THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Liberating Effects with Rows and Handlers

Citation for published version:

Hillerström, D & Lindley, S 2016, Liberating Effects with Rows and Handlers. in *TyDe 2016 Proceedings of the 1st International Workshop on Type-Driven Development*. ACM, pp. 15-27, 1st International Workshop on Type-Driven Development, Nara, Japan, 18/09/16. <https://doi.org/10.1145/2976022.2976033>

Digital Object Identifier (DOI):

[10.1145/2976022.2976033](https://doi.org/10.1145/2976022.2976033)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

TyDe 2016 Proceedings of the 1st International Workshop on Type-Driven Development

Publisher Rights Statement:

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 ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

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

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Liberating Effects with Rows and Handlers

Daniel Hillerström

The University of Edinburgh, UK
daniel.hillerstrom@ed.ac.uk

Sam Lindley

The University of Edinburgh, UK
sam.lindley@ed.ac.uk

Abstract

Algebraic effects and effect handlers provide a modular abstraction for effectful programming. They support user-defined effects, as in Haskell, in conjunction with direct-style effectful programming, as in ML. They also present a structured interface to programming with delimited continuations.

In order to be modular, it is necessary for an effect type system to support extensible effects. Row polymorphism is a natural abstraction for modelling extensibility at the level of types. In this paper we argue that the abstraction required to implement extensible effects and their handlers is exactly row polymorphism.

We use the Links functional web programming language as a platform to substantiate this claim. Links is a natural starting point as it uses row polymorphism for polymorphic variants, records, and its built-in effect types. It also has infrastructure for manipulating continuations. Through a small extension to Links we smoothly add support for effect handlers, making essential use of rows in the frontend and first-class continuations in the backend.

We demonstrate the usability of our implementation by modelling the mathematical game of Nim as an abstract computation. We interpret this abstract computation in a variety of ways, illustrating how rows and handlers support modularity and smooth composition of effectful computations.

We present a core calculus of row-polymorphic effects and handlers based on a variant of A-normal form used in the intermediate representation of Links. We give an operational semantics for the calculus and a novel generalisation of the CEK machine that implements the operational semantics, and prove that the two coincide.

Categories and Subject Descriptors D.3.3 [Language Constructs and Features]; F3.2 [Semantics of Programming Languages]

Keywords algebraic effects; effect handlers; effect typing; abstract machine semantics; operational semantics; delimited control

1. Introduction

Algebraic effects [29] and effect handlers [30] are a more modular alternative to monads for managing user-defined computational effects [6, 15, 17]. Effect handlers generalise exception handlers, providing a mechanism for interpreting arbitrary algebraic effects, and they present a structured interface to programming with delimited continuations.

As a simple example consider a choice effect given by a single effectful operation:

`Choose : Bool`

We can write an abstract computation M that invokes `Choose`, independently of specifying the meaning of `Choose`. We can then handle M in multiple ways. For instance, we can define a handler `allResults` that interprets the operation as nondeterministic choice, returning a list of every possible outcome of M . We can also define a different handler `coin` that interprets `Choose` as random choice, returning a single value that depends on all of the random choices made in M . Where effect handlers really come into their own as a programming abstraction is when we start composing them: we can handle some effects while forwarding all others, using row types to statically track the forwarded effects. We make extensive use of handler composition in Section 2.

Many existing implementations of effect handlers are Haskell libraries. Notable examples include the effect handlers library of Kammar et al. [15], the extensible effects library of Kiselyov et al. [17] and Kiselyov and Ishii [16], and implementations based on variants of Swierstra’s *data types à la carte* technique [34], such as the work of Wu et al. [36] on scoped effect handlers. Another notable effect handlers library is the Idris effects library [6]. Each of these libraries uses its own sophisticated encoding of an abstraction which amounts to a restricted form of row polymorphism. In this work we present the first, to our knowledge, implementation of effect handlers using actual Remy-style row polymorphism [33].

Links [8] is a functional programming language for building web applications. The defining feature of Links is that it provides a single source language that targets all three tiers of a web application: client, server, and database. Links source code is translated into an intermediate representation (IR) based on A-normal form [11]. For the client, the IR is compiled to JavaScript. For the server, the IR is interpreted using a variant of the CEK machine [10]. For the database, the IR is translated into an SQL query, taking advantage of the effect type system and the subformula property to guarantee query generation [22].

Links is a strict language with Hindley-Milner type inference. Links has a row type system for polymorphic variants, records, and its built-in effect types for concurrency and database integration [8, 22]. It also has support for manipulating first-class continuations, a feature that is central to implementing effect handlers.

The row-polymorphic effect type system and continuation support make Links a natural choice for experimenting with row-based algebraic effects and effect handlers. We have implemented an effect handlers extension to Links. Currently, it is supported only on the server-side. The frontend to our implementation makes essential use of row polymorphism, while the backend is implemented as a novel generalisation of the CEK machine.

Our main contributions are as follows.

- An implementation of effect handlers using Remy-style row polymorphism [33].

- A demonstration of the usability of our implementation illustrating how rows and handlers support modularity and smooth composition of effectful computations.
- A formalisation of our implementation including a small-step call-by-value operational semantics and an abstract machine semantics, based on a novel generalisation of the CEK machine to account for effect handlers.
- A strong correspondence proof between the small-step and abstract machine semantics: every reduction in the operational semantics corresponds to a sequence of administrative steps followed by a β -step in the abstract machine.

The rest of the paper is structured as follows. Section 2 gives a tutorial introduction to programming with handlers in Links. Section 3 presents a core calculus λ_{eff}^v along with a type-and-effect system and a small-step operational semantics. Section 4 relates the operational semantics to an abstract machine semantics, that captures the essence of our implementation. Section 5 discusses implementation details. Section 6 discusses related work. Finally, Section 7 concludes and discusses future work.

2. Programming with Handlers in Links

To demonstrate that rows and handlers provide an elegant and modular abstraction for effectful programming, we use a simplified version of the mathematical game Nim [5] as a running example.

Starting from an abstract representation of the game, we iteratively extend it with cheat detection and high score tracking capabilities through smooth composition of handlers, without needing to change the initial representation.

2.1 The Game of Nim and Effect Rows

The game of Nim is played between two players: Alice and Bob. The game begins with a heap of n sticks. The players alternate to take one, two, or three sticks from the heap. Alice makes the first move. The player who takes the last stick wins the game.

We abstract over the notion of making a move by defining it as an abstract effectful operation $\text{Move} : (\text{Player}, \text{Int}) \{\} \rightarrow \text{Int}$, where Player is variant type with two constructors Alice and Bob . The first parameter to Move is the active player, the second parameter is the current number of sticks on the heap. We will discuss the meaning of the braces $\{\}$ prefix on the arrow shortly. In Links, abstract operations like Move , are invoked using the do primitive, for instance

```
do Move(Alice, 3)
```

invokes the Move operation with values Alice and 3. Operation names, data constructors, and type aliases all begin with a capital letter. Records, variants, and effect signatures all have row types. All typing is structural in Links, thus it is unnecessary to declare a row occupant, such as an operation, before use. However, we consider it good practice to wrap the invocation of operations as functions. This is mainly because it lets us compose effects with functions seamlessly. Moreover, sometimes we want to do more than just invoking an operation. We will see an example of this in Section 2.4. We wrap Move as follows:

```
sig move :  
  (Player, Int) {Move: (Player, Int) \{\} \rightarrow Int | e \} \rightarrow Int  
fun move(p, n) {do Move(p, n)}
```

The syntax of Links is loosely based on that of JavaScript. The fun keyword begins a function definition (like function in JavaScript). Just as in JavaScript functions are n -ary, but they can also be curried. Unlike in JavaScript, functions are statically typed and the sig keyword begins a type signature. The function move invokes the operation Move with the parameters p and n .

In the type signature, the function arrow ($\rightarrow$) is prefixed by a row enclosed in curly braces. This row is the effect signature, or *effect row*, of the function. The presence of Move in the effect row indicates that the function may perform the Move operation. Furthermore, the effect row is equipped with an effect variable e , which can be instantiated with additional operations. This means that move may be invoked in the scope of additional effects. We say an effect row is *closed* if it has no effect variable, and *open* if it does. In general an effect row consists of an unordered collection of operation specifications and an optional effect variable. An operation specification either specifies that an operation is admissible (or *present*) and has a particular type signature, or that it is absent, or that it is polymorphic in its presence. We discuss the use of absence in Section 2.5.

The effect row on the type signature of the Move operation itself is empty, denoted by a pair of braces $\{\}$. This is always the case for abstract operations as any effects they might have are ultimately conferred by their handlers.

The Nim game is modelled as two mutually recursive functions aliceTurn and bobTurn . Here we show aliceTurn :

```
sig aliceTurn :  
  (Int) {Move: (Player, Int) \{\} \rightarrow Int | _ \} \rightarrow Player  
fun aliceTurn(n) {  
  if (n <= 0) Bob  
  else bobTurn(n - move(Alice, n)) }
```

The parameter n is the current number of sticks on the heap. If n is zero then Bob wins. Otherwise, Alice makes a move and it is now Bob's turn. The definition of bobTurn is completely symmetric, so we omit it here for brevity.

