}_i : A_i \rightarrow B_i\}_i \cup \{f_j\}_j \quad c \notin E \quad H = \{\mathbf{return} x \mapsto M\} \uplus \{\text{op}_i p k \mapsto N_i\}_i \quad [\Gamma, p : \{X \rightarrow [A_i]\}_{E'}, k : \{G(X \times B_i)\}_{E'} \vdash_{E'} N_i : G X\}_i \quad \Gamma, x : \{X \rightarrow [A]\}_{E'} \vdash_{E'} M : G X}{\Gamma \vdash H : A \xrightarrow{E} \xrightarrow{\infty}^{E'} G}$$

$$\boxed{\Gamma \vdash H : A \xrightarrow{E} \xrightarrow{!}^{E'} G}$$

$$\text{IDIOMHANDLER}$$

$$\frac{X \text{ fresh} \quad E = \{\text{op}_i : A_i \rightarrow B_i\}_i \cup \{f_j\}_j \quad c, d \notin E \quad H = \{\mathbf{return} x \mapsto M\} \uplus \{\text{op}_i p k \mapsto N_i\}_i \quad [\Gamma, p : A_i, k : \{G(X \times B_i)\}_{E'} \vdash_{E'} N_i : G X\}_i \quad \Gamma, x : \{X \rightarrow [A]\}_{E'} \vdash_{E'} M : G X}{\Gamma \vdash H : A \xrightarrow{E} \xrightarrow{!}^{E'} G}$$

Figure 9. Typing Rules for λ_{flow} Handlers

$$\begin{aligned} (\mathit{handle}_{\top}.[]) & \quad \mathbf{handle}_{\top} (\mathbf{return} V) \mathbf{with} H \longrightarrow H(\mathbf{return}, V) \\ (\mathit{handle}_{\top}.\text{op}) & \quad \mathbf{handle}_{\top} D[\text{op} V] \mathbf{with} H \longrightarrow H(\text{op}, V, \{\lambda x. \mathbf{handle} D[\mathbf{return} x] \mathbf{with} H\}) \\ (\mathit{handle}_{\infty}.[]) & \quad \mathbf{handle}_{\infty} (\lambda z. \mathbf{return} P) \mathbf{with} H \longrightarrow H(\mathbf{return}, \{\lambda z. \mathbf{return} P\}) \\ (\mathit{handle}_{\infty}.\text{op}) & \quad \mathbf{handle}_{\infty} (\lambda z. D[\text{op} V]) \mathbf{with} H \longrightarrow H(\text{op}, \{\lambda z. \mathbf{return} V\}, \{\mathbf{handle}_{\infty} (\lambda(z, x). D[\mathbf{return} x]) \mathbf{with} H\}) \\ (\mathit{handle}_!.[]) & \quad \mathbf{handle}_! (\lambda z. \mathbf{return} P) \mathbf{with} H \longrightarrow H(\mathbf{return}, \{\lambda z. \mathbf{return} P\}) \\ (\mathit{handle}_!.\text{op}) & \quad \mathbf{handle}_! (\lambda z. D[\text{op} V]) \mathbf{with} H \longrightarrow H(\text{op}, V, \{\mathbf{handle}_! (\lambda(z, x). D[\mathbf{return} x]) \mathbf{with} H\}) \end{aligned}$$

$$\begin{aligned} (\text{value or computation term}) & \quad P ::= V \mid M \\ (\text{delimited computation contexts}) & \quad D ::= [] \mid D V \mid \mathbf{let} x \leftarrow D \mathbf{in} N \\ (\text{evaluation contexts}) & \quad E ::= [] \mid E V \mid \mathbf{let} x \leftarrow E \mathbf{in} N \mid \mathbf{return} E \\ & \quad \mid \mathbf{handle}_{\top} E \mathbf{with} H \mid \mathbf{handle}_{\infty} (\lambda z.E) \mathbf{with} H \mid \mathbf{handle}_! (\lambda z.E) \mathbf{with} H \end{aligned}$$

Figure 10. Operational Semantics for λ_{flow} Handlers

of applying G to a type is the output type. In the continuation, the current input type is paired up with the return type of the operation.