Two observations are worth making about the effect signature of aliceTurn . First, the effect variable is anonymous ($_$): type (or effect) variables need not be named when they appear only once. Second, the function arrow is squiggly ($\rightarrow$), which is syntactic sugar for denoting that the computation has the *wild* effect. The wild effect captures all intrinsic effects such as I/O, randomness, divergence, etc. To some extent it is analogous to the *IO monad* of Haskell, though the wild effect is much stricter as without it general recursion is disallowed. In our current implementation intrinsic effects cannot be handled by a handler, instead they are given a predefined interpretation by the interpreter.

Links employs a strict evaluation strategy, so we thunk computations that we wish to handle, and define the following type alias:

```
typename Comp(e :: Row, a) = () { | e \} \rightarrow a;
```

The keyword typename is used to define type aliases. The Comp type captures our notion of abstract computation, it is an alias for a thunk with an open effect row and return type a . The notation $e :: \text{Row}$ denotes that the type variable e has kind Row , which is the kind of effect variables.

The game function begins a game with a given number of sticks. Alice starts:

```
sig game : (Int) \rightarrow  
  Comp({Move: (Player, Int) \{\} \rightarrow Int | _ }, Player)  
fun game(n) () {aliceTurn(n)}
```

2.2 Strategies and Handlers

In general, algebraic effects come with equations [29], but as with most other implementations of effect handlers, we do not consider equations. Thus, on their own, abstract operations have no meaning; handlers give them a semantics. We can use handlers to encode particular strategies for Alice and Bob by interpreting the operation Move . We start by considering the *perfect* strategy, defined by $ps(n) \stackrel{\text{def}}{=} \max\{1, n \bmod 4\}$, where n is the number of sticks left

in the game. If player p adopts the perfect strategy, then p is guaranteed to win if on p 's turn n is not divisible by four. We define a handler `pp` (short for *perfect-vs-perfect*), that assigns the perfect strategy to both players

```
1 sig pp : (Comp({Move:(Player,Int) {}-> Int |e}, a)) ->
2           Comp({Move-
3             |e}, a)
4 fun pp(m) () {
5   handle(m) {
6     case Return(x) -> x
7     case Move(p,n,k) -> k(maximum(1, n 'mod' 4)) }}
```

We describe the handler line by line.

Lines 1 and 2 give the type of `pp`: it takes a computation, which may invoke the `Move` operation, and yields another computation where the operation is absent (denoted by `Move-`). The computation returns a value of type `a`. We may omit this type signature altogether as the type system is capable of inferring the appropriate types.

Lines 3 and 4 begin the definition. The curried function `pp` wraps the actual handler, that is applied to the argument m using the `handle` construct, which specifies how to interpret abstract operations through a sequence of clauses.

Line 5 is a *return clause*. It defines how to handle the final return value of the input computation. In this case, this value is simply returned as is.

Line 6 is an *operation clause*. It expresses how to handle `Move`. In general, an operation clause takes the form $Op(p_1, \dots, p_n, k) \rightarrow M$, where $p_1, \dots, p_n$ are patterns that bind the operation parameters and k is a pattern that binds the continuation of the computation in M . In this case p and n are bound to the active player and number of sticks in the heap, respectively. The continuation is invoked with the perfect strategy, irrespective of the player. The input type of the continuation is given by the output type of `Move`. The return type of the continuation is the return type of the handler.

The handler `pp` returns a computation. For convenience we define an auxiliary function to *run* (or force) a thunk:

```
sig run : (Comp({}, a)) {}-> a
fun run(m) { m() }
```

We can now compute the winner of a game in which both players play the perfect strategy:

```
links> run(pp(game(7)));
Alice : Player
links> run(pp(game(12)));
Bob : Player
```

Syntactic Sugar The input computation in `pp` is immediately supplied to `handle`. This abstract-over-handle idiom arises frequently, so Links provides syntactic sugar for it. We can give a more succinct definition of `pp` using the `handler` keyword:

```
handler pp {
  case Return(x) -> x
  case Move(_,n,k) -> k(maximum(1,n 'mod' 4)) }
```

We may choose to name the input computation by adding a parameter to the `handler` keyword in square brackets, for instance, `handler[m] pp {...}`. Naming the computation can be useful if we want to invoke the wrapper function recursively with the computation from inside the handler.

2.3 Game Trees and Multi-shot Continuations

The handler `pp` computes the winner of a particular game. It only considers one scenario in which both players play the same strategy, but we can use handlers to compute other data about a game.

Figure 1: Game Tree Generated by `run(gametree(game(3)))`.

For instance, we can give an interpretation that computes the *game tree*. Figure 1 shows an example game tree. Each node represents the active player, and each edge corresponds to a possible move for that player. We define a game tree inductively:

```
typename GTree = [ |Take:(Player, [(Int, GTree)])
                  |Winner:(Player) | ];
```

The syntax `[...]` denotes a (polymorphic) variant type in Links in which components of the variant type are delimited by the pipe symbol `|`. A `Take` node includes the active player and a list of possible moves, where each move is paired with the subsequent game tree. A `Winner` leaf denotes the winner of a game.

We define a handler `gametree` that generates game trees:

```
sig gametree :
  (Comp({Move:(Player,Int) {}-> Int |e}, Player)) ->
  Comp({Move-
    |e}, GTree)
handler gametree {
  case Return(x) -> Winner(x)
  case Move(p,n,k) ->
    var subgames = map(k, validMoves(n));
    var subtrees = zip([1,2,3], subgames);
    Take(p, subtrees) }
```

The effect signatures of `gametree` and `pp` are identical, though their interpretations of `Move` differ. The return clause wraps the winning player x in a leaf node. The operation clause for `Move` reifies the move as a node in the game tree. The `var` keyword denotes a let binding. The crucial part is the invocation of `map` which applies the continuation multiple times, once for each valid move, enumerating every possible subgame. The function `validMoves` is a simple filter:

```
fun validMoves(n)
  { filter(fun(m) {m <= n}, [1,2,3]) }
```

Figure 1 shows the game tree generated by the handler when $n = 3$.

2.4 Cheating and Forwarding

Thus far we have considered a single operation `Move`, but in general we may allow arbitrary algebraic operations. We could define a monolithic handler that interprets every operation that may occur in a computation. However, a more modular alternative is to define a series of fine-grained, specialised handlers that each handle a particular operation, and then compose them together to fully interpret a computation. Fortunately, handlers compose seamlessly. Composed handlers can cooperate to interpret an abstract computation. Each handler operates on a subset of the abstract operations, leaving the remainder for other handlers. Consequently, we obtain a considerable amount of flexibility as it becomes possible to reinterpret computations by swapping in and out individual handlers.

We defer a full discussion of the role that row polymorphism plays during composition of handlers until Section 2.5. In this section we omit type signatures for handlers and instead focus on the dynamic semantics of handler composition by augmenting the game model with a cheat detection mechanism. A cheating strategy might remove all remaining sticks from the heap, winning in a single move. We introduce an additional operation `Cheat` to signal that a cheater has been detected. The operation is parameterised by the player, who was caught cheating:

```
sig cheat : (Player) {Cheat:(Player) {}-> Zero|_}-> _
fun cheat(p) { switch (do Cheat(p)) { } }
```

The `Cheat` operation can never return a value as its return type is the empty type `Zero`. Thus invoking `Cheat` amounts to raising an exception. Concretely, an operation clause for `Cheat` can never invoke the continuation. The `switch(e){...}` construct pattern matches on the expression `e`, through a possibly empty list of clauses. We define an *exception* handler that interprets `Cheat` by outputting an error message and exiting the program:

```
handler report {
  case Return(x)      -> x
  case Cheat(Alice,_) -> error("Alice cheated!")
  case Cheat(Bob,_)   -> error("Bob cheated!")
}
```

We implement the heart of the cheat detection machinery as a handler:

```
handler checker {
  case Return(x)      -> x
  case Move(p,n,k)    -> var m = move(p,n);
                      if (m 'elem' validMoves(n)) k(m)
                      else cheat(p) }
}
```

In order to detect cheating, the handler analyses the active player's move. If it is legal then the game continues. Otherwise, the `Cheat` operation is invoked to signal that cheating has occurred. We may compose `pp` with `report` and `checker` to give an interpretation of a game in which no player may cheat. To make handler composition syntactically lightweight we define a pipeline operator (`-<-`) for composing handlers and another operator (`-<`) for applying a computation to a pipeline of handlers. Infix binary operators are defined using the `op` keyword in Links:

```
op f -<- g {fun(m) {f(g(m))}}
op f -< m {f(m)}
```

The operators are meant to indicate that unhandled operations are forwarded from right to left in a pipeline. In order to run a pipeline of handlers, we may apply the function `run`:

```
links> run -<- pp -<- report -<- checker -< game(7);
Alice : Player
```

The `Cheat` operation is never invoked as both players play the same legal strategy. Let us define another handler that assigns the perfect strategy to `Alice` and a cheating strategy to `Bob`

```
handler pc {
  case Return(x)      -> x
  case Move(Alice,n,k) -> k(maximum(1, n 'mod' 4))
  case Move(Bob,n,k)  -> k(n) }
}
```

Now the cheat detection handler catches Bob:

```
links> run -<- pc -<- report -<- checker -< game(7);
*** Fatal error : Bob cheated!
```

The order of composition is important as `pc` and `checker` both handle moves. Bob gets away with cheating if we swap the two handlers:

```
links> run -<- pp -<- report -<- checker
          -<- pc -< game(7);
Bob : Player
```

Here we also use `pp`, because the type system does not know that `checker` is not performing any `Move` operations.