The operational semantics is similar to that for monadic handlers. The differences are all related to explicitly threading the context through the handler. When handling a return clause ($\mathit{handle}_{\infty}.[]$), the returned term is a function of the input. (Note that we must account for the possibility that P is a computation term M as z may appear free in M , in which case $\mathbf{return} M$ may be stuck.) When handling an operation ($\mathit{handle}_{\infty}.\text{op}$), the parameter is a function of the input, and the continuation extends the context with the return value of the operation.

It is not sound to apply an arrow handler to a monadic computation and thus the (ARROWHANDLER) rule prevents arrow handlers from being applied to computations with dynamic control flow. Concretely, this constraint ensures that closed terms do not get stuck. For instance, it disallows the term:

$$\mathbf{handle}_{\infty} (\lambda z. \{\mathbf{case}(z, x_1.M_1, x_2.M_2)\}!) \mathbf{with} H$$

which reduces to the stuck term:

$$\mathbf{handle}_{\infty} (\lambda z. \mathbf{case}(z, x_1.M_1, x_2.M_2)) \mathbf{with} H$$

As every abstract idiom computation is just a restricted abstract arrow computation, it is perfectly sound to handle an abstract idiom computation with an arrow handler.

Reifying Arrow Handlers Just as standard handlers can be reified as values, so can arrow handlers. In essence, for any type X , the arrow handler type $A \xrightarrow{E} \xrightarrow{E'} G$ behaves like a suspended function of type $\{\{X \rightarrow [A]\}_E \rightarrow GX\}_{E'}$. Indeed, for any type X , we can reify an arrow handler H as a value as follows:

$$\frac{\Gamma \vdash H : A \xrightarrow{E} \xrightarrow{E'} G}{\Gamma \vdash \{\lambda y. \mathbf{handle}_{\rightsquigarrow} (\lambda x. y! x) \mathbf{with} H\} : \{\{X \rightarrow [A]\}_E \rightarrow GX\}_{E'}}$$

Examples We can systematically transform any monad handler into a corresponding arrow handler. For instance, `evalState` can be rewritten as an arrow handler as follows.

$$\begin{aligned} \mathbf{evalState}_{\rightsquigarrow} &: A \text{ State } S \omega \{d\} \xrightarrow{\rightsquigarrow} \{c,d\} \lambda X. X \rightarrow S \rightarrow [A] \\ \mathbf{evalState}_{\rightsquigarrow} &= \mathbf{return} x \mapsto x! \\ &\quad \mathbf{get} p k \mapsto \lambda z. \lambda s. k! (z, s) s \\ &\quad \mathbf{put} p k \mapsto \lambda z. \lambda s. \mathbf{let} p' \leftarrow p z \mathbf{in} k! (z, ()) p' \end{aligned}$$

However, we can also provide other interpretations that are not possible with monad handlers. As a simple, albeit slightly contrived, example, we can write an alternative arrow handler for state that determines statically whether a computation performs a put operation or not.

$$\begin{aligned} \mathbf{puts}_{\rightsquigarrow} &: A \text{ State } S \omega \{d\} \xrightarrow{\rightsquigarrow} \{c,d\} \lambda X. [\mathbf{Bool}] \\ \mathbf{puts}_{\rightsquigarrow} &= \mathbf{return} x \mapsto \mathbf{return} \mathbf{False} \\ &\quad \mathbf{get} p k \mapsto k! \\ &\quad \mathbf{put} p k \mapsto \mathbf{return} \mathbf{True} \end{aligned}$$

As a computation handled by `puts` must be an abstract arrow computation, in which the operations performed are independent of the values returned by intermediate operations, it is possible to compute the result statically without actually supplying any state. This is not possible for any monad handler for state as the number of times that `put` is performed is in general a dynamic property in a monadic computation.

Let us now consider an effect signature for failure:

$$\mathbf{Fail} = \{\mathbf{fail} : \forall X. 1 \rightarrow X\}$$

The universal quantifier indicates that `fail` is a polymorphic operation. Effectively it defines an infinite family of operations parameterised by X . Correspondingly, an operation clause for `fail` defines an infinite family of operation clauses parameterised by X . Polymorphic operations and corresponding polymorphic operation clauses are a straightforward extension that applies just as well to λ_{flow} as it does λ_{eff} [8]. We could give `fail` the signature $1 \rightarrow 0$, but we choose to make it polymorphic because doing so is more convenient for writing example code. We can handle failure as an option type in the standard way:

$$\begin{aligned} \mathbf{maybe}_{\rightsquigarrow} &: A \text{ Fail} \omega \{d\} \xrightarrow{\rightsquigarrow} \{c,d\} \lambda X. [\{X \rightarrow A\}_{\{c,d\}} + 1] \\ \mathbf{maybe}_{\rightsquigarrow} &= \mathbf{return} x \mapsto \mathbf{return} (\mathbf{inj}_1 x) \\ &\quad \mathbf{fail} p k \mapsto \mathbf{return} (\mathbf{inj}_2 ()) \end{aligned}$$

Alternatively, we can write a non-standard handler that in the event of failure counts the total number of failures (a quantity which is fixed for an abstract arrow computation, but not for an abstract monad computation):

$$\begin{aligned} \mathbf{fails}_{\rightsquigarrow} &: A \text{ Fail} \omega \{d\} \xrightarrow{\rightsquigarrow} \{c,d\} \lambda X. [\{X \rightarrow A\}_{\{c,d\}} + \mathbf{Nat}] \\ \mathbf{fails}_{\rightsquigarrow} &= \mathbf{return} x \mapsto \mathbf{return} (\mathbf{inj}_1 x) \\ &\quad \mathbf{fail} p k \mapsto \mathbf{let} r \leftarrow k! \mathbf{in} \\ &\quad \quad \mathbf{case}(r, \\ &\quad \quad \quad x. \mathbf{return} (\mathbf{inj}_2 (S Z))) \\ &\quad \quad \quad n. \mathbf{return} (\mathbf{inj}_2 (S n)), \end{aligned}$$

where we assume a data type of natural numbers `Nat` with constructors `S` and `Z`. If we define:

$$\begin{aligned} \mathbf{twoFail} &: \{[A \times B]\}_{\mathbf{Fail} \omega \{d\}} \\ \mathbf{twoFail} &= \{\mathbf{let} x \leftarrow \mathbf{fail} () \mathbf{in} \mathbf{let} y \leftarrow \mathbf{fail} () \mathbf{in} \mathbf{return} (x, y)\} \end{aligned}$$

then:

$$(\mathbf{handle}_{\rightsquigarrow} (\lambda z. \mathbf{twoFail}!) \mathbf{with} \mathbf{maybe}_{\rightsquigarrow}) () \longrightarrow^* \mathbf{return} ()$$

and:

$$(\mathbf{handle}_{\rightsquigarrow} (\lambda z. \mathbf{twoFail}!) \mathbf{with} \mathbf{fails}_{\rightsquigarrow}) () \longrightarrow^* \mathbf{return} (S (S Z))$$

5.3 Idiom Handlers

The `IDIOMHANDLER` rule is similar to the `ARROWHANDLER` rule. The difference is that the context is not threaded through parameters in operation clauses — directly capturing the property that idiom computations do not have the data flow effect. The $(\mathbf{handle}_{e_1}. [])$ rule is identical to the $(\mathbf{handle}_{\rightsquigarrow}. [])$ rule. The $(\mathbf{handle}_{e_1}. \mathbf{op})$ rule is similar to the $(\mathbf{handle}_{\rightsquigarrow}. \mathbf{op})$ rule; the only difference is that the context is not threaded through operation parameters.

It is not sound to apply an idiom handler to an abstract monad or arrow computation and thus the `(IDIOMHANDLER)` rule prevents idiom handlers from being applied to computations with dynamic data flow. Concretely, this constraint ensures that closed terms do not become open. For instance, if we were to allow dynamic data flow then we could write:

$$M = \mathbf{handle}_{e_1} (\lambda y. \mathbf{op} y) \mathbf{with} H$$

Now suppose:

$$H(\mathbf{op}, y, \{\mathbf{handle}_{e_1} ((\lambda(z, x). D[\mathbf{return} x])) \mathbf{with} H\}) = \mathbf{return} y$$

then M reduces to the open term `return y`.