2.5 Composition and Row Polymorphism

In this section we discuss the typing of composed handlers. First, consider the type signature for the `report` handler:

```
sig report :
  (Comp({Cheat:(Player) {}-> Zero|e1}, a)) ->
  Comp({Cheat{p} |e1}, a)
```

In general, a handler accepts a computation as input and produces another computation as output. Moreover, handlers have *open* input and output effect rows, which both share the same effect variable. As a consequence both rows mention the same operation names. However, some of these operation names may be marked as absent or polymorphic in their presence. In the output effect row of `report`, the syntax `Cheat{p}` denotes that the operation is presence polymorphic. The type variable `p` can be instantiated to be either present with a particular type (`:A`) or absent (`-`). Presence polymorphism is useful for seamless composition of handlers. We illustrate why by type checking the composition:

```
var f = run -<- (pp -<- report);
```

Recall the type signature for `pp`:

```
sig pp : (Comp({Move:(Player,Int) {}-> Int|e2}, a)) ->
  Comp({Move- |e2}, a)
```

The output effects of `report` must be compatible with the input effects of `pp`, therefore the composition gives rise to the following unification constraint:

$$\{Move:(Player,Int) {}-> Int|e2\} \sim \{Cheat\{p\}|e1\}$$

which is solved by introducing a fresh effect variable `e3`, instantiating `e1` to `{Move:(Player,Int) {}-> Int|e2}`, and `e2` to `{Cheat{p}|e3}`. Thus the unified type is:

$$\{Move:(Player,Int) {}-> Int, Cheat\{p\}|e3\}$$

Note that with rows the order of operations is unimportant. The output of `pp` must be compatible with the input to `run`, giving rise to the constraint:

$$\{Move-, Cheat\{p\}|e3\} \sim \{\}$$

which is solved by instantiating `p` to `-` and `e3` to `\{\}`. Thus `f` has type:

$$(Comp(\{Move:(Player,Int) {}-> Int, Cheat:(Player) {}-> Zero\}, a)) {}-> a$$

Now we consider the type signature for the `checker` handler:

```
sig checker :
  (Comp({Cheat:(Player) {}-> Zero,
        Move:(Player,Int) {}-> Int|e}, a)) ->
  Comp({Cheat:(Player) {}-> Zero,
        Move:(Player,Int) {}-> Int|e}, a)
```

The reason `Cheat:(Player) {}-> Zero` appears in the input effect is because `Cheat` is not handled, so if a `Cheat` operation is forwarded it must have the correct type.

Remy's IIML' What we actually require for soundness is that if the `Cheat` effect is present then it must have type `(Player) {}-> Zero`, as that is the type it has in the output. In a slightly more refined system along the lines of Remy's IIML' [33], we could specify this as follows:

```
sig checker :
  (Comp({Cheat{_: (Player) {}-> Zero,
    Move: (Player,Int) {}-> Int|e},a)) ->
  Comp({Cheat: (Player) {}-> Zero,
    Move: (Player,Int) {}-> Int|e},a)
```

In the input row `Cheat` may or may not be present, but if it is then it must have type `(Player) {}-> Zero`.

2.6 Choice and Built-in Effects

In this section we implement the choice effect described in the introduction. We let Bob choose which strategy he will adopt. First, we define a wrapper for the choice operation.

```
sig choose : Comp({Choose:Bool|_, Bool})
fun choose() {do Choose}
```

Using this operation we define a strategy selecting function in which Bob decides between playing the perfect or cheating strategy

```
fun bobChooses(m)()
{ if (choose()) pc(m)() else pp(m)() }
```

We can give a nondeterministic interpretation of `Choose` that infuses Bob with oracular powers that enable him to explore both alternatives. We define it as a handler `allResults`

```
sig allResults : (Comp({Choose:Bool|e},a)) ->
  Comp({Choose{_|e},[a])
handler allResults {
  case Return(x) -> [x]
  case Choose(k) -> k(true) ++ k(false) }
```

The handler wraps the result of the input computation into a singleton list. In the `Choose`-clause the handler accumulates the results of both the alternatives by invoking the continuation twice.

Now, we can put everything together:

```
links> run -<- allResults -<- bobChooses -< game(7);
[Bob,Alice] : [Player]
```

Thus Bob wins only when he cheats.

Alternatively, we can replace Bob's oracular powers with a fair coin and let him perform a coin flip to decide which strategy to pick. We use Links' built-in random number generator, which returns a floating point value between 0.0 and 1.0 (both inclusive):

```
sig coin : (Comp({Choose:Bool|e}, a)) ->
  Comp({Choose{_|e}, a)
handler coin {
  case Return(x) -> x
  case Choose(k) -> if (random() > 0.5) k(true)
    else k(false) }
```

The handler uniformly interprets `Choose` as `true` or `false`. Thus, using this handler Bob will be equally likely to play honestly or dishonestly. The computation

```
links> run -<- coin -<- bobChooses -< game(7);
```

returns either `Alice` or `Bob`. Built-in effects interact smoothly with the rest of the system.

2.7 A Scoreboard and Parameterised Handlers

As a final extension, we add a scoreboard that accumulates the number of wins for each player. The scoreboard is updated after each game. We represent state as an effect with operations for reading (`Get : s`) and updating (`Put : s {}-> ()`) a state of type `s`. We wrap them in the usual way:

```
sig get : () {Get:s|_}-> s
fun get() {do Get}

sig put : (s) {Put:(s) {}-> ()|_}-> ()
fun put(s) {do Put(s)}
```

We use a parameterised handler to give an interpretation of state. In addition to supplying a computation to a parameterised handler, we also supply one or more parameters. In this instance we pass the state as an additional parameter `s`

```
sig state : (s) -> (Comp({Get:s,Put:(s) {}-> ()|e},a))->
  Comp({Get{_|_},Put{_|_} |e},a)

handler state(s) {
  case Return(x) -> x
  case Get(k) -> k(s)(s)
  case Put(p,k) -> k(())(p) }
```

The main difference compared to an unparameterised handler is that the continuation `k` is a curried function that takes a return value followed by the handler parameters. In the `Get` clause, we return the state and also pass it unmodified to any subsequent invocations of the handler. Similarly, in the `Put` clause, we return unit and update the state.

We represent high scores as an association list and refer to a value of this type as the game state:

```
typename GState = [(Player,Int)];
```

We define an initial state `s0 = [(Alice,0),(Bob,0)]`. We now need a mechanism to update the game state when a game finishes. Recall that `game(n)` returns a computation whose type is:

```
Comp({Move:(Player,Int) {}-> Int|_, Player}).
```

The computation returns the winner of the game. We may exploit the fact that the return clauses of handlers are invoked in the order of composition. Therefore, we define a simple post-processing handler that contains only a `Return` case to update the scoreboard:

```
sig scoreUpdater :
  (Comp({Get:GState,Put:(GState) {}-> ()|e}, Player)) ->
  Comp({Get:GState,Put:(GState) {}-> ()|e}, Player)
handler scoreUpdater {
  case Return(x) -> var s = updateScore(x, get());
    put(s); x }
```

The function `updateScore` is pure it simply returns a copy of the given game state, in which the number of wins for the given player `p` has been incremented by one. The handler then reads and updates the game state. The composition `scoreUpdater(game(n))` causes the effect row to grow accordingly:

```
Comp({Move:(Player,Int) {}-> Int,
  Get:GState,Put:(GState) {}-> ()|_, Player}).
```

In a similar fashion, we define a handler that prints the scoreboard:

```
sig printer : (Comp({Get:GState|e}, a)) ->
  Comp({Get:GState|e}, a)
handler printer
{ case Return(x) -> printBoard(get()); x }
```

The function `printBoard` is impure as it prints an ASCII representation of the given game state to standard out. To make matters more interesting, we add replay functionality, which we implement by invoking a handler recursively on its input computation:

```
sig replay : (Int) -> (Comp({_|e}, a)) -> Comp({_|e}, a)
handler[m] replay(n)
{case Return(x) -> if (n <= 1) x else replay(n-1)(m)() }
```

The `replay` handler reevaluates the computation `m` precisely `n` times. Note, that the handler's effect signature is an empty, open row which means that it forwards every operation that might occur to subsequent handlers. Now we can wire everything together:

```
links> run -<- state(s0) -<- printer -<- replay(10) -<-
  coin -<- bobChooses -<- scoreUpdater -< game(7);
```

Figure 2 shows a possible output. In the same manner, we can effortlessly merge the cheating infrastructure into the pipeline without changing the underlying computation.

NIM HIGHSCORE	
Player	#Wins
Alice	7
Bob	3

The Nim scoreboard after 10 games with $n = 7$, where Alice plays the perfect strategy and Bob chooses between the perfect and cheating strategies.