Reifying Idiom Handlers Just as monadic and arrow handlers can be reified as values, so can idiom handlers. For any type X , we can reify an idiom handler H as a value as follows:

$$\frac{\Gamma \vdash H : A \xrightarrow{E} \xrightarrow{E'} G}{\Gamma \vdash \{\lambda y. \mathbf{handle}_{e_1} (\lambda x. y! x) \mathbf{with} H\} : \{\{X \rightarrow [A]\}_E \rightarrow GX\}_{E'}}$$

Examples Just as we can transform any monad handler into an corresponding arrow handler, we can systematically transform any arrow handler (or monad handler) into a corresponding idiom handler. For instance, `evalState` and `puts` can be written as idiom handlers as follows.

$$\begin{aligned} \mathbf{evalState}_{e_1} &: A \text{ State } S \xrightarrow{e_1} \{c,d\} \lambda X. X \rightarrow S \rightarrow [A] \\ \mathbf{evalState}_{e_1} &= \mathbf{return} x \mapsto \lambda z. \lambda s. x! z \\ &\quad \mathbf{get} p k \mapsto \lambda z. \lambda s. k! (z, s) s \\ &\quad \mathbf{put} p k \mapsto \lambda z. \lambda s. k! (z, ()) p \end{aligned}$$

$$\begin{aligned} \mathbf{puts}_{e_1} &: A \text{ State } S \xrightarrow{e_1} \{c,d\} \lambda X. [\mathbf{Bool}] \\ \mathbf{puts}_{e_1} &= \mathbf{return} x \mapsto \mathbf{return} \mathbf{False} \\ &\quad \mathbf{get} p k \mapsto k! \\ &\quad \mathbf{put} p k \mapsto \mathbf{return} \mathbf{True} \end{aligned}$$

The only difference between the definition of `evalState` and `evalStatee1` is that the latter does not apply p to the context in the clause that handles `put`, as it cannot depend on the context. The definitions of `puts` and `putse1` are identical because the operation parameters are never used in either.

Here is another idiom handler for state that statically sums all of the values that will be written to the state cell by an abstract computation:

$$\begin{aligned} \mathbf{putSum}_{e_1} &: A \text{ State } \mathbf{Nat} \xrightarrow{e_1} \{c,d\} \lambda X. [\mathbf{Nat}] \\ \mathbf{putSum}_{e_1} &= \mathbf{return} x \mapsto \mathbf{return} Z \\ &\quad \mathbf{get} p k \mapsto k! \\ &\quad \mathbf{put} p k \mapsto \mathbf{let} x \leftarrow k! \mathbf{in} \mathbf{return} p + x \end{aligned}$$

This is not possible for any arrow handler for state, as in an abstract arrow computation the parameters to put operations can be dynamic.

For the following example, we assume λ_{flow} has been extended with a List data type with constructors Nil and Cons and String \equiv List char. Let us consider a variation on failure in which an error message is reported alongside failure:

$$\text{Error} = \{\text{Err} : \forall X. \text{String} \rightarrow X\}$$

We can straightforwardly write idiom handlers that have analogous behaviour to `maybe~` and `fails~`. We can also define a handler that, in the case of error, aggregates all possible error messages.

$$\begin{aligned} \text{errs}_1 &: A \xrightarrow{\text{Error}}_{\{c,d\}} \lambda X. [\{X \rightarrow A\}_{\{c,d\}} + (\text{List String})] \\ \text{errs}_1 &= \text{return } x \mapsto \text{return } (\text{inj}_1 x) \\ \text{Err } s \ k &\mapsto \text{let } r \leftarrow k! \text{ in} \\ &\quad \text{case}(r, \\ &\quad \quad ss. \text{return } (\text{inj}_2 (\text{Cons } s \ ss)), \\ &\quad \quad x. \text{return } (\text{inj}_2 (\text{Cons } s \ \text{Nil}))) \end{aligned}$$

This is not possible for any arrow handler for state, as in an abstract arrow computation the parameters to Err operations can be dynamic.

5.4 Forwarding for Arrow and Idiom Handlers

As the continuation of an operation handled by an arrow or an idiom handler may not be in the standard function space, we cannot use the universal definition of forwarding that works for monadic handlers. One possibility is to add a special forwarding clause:

$$\text{default } p \ k \mapsto N'$$

where

$$H(\text{op}, V, W) = N'[\{\lambda x. \text{op } x\} / \text{default}, V/p, W/k], \\ \text{op} \neq \text{op}_i \text{ for any } i$$

We leave it to future work to investigate how this idea pans out in practice.