Figure 2: Nim High Score Output

Value types	$A, B ::= A \rightarrow C \mid \forall \alpha^K. C$
Computation types	$C, D ::= A!E$
Effect types	$E ::= \{R\}$
Row types	$R ::= \ell : P; R \mid \rho \mid \cdot$
Presence types	$P ::= \text{Pre}(A) \mid \text{Abs} \mid \theta$
Handler types	$F ::= C \Rightarrow D$
Types	$T ::= A \mid C \mid E \mid R \mid P \mid F$
Kinds	$K ::= \text{Type} \mid \text{Row}_{\mathcal{L}} \mid \text{Presence}$ $\mid \text{Comp} \mid \text{Effect} \mid \text{Handler}$
Label sets	$\mathcal{L} ::= \emptyset \mid \{\ell\} \uplus \mathcal{L}$
Type environments	$\Gamma ::= \cdot \mid \Gamma, x : A$
Kind environments	$\Delta ::= \cdot \mid \Delta, \alpha : K$

Figure 3: Types, Effects, Kinds, and Environments

3. A Calculus of Handlers and Rows

In this section, we present a type and effect system and a small-step operational semantics for $\lambda_{\text{eff}}^{\rho}$ (pronounced “lambda-eff-row”), a Church-style row-polymorphic call-by-value calculus for effect handlers. This core calculus captures the essence of the Links IR. We prove that the operational semantics is sound with respect to the type and effect system.

A key advantage of row polymorphism is that it integrates rather smoothly with Hindley-Milner type inference. We concern ourselves only with the explicitly-typed core language, as the treatment of type inference is quite standard.

The design of $\lambda_{\text{eff}}^{\rho}$ is inspired by the λ -calculi of Kammar et al. [15], Pretnar [32], and Lindley and Cheney [22]. As in the work of Kammar et al. [15], each handler can have its own effect signature. As in the work of Pretnar [32], the underlying formalism is fine-grain call-by-value [20], which names each intermediate computation like in A-normal form [11], but unlike A-normal form is closed under β -reduction. As in the work of Lindley and Cheney [22], the effect system is based on row polymorphism.

3.1 Types

The grammars of types, kinds, label sets, and type and kind environments are given in Figure 3.

Value Types The function type $A \rightarrow C$ represents functions that map values of type A to computations of type C . The polymorphic type $\forall \alpha^K. C$ is parameterised by a type variable α of kind K . The record type $\langle R \rangle$ represents records with fields constrained by row R . Dually, the variant type $[R]$ represents tagged sums constrained by row R . The handler type $C \Rightarrow D$ represents handlers that transform computations of type C into computations of type D .

Values	$V, W ::= x \mid \lambda x^A. M \mid \Lambda \alpha^K. M$ $\mid \langle \rangle \mid \langle \ell = V; W \rangle \mid (\ell V)^R$
Computations	$M, N ::= V W \mid V A$ $\mid \text{let } \langle \ell = x; y \rangle = V \text{ in } N$ $\mid \text{case } V \{ \ell x \mapsto M; y \mapsto N \} \mid \text{absurd}^C V$ $\mid \text{return } V$ $\mid \text{let } x \leftarrow M \text{ in } N$ $\mid (\text{do } \ell V)^E$ $\mid \text{handle } M \text{ with } H$
Handlers	$H ::= \{\text{return } x \mapsto M\}$ $\mid \{ \ell x k \mapsto M \} \uplus H$

Figure 4: Term Syntax

Computation Types A computation type $A!E$ is given by a value type A and an effect E , which specifies the operations that the computation may perform.

Row Types Effect types, records and variants are defined in terms of rows. A row type embodies a collection of distinct labels, each of which is annotated with a presence type. A presence type indicates whether a label is *present* with some type A ($\text{Pre}(A)$), *absent* (Abs) or *polymorphic* in its presence (θ).

Row types are either *closed* or *open*. A closed row type ends in $\cdot$, whilst an open row type ends with a *row variable* ρ . Furthermore, a closed row term can have only the labels explicitly mentioned in its type. Conversely, the row variable in an open row can be instantiated with additional labels. We identify rows up to reordering of labels. For instance, we consider the following two rows equivalent:

$$\ell_1 : P_1; \dots; \ell_n : P_n \equiv \ell_n : P_n; \dots; \ell_1 : P_1.$$

The unit and empty type are definable in terms of row types. We define the unit type as the empty, closed record, that is, $\langle \cdot \rangle$. Similarly, we define the empty type as the empty, closed variant $[\cdot]$. Usually, we usually omit the $\cdot$ for closed rows.

Handler Types A handler type $C \Rightarrow D$ is given by an input computation type C and an output computation type D .

Kinds We have six kinds: Type , Comp , Effect , $\text{Row}_{\mathcal{L}}$, Presence , Handler , which classify value types, computation types, effect types, row types, presence types, and handler types, respectively. Row kinds are annotated with a set of labels $\mathcal{L}$. The kind of a complete row is $\text{Row}_{\emptyset}$. More generally, the kind $\text{Row}_{\mathcal{L}}$ denotes a partial row that cannot mention the labels in $\mathcal{L}$.

Type Variables We let α, ρ and θ range over type variables. By convention we use α for value type variables or for type variables of unspecified kind, ρ for type variables of row kind, and θ for type variables of presence kind.

Type and Kind Environments Type environments map term variables to their types and kind environments map type variables to their kinds.

3.2 Terms

The terms are given in Figure 4. We let x, y, z, k range over term variables. By convention, we use k to denote continuation names.

The syntax partitions terms into values, computations and handlers. Value terms comprise variables (x), lambda abstraction ($\lambda x^A. M$), type abstraction ($\Lambda \alpha^K. M$), and the introduction forms for records and variants. Records are introduced using the empty record $\langle \rangle$ and record extension $\langle \ell = V; W \rangle$, whilst variants are introduced using injection $(\ell V)^R$, which injects a field with label ℓ

and value V into a row whose type is R . We include the row type annotation in order to support bottom-up type reconstruction.

All elimination forms are computation terms. Abstraction and type abstraction are eliminated using application (VW) and type application (VA) respectively. The record eliminator ($\text{let } \langle \ell = x; y \rangle = V \text{ in } N$) splits a record V into x , the value associated with ℓ , and y , the rest of the record. Non-empty variants are eliminated using the case construct ($\text{case } V \{ \ell \ x \mapsto M; y \mapsto N \}$), which evaluates the computation M if the tag of V matches ℓ . Otherwise it falls through to y and evaluates N . The elimination form for empty variants is ($\text{absurd}^C V$). A trivial computation ($\text{return } V$) returns value V . The expression ($\text{let } x \leftarrow M \text{ in } N$) evaluates M and binds the result value to x in N .

The construct ($\text{do } \ell \ V \rangle^E$) invokes an operation ℓ with value argument V . The handle construct ($\text{handle } M \text{ with } H$) runs a computation M with handler definition H . A handler definition H consists of a return clause $\text{return } x \mapsto M$ and a possibly empty set of operation clauses $\{ \ell_i \ x_i \ k_i \mapsto M_i \}_i$. The return clause defines how to handle the final return value of the handled computation, which is bound to x in M . The i -th operation clause binds the operation parameter to x_i and the continuation k_i in M_i .

We write $\text{Id}(M)$ for $\text{handle } M \text{ with } \{\text{return } x \mapsto x\}$. We write $H(\text{return})$ for the return clause of H and $H(\ell)$ for the set of either zero or one operation clauses in H that handle the operation ℓ . We write $\text{dom}(H)$ for the set of operations handled by H . As our calculus is Church-style, we annotate various term forms with type or kind information (term abstraction, type abstraction, injection, operations, and empty cases); we sometimes omit these annotations.

3.3 Static Semantics

The kinding rules are given in Figure 5 and the typing rules are given in Figure 6.

The kinding judgement $\Delta \vdash T : K$ states that type T has kind K in kind environment Δ . The value typing judgement $\Delta; \Gamma \vdash V : A$ states that value term V has type A under kind environment Δ and type environment Γ . The computation typing judgement $\Delta; \Gamma \vdash M : C$ states that term M has computation type C under kind environment Δ and type environment Γ . The handler typing judgement $\Delta; \Gamma \vdash H : C \Rightarrow D$ states that handler H has type $C \Rightarrow D$ under kind environment Δ and type environment Γ . In the typing judgements, we implicitly assume that Γ , A , C , and D , are well-kinded with respect to Δ . We define the functions $\text{FTV}(\Gamma)$ to be the set of free type variables in Γ .

The kinding and typing rules are mostly straightforward. The interesting typing rules are T-HANDLE and the two handler rules. The T-HANDLE rule states that $\text{handle } M \text{ with } H$ produces a computation of type D given that the computation M has type C , and that H is a handler that transforms a computation of type C into another computation of type D .