6. Example: Parser Combinators

As a larger example, we now build a collection of generic effectful parser combinators. We choose parser combinators [24] as there exist monad, arrow and idiom variations, each with different practical tradeoffs. Monad parser combinators are the most general, admitting fully context-sensitive grammars. Arrow parser combinators admit a small amount of context-sensitivity, for instance supporting parsing of XML, where a close tag must match the corresponding open tag. Idiom parser combinators capture context-free grammars.

In λ_{flow} each of these can be constructed using the same signature of parser operations. An important benefit of restricting the power of parser combinators is that doing so admits more efficient implementations. Rather than focusing on the intricacies of such implementations (which are discussed in detail elsewhere [24]), we present the necessary infrastructure to support generic parsers, without committing to a particular concrete implementation.

We make use of a number of extensions to λ_{flow} in order to make the example more realistic: recursive functions, pattern matching, effect polymorphism, and effect forwarding. We also make use of additional syntactic sugar for writing code in a more concise call-by-value style.

- In the presence of control and data flow effects we automatically activate the entire type environment (so we never need to write $\{M\}!$).
- We allow `return` V to be written as simply V when doing so is not ambiguous.

- We allow forcing to be omitted when applying a function to an argument, writing $V \ W$ for $V! \ W$.
- We allow top-level definitions to be written directly without curly braces. For instance: `id x = x` means `id = $\{\lambda x. x\}$` .
- We allow computation terms to appear anywhere a value term is normally required within a computation term, in which case we desugar by binding such computations to values in left-to-right order. For instance $M \ N$ means `let x \leftarrow N in M x` and (M, N) means `let x \leftarrow M in let y \leftarrow N in return (x, y)`, where x, y are fresh variables.

Let us begin with some primitive parsing operations:

$$\text{any} : 1 \rightarrow \text{Char} \quad \text{char} : \text{Char} \rightarrow \text{Char}$$

The any operation parses any character and the char operation parses a specific character. We do not need to define a special operation for the constant parser as it is simply `return`. Similarly, we do not need a special operation for sequential composition of parsers, as let binding gives us this functionality for free. For instance, the following invokes parsers p and q in sequence and then combines their results by function application:

$$\text{let } f \leftarrow p \text{ in let } x \leftarrow q \text{ in return } (f \ x)$$

We make use of a failure operation:

$$\text{fail} : \forall X. 1 \rightarrow X$$

along with an operation to choose between two parsers:

$$\text{choose} : \forall X. (\{[X]\}_E, \{[X]\}_E) \rightarrow X$$

where we will need to instantiate E appropriately. The idea is that this operation takes two parsers returning type X as arguments, tries running both, and then returns an answer of type X ; exactly how it does so depends on how the operation is handled. A handler might return a list of all matches, or it might just return the first match; it might run the parsers sequentially, or it might run them in parallel. The effect signature E must include all of the operations we intend to support for parsing — including `choose` itself. (The `fail` and `choose` operations together allow us to variously implement the behaviour of the Haskell type classes `Alternative`, `ArrowChoice`, and `MonadPlus`, which add additional monoidal structure respectively to idioms, arrows, and monads.)

Let us now define a suitable effect signature, `Parse`:

$$\begin{aligned} \text{Parse } \varepsilon &= \{\text{any} : 1 \rightarrow \text{Char}, \\ &\quad \text{char} : \text{Char} \rightarrow \text{Char} \\ &\quad \text{fail} : \forall X. 1 \rightarrow X \\ &\quad \text{choose} : \forall X. (\{[X]\}_{\text{Parse } \varepsilon}, \{[X]\}_{\text{Parse } \varepsilon}) \rightarrow X\} \\ &\quad \uplus \varepsilon \end{aligned}$$

It is parameterised by an effect variable ε that can be instantiated with any of the flow effects in order to choose between idiom, arrow, and monad interpretations.

Using the `choose` operation and a recursive function call, we can define a composite parser that applies its argument zero or more times:

$$\begin{aligned} \text{many} &: \{\{[X]\}_{\text{Parse } \varepsilon} \rightarrow [\text{List } X]\}_{\text{Parse } \varepsilon} \\ \text{many } p &= \text{choose} (\text{Nil}, \text{Cons } p \ (\text{many } p)) \end{aligned}$$

We parse digits as follows:

$$\begin{aligned} \text{digit} &: \{\{\text{char}\}\}_{\text{Parse } \varepsilon} \\ \text{digit} &= \text{choose}(\text{char}'0', \\ &\quad \text{choose}(\text{char}'1', \dots, \text{choose}(\text{char}'9') \dots)) \end{aligned}$$

Let us suppose that we have a function for converting a sequence of digits into a natural number:

$$\text{digitsToNat} : \{\text{String} \rightarrow [\text{Nat}]\}_{\{c,d\}}$$

The following parser parses natural numbers:

```
num : { [Nat] }Parse ε
num = let x ← many digit in digitsToNat
```

Here is a function that parses a comma delimited list of natural numbers terminated by an at sign (@):

```
nums : { Nat → [List Nat] }Parse {c,d}
nums Z = char '0'; Nil
nums (S n) = let x ← num () in char ','; Cons x (nums n)
```

Because it performs a case split on its argument, `nums` is context-sensitive and can only be handled by a monad handler. However, by separating out the computation into two stages, we can define a variant that generates a context-free parser computation that can be handled by arrow or even idiom handlers:

```
nums' : { Nat → [[List Nat]] }Parse ε
nums' Z = {char '0'; Nil}
nums' (S n) = let ys ← nums' n in
  {let x ← num () in char ','; Cons x (ys!)}
```

The outer computation performs the recursion on n . The generated parser computation is independent of n .

Here is a context-sensitive parser that parses a single number delimited by a character read dynamically from the input:

```
delim : { [Nat] }Parse {d}
delim = let c ← any () in
  char ' '; let x ← num () in char ' '; char c; x
```

As c does not affect control flow, `delim` can be handled by an arrow handler.

Here is a context-sensitive parser that parses a number n followed by a list of n numbers:

```
nnums : { [Nat] }Parse {c,d}
nnums = let x ← num () in char ':'; nums n
```

This is a truly dynamic parser, that can only be handled by a monad handler.

In order to implement a handler for abstract parser computations, we will make use of effect signatures for reading characters, state, and failure:

```
Read = {getc : 1 → char}
State S = {get : 1 → S, put : S → 1}
Fail = {fail : ∀X. 1 → X}
```

Let us define an idiom handler for parsing (reified as a function):

```
parse1 : { { [X] }Parse ∅ → [X] }Read⊔Fail⊔{c,d}
parse1 =
  handle1 (λm.m!) with
    return x      ↦ x!
    any () k      ↦ let c ← getc () in λz.k (z, c)
    char c k      ↦ if c == getc () then λz.k (z, c)
                  else λz.fail ()
    fail () k     ↦ λz.fail ()
    choose (p, q) k ↦ case maybe (parse1 p) of
      Just v → λz.k (z, v)
      Nothing →
        case maybe (parse1 q) of
          Just v → λz.k (z, v)
          Nothing → λz.fail ()
```

The most interesting operation clause is the one for `choose`. Given two parsers p and q , we first try to parse with p , and if that fails then we try with q . In order to test for failure, we recursively invoke `parse` and handle the result with `maybe`, a straightforward handler that interprets failure as an option data type `Maybe X` with

constructors `Just` and `Nothing`:

```
maybe : { {X} }Fail⊔{c,d}⊔ε → [Maybe X] }{c,d}⊔ε
maybe m = handle1 m! with
  return x ↦ Just x
  fail () k ↦ Nothing
```

We now define a handler for reading characters:

```
read : { { [X] }Read⊔{c,d}⊔ε → [X] }State String⊔Fail{c,d}⊔ε
read cs m = handle1 m! with
  return x ↦ x
  getc () k ↦ case get () of Nil      ↦ fail ()
                  Cons c cs ↦ put cs; c
```

This handler interprets `getc` in terms of state and failure. We can handle state with `evalState`:

```
state : { {X} }State S⊔{c,d}⊔ε → S → [X] }{c,d}
state m = handle1 m with evalState
```

and use `maybe` for failure. Finally we compose all of these handlers together:

```
runIdiomParser : { { [X] }Parse ∅ → String → [Maybe X] }{c,d}
runIdiomParser p cs = state (maybe (read (parse1 p))) cs
```

Now:

```
let p ← nums' 3 in runIdiomParser p "42,21,7,@"
```

evaluates to `Just (Cons 42 (Cons 21 (Cons 7 Nil)))` and:

```
let p ← nums' 4 in runIdiomParser p "42,21,7,@"
```

evaluates to `Nothing`.