The T-HANDLER rule is crucial. The effect rows on the computation type C and the output computation type D must share the same suffix R . This means that the effect row of D must explicitly mention each of the operations ℓ_i , whether that be to say that an ℓ_i is present with a given type signature, absent, or polymorphic in its presence. The row R describes the operations that are forwarded. It may include a row-variable, in which case an arbitrary number of effects may be forwarded by the handler. The typing of the return clause is straightforward. In the typing of each operation clause, the continuation returns the output computation type D . Thus, we are here defining *deep* handlers [15] in which the handler is implicitly wrapped around the continuation, such that any subsequent operations are handled uniformly by the same handler. The Links implementation also supports *shallow* handlers [15], in which the continuation is instead annotated with the input effect and one has

$\frac{\text{TYVAR}}{\Delta, \alpha : K \vdash \alpha : K}$	$\frac{\text{COMP} \quad \Delta \vdash A : \text{Type} \quad \Delta \vdash E : \text{Effect}}{\Delta \vdash A!E : \text{Comp}}$	
$\frac{\text{FUN} \quad \Delta \vdash A : \text{Type} \quad \Delta \vdash C : \text{Comp}}{\Delta \vdash A \rightarrow C : \text{Type}}$	$\frac{\text{FORALL} \quad \Delta, \alpha : K \vdash C : \text{Comp}}{\Delta \vdash \forall \alpha^K. C : \text{Type}}$	
$\frac{\text{RECORD} \quad \Delta \vdash R : \text{Row}_\emptyset}{\Delta \vdash \langle R \rangle : \text{Type}}$	$\frac{\text{VARIANT} \quad \Delta \vdash R : \text{Row}_\emptyset}{\Delta \vdash [R] : \text{Type}}$	$\frac{\text{EFFECT} \quad \Delta \vdash R : \text{Row}_\emptyset}{\Delta \vdash \{R\} : \text{Effect}}$
$\frac{\text{PRESENT} \quad \Delta \vdash A : \text{Type}}{\Delta \vdash \text{Pre}(A) : \text{Presence}}$	$\frac{\text{ABSENT}}{\Delta \vdash \text{Abs} : \text{Presence}}$	
$\frac{\text{EMPTYROW}}{\Delta \vdash \cdot : \text{Row}_\mathcal{L}}$	$\frac{\text{EXTENDROW} \quad \Delta \vdash P : \text{Presence} \quad \Delta \vdash R : \text{Row}_{\mathcal{L} \uplus \{\ell\}}}{\Delta \vdash \ell : P; R : \text{Row}_\mathcal{L}}$	
$\frac{\text{HANDLER} \quad \Delta \vdash C : \text{Comp} \quad \Delta \vdash D : \text{Comp}}{\Delta \vdash C \Rightarrow D : \text{Handler}}$		

Figure 5: Kinding Rules

to explicitly reinvoke the handler after applying the continuation inside an operation clause.

3.4 Operational Semantics

We give a small-step operational semantics for $\lambda_{\text{eff}}^\rho$. Figure 7 displays the operational rules. The reduction relation $\rightsquigarrow$ is defined on computation terms. The statement $M \rightsquigarrow M'$ reads: term M reduces to term M' in a single step. Most of the rules are standard. We use evaluation contexts to focus on the active expression. The interesting rules are the handler rules.

We write $\text{BL}(\mathcal{E})$ for the set of operation labels bound by $\mathcal{E}$.

$$\begin{aligned} \text{BL}([\]) &= \emptyset \\ \text{BL}(\text{let } x \leftarrow \mathcal{E} \text{ in } N) &= \text{BL}(\mathcal{E}) \\ \text{BL}(\text{handle } \mathcal{E} \text{ with } H) &= \text{BL}(\mathcal{E}) \cup \text{dom}(H) \end{aligned}$$

The rule S-HANDLE-RET invokes the return clause of a handler. The rule S-HANDLE-OP handles an operation by invoking the appropriate operation clause. The constraint $\ell \notin \text{BL}(\mathcal{E})$ ensures that no inner handler inside the evaluation context is able to handle the operation: thus a handler is able to reach past any other inner handlers that do not handle ℓ . In our abstract machine semantics we realise this behaviour using explicit forwarding operations, but more efficient implementations are perfectly feasible.

Remark The attentive reader may have noticed that the two rules S-LET and S-HANDLE-RET are strikingly similar. Indeed, we could omit let bindings altogether since any let binding

$$\text{let } x \leftarrow M \text{ in } N$$

may be simulated by a trivial handler:

$$\text{handle } M \text{ with } \{\text{return } x \mapsto N\}$$

Nevertheless, we elect to keep let bindings in $\lambda_{\text{eff}}^\rho$. As we discuss in Section 4, one reason for doing so is in order to facilitate more efficient implementations.

We write R^+ for the transitive closure of relation R . Subject reduction and type soundness for $\lambda_{\text{eff}}^\rho$ are standard.

Values

$$\begin{array}{c}
\text{T-VAR} \\
\frac{x : A \in \Gamma}{\Delta; \Gamma \vdash x : A} \\
\\
\text{T-LAM} \\
\frac{\Delta; \Gamma, x : A \vdash M : C}{\Delta; \Gamma \vdash \lambda x^A. M : A \rightarrow C} \\
\\
\text{T-POLYLAM} \\
\frac{\Delta, \alpha : K; \Gamma \vdash M : C \quad \alpha \notin FTV(\Gamma)}{\Delta; \Gamma \vdash \Lambda \alpha^K. M : \forall \alpha^K. C} \\
\\
\text{T-UNIT} \\
\frac{}{\Delta; \Gamma \vdash \langle \rangle : \langle \rangle} \\
\\
\text{T-EXTEND} \\
\frac{\Delta; \Gamma \vdash V : A \quad \Delta; \Gamma \vdash W : \langle \ell : \text{Abs}; R \rangle}{\Delta; \Gamma \vdash \langle \ell = V; W \rangle : \langle \ell : \text{Pre}(A); R \rangle} \\
\\
\text{T-INJECT} \\
\frac{\Delta; \Gamma \vdash V : A}{\Delta; \Gamma \vdash (\ell V)^R : [\ell : \text{Pre}(A); R]}
\end{array}$$

Computations

$$\begin{array}{c}
\text{T-APP} \\
\frac{\Delta; \Gamma \vdash V : A \rightarrow C \quad \Delta; \Gamma \vdash W : B}{\Delta; \Gamma \vdash VW : C} \\
\\
\text{T-POLYAPP} \\
\frac{\Delta; \Gamma \vdash V : \forall \alpha^K. C \quad \Delta \vdash A : K}{\Delta; \Gamma \vdash VA : C[A/\alpha]} \\
\\
\text{T-SPLIT} \\
\frac{\Delta; \Gamma \vdash V : \langle \ell : \text{Pre}(A); R \rangle \quad \Delta; \Gamma, x : A, y : \langle \ell : \text{Abs}; R \rangle \vdash N : C}{\Delta; \Gamma \vdash \text{let } \langle \ell = x; y \rangle = V \text{ in } N : C} \\
\\
\text{T-CASE} \\
\frac{\Delta; \Gamma \vdash V : [\ell : \text{Pre}(A); R] \quad \Delta; \Gamma, x : A \vdash M : C \quad \Delta; \Gamma, y : [\ell : \text{Abs}; R] \vdash N : C}{\Delta; \Gamma \vdash \text{case } V \{ \ell x \mapsto M; y \mapsto N \} : C} \\
\\
\text{T-ABSURD} \\
\frac{\Delta; \Gamma \vdash V : \square}{\Delta; \Gamma \vdash \text{absurd}^C V : C} \\
\\
\text{T-RETURN} \\
\frac{\Delta; \Gamma \vdash V : A}{\Delta; \Gamma \vdash \text{return } V : A!E} \\
\\
\text{T-LET} \\
\frac{\Delta; \Gamma \vdash M : A!E \quad \Delta; \Gamma, x : A \vdash N : C}{\Delta; \Gamma \vdash \text{let } x \leftarrow M \text{ in } N : C} \\
\\
\text{T-DO} \\
\frac{\Delta; \Gamma \vdash V : A \quad E = \{ \ell : A \rightarrow B; R \}}{\Delta; \Gamma \vdash (\text{do } \ell V)^E : B!E} \\
\\
\text{T-HANDLE} \\
\frac{\Delta; \Gamma \vdash M : C \quad \Delta; \Gamma \vdash H : C \Rightarrow D}{\Delta; \Gamma \vdash \text{handle } M \text{ with } H : D}
\end{array}$$

Handlers

$$\begin{array}{c}
\text{T-HANDLER} \\
\frac{C = A! \{ (\ell_i : A_i \rightarrow B_i); R \} \quad D = B! \{ (\ell_i : P_i); R \} \quad H = \{ \text{return } x \mapsto M \} \uplus \{ \ell_i y k \mapsto N_i \}_i}{[\Delta; \Gamma, y : A_i, k : B_i \rightarrow D \vdash N_i : D]_i \quad \Delta; \Gamma, x : A \vdash M : D} \\
\Delta; \Gamma \vdash H : C \Rightarrow D
\end{array}$$

Figure 6: Typing Rules

$$\begin{array}{c}
\text{S-APP} \quad (\lambda x^A. M)V \rightsquigarrow M[V/x] \\
\text{S-TYAPP} \quad (\Lambda \alpha^K. M)A \rightsquigarrow M[A/\alpha] \\
\text{S-SPLIT} \quad \text{let } \langle \ell = x; y \rangle = \langle \ell = V; W \rangle \text{ in } N \rightsquigarrow N[V/x, W/y] \\
\text{S-CASE}_1 \quad \text{case } (\ell V)^R \{ \ell x \mapsto M; y \mapsto N \} \rightsquigarrow M[V/x] \\
\text{S-CASE}_2 \quad \text{case } (\ell V)^R \{ \ell' x \mapsto M; y \mapsto N \} \rightsquigarrow N[(\ell V)^R/y], \quad \text{if } \ell \neq \ell' \\
\text{S-LET} \quad \text{let } x \leftarrow \text{return } V \text{ in } N \rightsquigarrow N[V/x] \\
\text{S-HANDLE-RET} \quad \text{handle } (\text{return } V) \text{ with } H \rightsquigarrow N[V/x], \quad \text{where } \{ \text{return } x \mapsto N \} \in H \\
\text{S-HANDLE-OP} \quad \text{handle } \mathcal{E}[\text{do } \ell V] \text{ with } H \rightsquigarrow N[V/x, \lambda y. \text{handle } \mathcal{E}[\text{return } y] \text{ with } H/k], \\
\quad \text{where } \ell \notin BL(\mathcal{E}) \text{ and } \{ \ell x k \mapsto M \} \in H \\
\quad \text{if } \mathcal{E}[M] \rightsquigarrow \mathcal{E}[N] \\
\\
\text{S-LIFT} \quad M \rightsquigarrow N, \quad \text{if } \mathcal{E}[M] \rightsquigarrow \mathcal{E}[N]
\end{array}$$

$$\text{Evaluation contexts } \mathcal{E} ::= [] \mid \text{let } x \leftarrow \mathcal{E} \text{ in } N \mid \text{handle } \mathcal{E} \text{ with } H$$

Figure 7: Small-Step Operational Semantics

Theorem 3.1 (Subject Reduction). *If $\Delta; \Gamma \vdash M : A!E$ and $M \rightsquigarrow M'$, then $\Delta; \Gamma \vdash M' : A!E$.*

There are two ways in which a computation can terminate. It can either successfully return a value, or it can get stuck on an unhandled operation.

Definition 3.2. *We say that computation term N is normal with respect to effect E , if N is either of the form $\text{return } V$, or $\mathcal{E}[\text{do } \ell W]$, where $\ell \in E$ and $\ell \notin BL(\mathcal{E})$.*

If N is normal with respect to the empty effect $\{\cdot\}$, then N has the form $\text{return } V$.

Theorem 3.3 (Type Soundness). *If $\vdash M : A!E$, then there exists $\vdash N : A!E$, such that $M \rightsquigarrow^+ N \not\rightsquigarrow$, and N is normal with respect to effect E .*

4. Abstract Machine Semantics

In this section we present an abstract machine semantics for $\lambda_{\text{eff}}^{\rho}$ that is closely related to the actual implementation of effect handlers in Links. We prove that the abstract machine simulates the operational semantics in the sense that each reduction in the small step semantics corresponds exactly to a finite sequence of one or more steps of the abstract machine.

The Links interpreter is based on a *CEK*-style abstract machine [10] and operates directly on ANF terms [11]. The standard

Configurations	$C ::= \langle M \mid \gamma \mid \kappa \rangle$ $\quad \mid \langle M \mid \gamma \mid \kappa \mid \kappa' \rangle_{\text{op}}$
Value environments	$\gamma ::= \emptyset \mid \gamma[x \mapsto v]$
Values	$v, w ::= (\gamma, \lambda x^A. M) \mid (\gamma, \Lambda \alpha^K. M)$ $\quad \mid \langle \rangle \mid \langle \ell = v; w \rangle \mid (\ell v)^R \mid \kappa^A$
Continuations	$\kappa ::= [] \mid \delta :: \kappa$
Continuation frames	$\delta ::= (\sigma, \chi)$
Pure continuations	$\sigma ::= [] \mid \phi :: \sigma$
Pure continuation frames	$\phi ::= (\gamma, x, N)$
Handler closures	$\chi ::= (\gamma, H)$

Figure 8: Abstract Machine Syntax

CEK machine operates on configurations which are triples of the form $\langle C \mid E \mid K \rangle$.

- The control C is the expression currently being evaluated.
- The environment E binds the free variables.
- The continuation K instructs the machine what to do once it is done evaluating the current term in the C component.

In order to accommodate handlers we generalise the CEK machine. The syntax of abstract machine states is given in Figure 8. Just like in the standard CEK machine, a standard configuration $C = \langle M \mid \gamma \mid \kappa \rangle$ of our abstract machine is a triple of a computation term M , an environment γ mapping free variables to values, and a continuation κ . However, our continuations differ from the standard machine. On the one hand, they are somewhat simplified, due to our strict separation between computations and values. On the other hand, they have considerably more structure in order to accommodate effects and handlers. In order to account for forwarding of unhandled operations, configurations occasionally gain an additional continuation argument.

Values consist of function closures, type function closures, records, variants, and captured continuations. A continuation κ consists of a stack of continuation frames $[\delta_1, \dots, \delta_n]$. We choose to annotate captured continuations with their input type in order to make the results of Section 4.1 easier to state. Intuitively, each continuation frame $\delta = (\sigma, \chi)$ represents the pure continuation σ , corresponding to a sequence of let bindings, inside a particular handler closure χ . A pure continuation is a stack of pure continuation frames. A pure continuation frame (γ, x, N) closes a let-binding **let** $x = []$ **in** N over environment γ . A handler closure (γ, H) closes a handler definition H over environment γ . We write $[]$ for an empty stack, $x :: s$ for the result of pushing x on top of stack s , and $s \mathbin{++} s'$ for the concatenation of stack s on top of s' . We use pattern matching to deconstruct stacks.

The abstract machine semantics is given in Figure 9. The transition function is given by $\longrightarrow$. This depends on an interpretation function $\llbracket - \rrbracket$ for values. The machine is initialised (M-INIT) by placing a term in a configuration alongside the empty environment and identity continuation κ_0 . The rules (M-APP), (M-TYAPP), (M-SPLIT), and (M-CASE) enact the elimination of values. Note that (M-APP) handles application of both closures and of captured continuations. The rules (M-LET) and (M-HANDLE) extend the current continuation with let bindings and handlers respectively. The rule (M-RETCONT) binds a returned value if there is a pure continuation in the current continuation frame. The rule (M-RETHANDLER) invokes the return clause of a handler if there is no pure continuation in the current continuation frame, but there is a handler. The rule (M-RETTOP) returns a final value if the continuation is empty. The rule (M-OP) switches to a special four place configuration in order to handle an operation. The fourth compo-

nent of the configuration is an auxiliary forwarding continuation, which keeps track of the continuation frames through which the operation has been forwarded. It is initialised to be empty. The rule (M-OP-HANDLE) uses the current handler to handle an operation if the label matches one of the operation clauses of the current handler. The captured continuation is assigned the forwarding continuation with the current continuation frame appended to the bottom of it. The rule (M-OP-FORWARD) appends the current continuation frame onto the bottom of the forwarding continuation. Notice that if the main continuation is empty then the machine gets stuck. This occurs when an operation is unhandled, and the forwarding continuation describes the succession of handlers that have failed to handle the operation along with any pure continuations that were encountered along the way.

Assuming the input is a well-typed closed computation term $\vdash M : A!E$, the machine will either return a value of type A , or it will get stuck failing to handle an operation appearing in E . We now make the correspondence between the operational semantics and the abstract machine more precise.

Remark Taking advantage of the observation that let bindings can be simulated by handlers, it is straightforward to define a variant of our abstract machine with flat continuations in which each continuation frame is just a handler closure without a pure continuation. In practice, we believe it is helpful to distinguish between pure and impure parts of the continuation. Apart from anything else, it supports more efficient implementations by reducing the number of continuation frames that a forwarded effect must bubble through.

4.1 Correctness

The (M-INIT) rule immediately provides a canonical way to map a computation term onto the abstract machine. A more interesting question is how to map an arbitrary configuration to a computation term. Figure 10 describes such a mapping $\llbracket - \rrbracket$ from configurations to terms via a collection of mutually recursive functions defined on configurations, continuations, computation terms, handler definitions, value terms, and values. We write $\text{dom}(\gamma)$ for the domain of γ , and $\gamma \setminus \{x_1, \dots, x_n\}$ for the restriction of environment γ to $\text{dom}(\gamma) \setminus \{x_1, \dots, x_n\}$.

The $\llbracket - \rrbracket$ function enables us to classify the abstract machine reduction rules according to how they relate to the operational semantics. The rules (M-INIT) and (M-RETTOP) are concerned only with initial input and final output, neither of which is a feature of the operational semantics, so we can ignore them. The rules (M-APPCONT), (M-LET), (M-HANDLE), (M-OP), and (M-OP-FORWARD) are administrative in the sense that $\llbracket - \rrbracket$ is invariant under these rules. This leaves the β -rules (M-APP), (M-TYAPP), (M-SPLIT), (M-CASE), (M-RETCONT), (M-RETHANDLER), and (M-OP-HANDLE). Each of these corresponds directly with performing a reduction in the operational semantics. We write $\longrightarrow_a$ for administrative steps, $\longrightarrow_\beta$ for β -steps, and $\Longrightarrow$ for a sequence of steps of the form $\longrightarrow_a^* \longrightarrow_\beta$.

The following lemma describes how we can simulate each reduction in the operational semantics by a sequence of administrative steps followed by one β -step in the abstract machine. The idea is to represent a computation term M by the equivalence class of configurations C such that $\llbracket C \rrbracket = \text{Id}(M)$. The Id wrapper captures the top-level identity handler.