We can abstract over the `parse1` handler, in particular allowing it to be replaced by arrow or monad handlers:

```
runParser : { { { [X] }Parse ε → [X] }Read⊔Fail⊔{c,d} →
  { [A] }Parse ε → String → [Maybe X] }{c,d}
runParser h p cs = state (maybe (read (h p))) cs
```

Let us define arrow and monad analogues of `parse1`:

```
parse↪ : { { [X] }Parse {d} → [X] }Read⊔Fail⊔{c,d}
parse↪ =
  handle↪ (λm.m!) with
    return x      ↦ x!
    any () k      ↦ let c ← getc () in λz.k (z, c)
    char c k      ↦ λz.if c == getc () then k (z, c)
                  else fail ()
    fail () k     ↦ λz.fail ()
    choose (p, q) k ↦ λz.case maybe (parse↪ (p z)) of
      Just v → λz.k (z, v)
      Nothing →
        case maybe (parse↪ (q z)) of
          Just v → λz.k (z, v)
          Nothing → λz.fail ()
```

```
parse⊔ : { { [X] }Parse {c,d} → [X] }Read⊔Fail⊔{c,d}
parse⊔ = λm.
```

```
  handle⊔ m! with
    return x      ↦ x!
    any () k      ↦ let c ← getc () in k c
    char c k      ↦ if c == getc () then k c
                  else fail ()
    fail () k     ↦ fail ()
    choose (p, q) k ↦ case maybe (parse⊔ p) of
      Just v → k v
      Nothing →
        case maybe (parse⊔ q) of
          Just v → k v
          Nothing → fail ()
```

Now, both:

```
runParser parse↪ delim “* 42,21,7 *”
```

and:

```
runParser parse⊥ nnums “3:42,21,7,@”
```

type check and evaluate to Just (Cons 42 (Cons 21 (Cons 7 Nil))).

We can use `parse⊥` to handle any arrow or idiom parsers (such as the `nums'` and `delim` examples above). We can use `parse↪` to handle any idiom parsers (such as the `nums'` example above). In practice, we would write optimised handlers for arrow and idiom parsers [24]. A key advantage of using a language like λ_{flow} is that we can write code that is generic in the choice of concrete implementation while using a uniform direct-style syntax.

7. Related Work

There has been a recent spate of work on practical languages and libraries for effect handlers. Apart from our own libraries, Kiselyov et al [9] has a similar library for Haskell, and Brady [3] has an effect handlers library for his dependently-typed language Idris. Two programming languages that build in algebraic effects and handlers as primitives are Bauer and Pretnar's Eff [1, 2] and McBride's Frank [15, 16]. None of these systems support algebraic effects or handlers for idioms or arrows.

Capriotti and Kaposi explore free idioms [4]. Free idioms correspond to abstract idiom computations. Yallop's thesis [26, Chapter 2] provides an in-depth analysis of idioms, arrows, and monads, expanding on the work of Lindley et al [14], and characterising the normal forms for idioms and arrows. We have implemented both free idiom and free arrow constructions in Haskell [11] directly inspired by the normal forms of Yallop.

Petricek and Syme [20] describe a novel use of F# computation expression syntax to write idiom computations using `let` notation, which is partly inspired by syntax for formlets [5], an abstraction for building web forms that is an idiom.

8. Future Work

This paper focuses on the theory of algebraic effects and handlers for arrows and idioms. In order to evaluate the practice of algebraic effects and handlers for arrows and idioms we would like to build an implementation.

One can implement handlers on top of our existing free idiom and free arrow constructions in Haskell. However, programming with free idioms and free arrows in Haskell requires the programmer to use a different syntax. Idioms only support a pointless syntax. Arrows support a direct-style syntax, but it is not quite the same as the `do` notation used for monads. Given that F# computation expressions are already expressive enough to cover a range of computation types including monads and idioms, it might be interesting to try to use computation expressions as a basis for building a source language for λ_{flow} . It may, however, be difficult to adequately encode an effect type system on top of F#. Ultimately, we expect the most fruitful path may be to build a new language, or extend a custom language like Frank or Eff.