Lemma 4.1. *If $M \rightsquigarrow N$, then for any C such that $\llbracket C \rrbracket = \text{Id}(M)$ there exists C' such that $C \Longrightarrow C'$ and $\llbracket C' \rrbracket = \text{Id}(N)$.*

Proof. By induction on the derivation of $M \rightsquigarrow N$. If $\llbracket C \rrbracket = \text{Id}(M)$, then the underlying structure of the term in the configuration C must be the same as M , as $\llbracket - \rrbracket$ is homomorphic on computation terms. Some value subterms of M may appear in the environment,

Identity continuation

$$\kappa_0 = [([], (\emptyset, \{\mathbf{return} \ x \mapsto x\}))]$$

Transition function

M-INIT	$M \longrightarrow \langle M \mid \emptyset \mid \kappa_0 \rangle$
M-APP	$\langle V \ W \mid \gamma \mid \kappa \rangle \longrightarrow \langle M \mid \gamma'[x \mapsto \llbracket W \rrbracket \gamma] \mid \kappa \rangle, \quad \text{if } \llbracket V \rrbracket \gamma = (\gamma', \lambda x^A. M)$
M-APPCONT	$\langle V \ W \mid \gamma \mid \kappa \rangle \longrightarrow \langle \mathbf{return} \ W \mid \gamma \mid \kappa' ++ \kappa \rangle, \quad \text{if } \llbracket V \rrbracket \gamma = (\kappa')^A$
M-TYAPP	$\langle M \ A \mid \gamma \mid \kappa \rangle \longrightarrow \langle M[A/\alpha] \mid \gamma' \mid \kappa \rangle, \quad \text{if } \llbracket V \rrbracket \gamma = (\gamma', \Lambda \alpha^K. M)$
M-SPLIT	$\langle \mathbf{let} \ \langle \ell = x; y \rangle = V \ \mathbf{in} \ N \mid \gamma \mid \kappa \rangle \longrightarrow \langle N \mid \gamma[x \mapsto v, y \mapsto w] \mid \kappa \rangle, \quad \text{if } \llbracket V \rrbracket \gamma = \langle \ell = v; w \rangle$
M-CASE	$\langle \mathbf{case} \ V \ \{ \ell \ x \mapsto M; y \mapsto N \} \mid \gamma \mid \kappa \rangle \longrightarrow \begin{cases} \langle M \mid \gamma[x \mapsto v] \mid \kappa \rangle, & \text{if } \llbracket V \rrbracket \gamma = \ell \ v \\ \langle N \mid \gamma[y \mapsto \ell' v] \mid \kappa \rangle, & \text{if } \llbracket V \rrbracket \gamma = \ell' v \text{ and } \ell \neq \ell' \end{cases}$
M-LET	$\langle \mathbf{let} \ x \leftarrow M \ \mathbf{in} \ N \mid \gamma \mid (\sigma, \chi) :: \kappa \rangle \longrightarrow \langle M \mid \gamma \mid ((\gamma, x, N) :: \sigma, \chi) :: \kappa \rangle$
M-HANDLE	$\langle \mathbf{handle} \ M \ \mathbf{with} \ H \mid \gamma \mid \kappa \rangle \longrightarrow \langle M \mid \gamma \mid ([], (\gamma, H)) :: \kappa \rangle$
M-RETCONT	$\langle \mathbf{return} \ V \mid \gamma \mid ((\gamma', x, N) :: \sigma, \chi) :: \kappa \rangle \longrightarrow \langle N \mid \gamma'[x \mapsto \llbracket V \rrbracket \gamma] \mid (\sigma, \chi) :: \kappa \rangle$
M-RETHANDLER	$\langle \mathbf{return} \ V \mid \gamma \mid ([], (\gamma', H)) :: \kappa \rangle \longrightarrow \langle M \mid \gamma'[x \mapsto \llbracket V \rrbracket \gamma] \mid \kappa \rangle, \quad \text{if } H(\mathbf{return}) = \{\mathbf{return} \ x \mapsto M\}$
M-RETOP	$\langle \mathbf{return} \ V \mid \gamma \mid [] \rangle \longrightarrow \llbracket V \rrbracket \gamma$
M-OP	$\langle (\mathbf{do} \ \ell \ V)^E \mid \gamma \mid \kappa \rangle \longrightarrow \langle (\mathbf{do} \ \ell \ V)^E \mid \gamma \mid \kappa \mid [] \rangle_{\text{op}}$
M-OP-HANDLE	$\langle (\mathbf{do} \ \ell \ V)^E \mid \gamma \mid (\sigma, (\gamma', H)) :: \kappa \mid \kappa' \rangle_{\text{op}} \longrightarrow \langle M \mid \gamma'[x \mapsto \llbracket V \rrbracket \gamma, k \mapsto (\kappa' ++ [(\sigma, (\gamma', H))])^B] \mid \kappa \rangle, \quad \text{if } \ell : A \rightarrow B \in E \text{ and } H(\ell) = \{\ell \ x \ k \mapsto M\}$
M-OP-FORWARD	$\langle (\mathbf{do} \ \ell \ V)^E \mid \gamma \mid (\sigma, (\gamma', H)) :: \kappa \mid \kappa' \rangle_{\text{op}} \longrightarrow \langle (\mathbf{do} \ \ell \ V)^E \mid \gamma \mid \kappa \mid \kappa' ++ [(\sigma, (\gamma', H))] \rangle_{\text{op}}, \quad \text{if } H(\ell) = \emptyset$

Value interpretation

$$\begin{aligned} \llbracket x \rrbracket \gamma &= \gamma(x) & \llbracket \lambda x^A. M \rrbracket \gamma &= (\gamma, \lambda x^A. M) & \llbracket \Lambda \alpha^K. M \rrbracket \gamma &= (\gamma, \Lambda \alpha^K. M) \\ \llbracket \langle \rangle \rrbracket \gamma &= \langle \rangle & \llbracket \langle \ell = V; W \rangle \rrbracket \gamma &= \langle \ell = \llbracket V \rrbracket \gamma; \llbracket W \rrbracket \gamma \rangle & \llbracket \langle \ell \ V \rangle^R \rrbracket \gamma &= (\ell \ \llbracket V \rrbracket \gamma)^R \end{aligned}$$

Figure 9: Abstract Machine Semantics

Configurations

$$\langle \langle M \mid \gamma \mid \kappa \rangle \rangle = \langle \kappa \rangle (\langle M \rangle \gamma) \quad \langle \langle M \mid \gamma \mid \kappa \mid \kappa' \rangle_{\text{op}} \rangle = \langle \kappa' ++ \kappa \rangle (\langle M \rangle \gamma) = \langle \kappa' \rangle (\langle \kappa \rangle (\langle M \rangle \gamma))$$

Continuations

$$\begin{aligned} \langle [] \rangle M &= M \\ \langle ((\gamma, x, N) :: \sigma, \chi) :: \kappa \rangle M &= \langle (\sigma, \chi) :: \kappa \rangle (\mathbf{let} \ x \leftarrow M \ \mathbf{in} \ \langle N \rangle (\gamma \setminus \{x\})) \\ \langle ([], (\gamma, H)) :: \kappa \rangle M &= \langle \kappa \rangle (\mathbf{handle} \ M \ \mathbf{with} \ \langle H \rangle \gamma) \end{aligned}$$

Computation terms

$$\begin{aligned} \langle V \ W \rangle \gamma &= \langle V \rangle \gamma \ \langle W \rangle \gamma \\ \langle V \ A \rangle \gamma &= \langle V \rangle \gamma \ A \\ \langle \mathbf{let} \ \langle \ell = x; y \rangle = V \ \mathbf{in} \ N \rangle \gamma &= \mathbf{let} \ \langle \ell = x; y \rangle = \langle V \rangle \gamma \ \mathbf{in} \ \langle N \rangle (\gamma \setminus \{x, y\}) \\ \langle \mathbf{case} \ V \ \{ \ell \ x \mapsto M; y \mapsto N \} \rangle \gamma &= \mathbf{case} \ \langle V \rangle \gamma \ \{ \ell \ x \mapsto \langle M \rangle (\gamma \setminus \{x\}); y \mapsto \langle N \rangle (\gamma \setminus \{y\}) \} \\ \langle \mathbf{return} \ V \rangle \gamma &= \mathbf{return} \ \langle V \rangle \gamma \\ \langle \mathbf{let} \ x \leftarrow M \ \mathbf{in} \ N \rangle \gamma &= \mathbf{let} \ x \leftarrow \langle M \rangle \gamma \ \mathbf{in} \ \langle N \rangle (\gamma \setminus \{x\}) \\ \langle \mathbf{do} \ \ell \ V \rangle \gamma &= \mathbf{do} \ \ell \ \langle V \rangle \gamma \\ \langle \mathbf{handle} \ M \ \mathbf{with} \ H \rangle \gamma &= \mathbf{handle} \ \langle M \rangle \gamma \ \mathbf{with} \ \langle H \rangle \gamma \end{aligned}$$

Handler definitions

$$\begin{aligned} \langle \{\mathbf{return} \ x \mapsto M\} \rangle \gamma &= \{\mathbf{return} \ x \mapsto \langle M \rangle (\gamma \setminus \{x\})\} \\ \langle \{\ell \ x \ k \mapsto M\} \uplus H \rangle \gamma &= \{\ell \ x \ k \mapsto \langle M \rangle (\gamma \setminus \{x, k\})\} \uplus \langle H \rangle \gamma \end{aligned}$$