On the theoretical side, it would be interesting to explore denotational semantics for λ_{flow} and to consider how the story is affected by reintroducing equations to the picture. Another direction is to explore algebraic effects and handlers for other variations on the basic theme, such as for linear and dependent types.

Acknowledgments

We would like to thank Bob Atkey, Pierre-Evariste Dagand, and the anonymous referees for helpful feedback. This work was supported by EPSRC grants EP/J014591/1 and EP/K034413/1.

References

- [1] A. Bauer and M. Pretnar. Programming with algebraic effects and handlers. *CoRR*, abs/1203.1539, 2012.
- [2] 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.
- [3] E. Brady. Programming and reasoning with algebraic effects and dependent types. In *ICFP*. ACM, 2013.
- [4] P. Capriotti and A. Kaposi. Free applicative functors. *CoRR*, abs/1403.0749, 2014.
- [5] E. Cooper, S. Lindley, P. Wadler, and J. Yallop. The essence of form abstraction. In G. Ramalingam, editor, *APLAS*, volume 5356 of *Lecture Notes in Computer Science*, pages 205–220. Springer, 2008.
- [6] B. C. d. S. Oliveira, T. Schrijvers, and W. R. Cook. MRI: Modular reasoning about interference in incremental programming. *J. Funct. Program.*, 22(6):797–852, 2012.
- [7] J. Hughes. Generalising monads to arrows. *Sci. Comput. Program.*, 37(1-3):67–111, 2000.
- [8] O. Kammar, S. Lindley, and N. Oury. Handlers in action. In G. Morrisett and T. Uustalu, editors, *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] P. B. Levy. *Call-By-Push-Value: A Functional/Imperative Synthesis*, volume 2 of *Semantics Structures in Computation*. Springer, 2004.
- [11] S. Lindley. Free idioms and free arrows in Haskell, 2013. <https://github.com/slindley/dependent-haskell/tree/master/Free>.
- [12] S. Lindley. Aspect oriented programming with handlers, 2013. <https://github.com/slindley/effect-handlers/blob/master/Examples/AOP.hs>.
- [13] S. Lindley, P. Wadler, and J. Yallop. The arrow calculus. *J. Funct. Program.*, 20(1):51–69, 2010.
- [14] S. Lindley, P. Wadler, and J. Yallop. Idioms are oblivious, arrows are meticulous, monads are promiscuous. *Electr. Notes Theor. Comput. Sci.*, 229(5):97–117, 2011.
- [15] C. McBride. How might effectful programs look? In *Workshop on Effects and Type Theory*, 2007. <http://cs.ioc.ee/efftt/mcbride-slides.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] E. Moggi. Computational lambda-calculus and monads. In *LICS*, pages 14–23. IEEE Computer Society, 1989.
- [19] R. Paterson. A new notation for arrows. In B. C. Pierce, editor, *ICFP*, pages 229–240. ACM, 2001.
- [20] T. Petricek and D. Syme. The F# computation expression zoo. In M. Flatt and H.-F. Guo, editors, *PADL*, volume 8324 of *Lecture Notes in Computer Science*, pages 33–48. Springer, 2014.
- [21] G. D. Plotkin and J. Power. Adequacy for algebraic effects. In F. Honsell and M. Miculan, editors, *FoSSaCS*, volume 2030 of *Lecture Notes in Computer Science*, pages 1–24. Springer, 2001.
- [22] G. D. Plotkin and J. Power. Semantics for algebraic operations. *Electr. Notes Theor. Comput. Sci.*, 45:332–345, 2001.
- [23] G. D. Plotkin and M. Pretnar. Handling algebraic effects. *Logical Methods in Computer Science*, 9(4), 2013.
- [24] S. D. Swierstra. Combinator parsing: A short tutorial. In *LerNet ALFA Summer School*, volume 5520 of *Lecture Notes in Computer Science*, pages 252–300. Springer, 2008.
- [25] W. Swierstra. Data types à la carte. *J. Funct. Program.*, 18(4):423–436, 2008.
- [26] J. Yallop. *Abstraction for web programming*. PhD thesis, The University of Edinburgh, 2010.