Value terms and values

$$\begin{aligned} \langle x \rangle \gamma &= \langle v \rangle, \quad \text{if } \gamma(x) = v \\ \langle x \rangle \gamma &= x, \quad \text{if } x \notin \text{dom}(\gamma) \\ \langle \lambda x^A. M \rangle \gamma &= \lambda x^A. \langle M \rangle (\gamma \setminus \{x\}) & \langle (\gamma, \lambda x^A. M) \rangle &= \lambda x^A. \langle M \rangle (\gamma \setminus \{x\}) \\ \langle \Lambda \alpha^K. M \rangle \gamma &= \Lambda \alpha^K. \langle M \rangle \gamma & \langle (\gamma, \Lambda \alpha^K. M) \rangle &= \Lambda \alpha^K. \langle M \rangle \gamma \\ \langle \langle \rangle \rangle \gamma &= \langle \rangle & \langle \langle \rangle \rangle &= \langle \rangle \\ \langle \langle \ell = V; W \rangle \rangle \gamma &= \langle \ell = \langle V \rangle \gamma; \langle W \rangle \gamma \rangle & \langle \langle \ell = v; w \rangle \rangle &= \langle \ell = \langle v \rangle; \langle w \rangle \rangle \\ \langle \langle \ell \ V \rangle^R \rangle \gamma &= (\ell \ \langle V \rangle \gamma)^R & \langle \langle \ell \ v \rangle^R \rangle &= (\ell \ \langle v \rangle)^R \\ & & \langle \kappa^A \rangle &= \lambda x^A. \langle \kappa \rangle (\mathbf{return} \ x) \end{aligned}$$

Figure 10: Mapping from Abstract Machine Configurations to Terms

and part of the evaluation context of M may appear in the continuation. Administrative reductions update C by growing the continuation, whilst maintaining the invariant that $\langle C \rangle = Id(M)$. This process corresponds directly to traversing an evaluation context. It is straightforward to see that only a finite number of administrative reductions can occur consecutively as they either reduce the size of M or leave it unchanged and reduce the size of κ in the case of forwarding. Eventually the control part of the continuation will contain a redex corresponding to the active redex in M and it will then transition via a β -rule to a configuration C' , such that $\langle C' \rangle = Id(N)$. $\square$

The correspondence here is rather strong: there is a one-to-one mapping between $\rightsquigarrow$ and $\implies$. The inverse of the lemma is straightforward as the semantics is deterministic. Notice that Lemma 4.1 does not require that M be well-typed. We have chosen here not to perform type-erasure, but it is straightforward to adapt the results to semantics in which all polymorphism is erased and all type annotations are erased.

Theorem 4.2 (Simulation). *If $M : A!E$, and $M \rightsquigarrow^+ N$, such that N is normal with respect to E , then $M \longrightarrow^+ C$, such that $\langle C \rangle = N$.*

Proof. By repeated application of Lemma 4.1. $\square$

5. Implementation

Our implementation of handlers is based on a mild syntactic extension to Links: the syntax is extended with `do` for invoking operations and `handle(m) { ... }` for handling abstract computations.

Syntactic Sugar We provide syntactic sugar to make it more convenient to program with handlers. The function $\mathcal{D}$ is a source-to-source translation that elaborates the sugar. Figure 11 shows the cases for handlers only: $\mathcal{D}$ is a homomorphism on the other syntax constructors. Crucially, handlers are desugared into a function, that returns a thunk. This is important to ensure that handlers compose smoothly. For the same reason a parameterised handler desugars into a curried function, where the parameters precede the computation argument m . The parameters are passed around by enclosing each operation clause by a function. Thus, the initial parameter values are applied directly to the `handle` expression.

Backend The Links interpreter is based on a CEK machine for ANF expressions. We have generalised this machine to support handlers based on the abstract machine of Section 4. The interpreter maintains a stack of handlers with first-in-last-out semantics, which makes it straightforward to implement effect forwarding. The invocation of an operation causes the interpreter to unwind the stack to find a suitable handler for the operation.

Row Polymorphism We have extended Links with support for user-defined operations, making use of the existing row type system. The current row type system is based on that of Rémy [33], adapted to support effect typing in a similar manner to the work of Leroy and Pessaux [19] and Blume et al. [4] on typing exceptions. Fields in a record can be absent, present at a particular type, or polymorphic in their presence. An earlier version of Links [22] was based on a slightly more refined variant of Rémy’s system, $\Pi ML'$ [33], in which the type of a label is independent of its presence. This system was abandoned because it seemed somewhat counterintuitive for the purposes of record typing, but (as discussed in Section 2.5) it may offer advantages for effect typing.

Subtyping and Row Typing

Subtyping is in fact a poor man’s row polymorphism.

— Andreas Rossberg¹

Subtyping (or subeffecting) and row typing address similar concerns. However, they are not the same. In one dimension row polymorphism is more expressive than subtyping and in another dimension subtyping is more expressive than row polymorphism. Row polymorphism allows part of an effect to be named and reused in several places, which is essential for typing polymorphic functions such as `map`. Row polymorphism can also be used whenever one might otherwise use subtyping in a first order manner. In terms of effects, this amounts to always keeping functions polymorphic in their effects such that the effect variable can be instantiated in order to simulate an upcast. On the other hand, subtyping applies at higher-order whereas row typing does not.

Links does not support subsumption (implicitly inferred subtyping), but does support subtyping via explicit upcasts. Anecdotally, such casts seem to be rarely needed in practice and can often be replaced by first-class polymorphism (another Links feature) instead.

Shallow Handlers We have not covered them in much detail here, but our implementation also supports shallow handlers [15]. These are indicated by using the `shallowhandler` keyword in place of `handler`. Whereas a deep handler performs a fold over a computation, a shallow handler merely performs a case-split. This means that one must explicitly reinvoke the handler each time the continuation is implied inside an operation clause. An advantage is that it makes it easy to switch to a different handler midway through a computation. A disadvantage is that shallow handlers are not much use without an external notion of recursion, and they are less easy to optimise than deep handlers [35]. The changes to the typing rules and operational semantics to accommodate shallow handlers are standard [15]. The modifications to the abstract machine are modest; the M-OP-HANDLE rule is adapted to drop the current handler and append the pure continuation on to its successor:

$$\langle (\text{do } \ell \ V) \mid \gamma \mid (\sigma, \chi) :: (\sigma', \chi') :: \kappa \mid \kappa' \rangle_{\text{op}} \longrightarrow \langle M \mid \gamma'[x \mapsto \llbracket V \rrbracket \gamma, k \mapsto (\kappa')^B] \mid (\sigma \uparrow \sigma', \chi') :: \kappa \rangle$$

6. Related Work

Faking Row Polymorphism in Haskell Haskell provides a rather rich type system that allows one to simulate many aspects of row polymorphism. Perhaps the biggest mismatch between rows and other typing features is that rows are inherently unordered, whereas other typing features are usually inherently ordered.

One approach is Swierstra’s *data types a la carte* technique [34]. This amounts to encoding a row type as a sum, and then leveraging the type class system to automatically navigate through the sum type as if it were unordered. In practice, this encoding is a little fragile (e.g. sometimes additional type annotations are required), although recent improvements can make it somewhat more robust [25], particularly if one adds support for instance chains [26].

Another approach is to take advantage of the property that type class constraints are unordered. Early work on monad transformers [21] uses this idea to write modular abstract computations, as do Kammar et al. [15], Kiselyov et al. [17], and Kiselyov and Ishii [16] in their handler libraries. However, without some form of higher-order constraint solving (not supported by Haskell), one must still materialise ordered lists of effects when composing handlers. For many useful examples this is not a problem, but if we wish to build a list of handlers from disparate sources, then we must carefully ensure that their effects are composed in the same order.

Orchard and Petricek [28] make some progress towards encoding unordered effect rows by performing a sorting algorithm at the level of types, taking advantage of GHC’s support for dependently-typed programming [37]. However, this approach can fail in practice, as the type system cannot always infer that two types are equivalent in the presence of effect polymorphism.

¹<http://lambda-the-ultimate.org/node/3711#comment-52984>

Handler

$\text{handler } h(\overline{p}) \equiv \text{handler } [m] \ h(\overline{p}), \text{ where } m \text{ is fresh.}$

$$\mathcal{D} \left(\text{handler } [m] \ h(\overline{p}) \ \{ \bar{c} \} \right) = \begin{cases} \text{fun } h(m) () \ \{ \text{handle}(m) \ \{ \mathcal{D}(\bar{c}) \} \} & \text{if } |\overline{p}| = 0 \\ \text{fun } h(\overline{p})(m) () \ \{ \text{handle}(m) \ \{ \mathcal{D}_{(\overline{p})}(\bar{c}) \}(\overline{p}) \} & \text{otherwise} \end{cases}$$

Handler cases

$$\mathcal{D}_{(\overline{p})}(\bar{c}) = \mathcal{D}_{(\overline{p})}(c_1) \cdots \mathcal{D}_{(\overline{p})}(c_n) \quad \mathcal{D}_{(\overline{p})}(\text{case } q \rightarrow M) = \begin{cases} \text{case } q \rightarrow \mathcal{D}(M) & \text{if } |\overline{p}| = 0 \\ \text{case } q \rightarrow \text{fun}(\overline{p}) \ \{ \mathcal{D}(M) \} & \text{otherwise} \end{cases}$$

Figure 11: Desugaring Handlers

Implementations Any signature of abstract operations can be understood as a free algebra and represented as a functor. In particular, every such functor gives rise to a free monad. Thus, free monads provide a natural basis for implementing effect handlers. Many of the library implementations of effect handlers include implementations based on free monads [6, 15–17, 36].

Kammar et al. [15] provide an implementation of effect handlers using a continuation monad, which avoids materialising any data constructors. Wu and Schrijvers [35] explain how it works, by taking advantage of Haskell’s fusion optimisations. This approach does appear to depend rather critically on the handlers being deep rather than shallow, and in Haskell it relies on them being type classes, and hence not really first class.

The Idris effects library [6] takes advantage of dependent types to provide effect handlers for a form of effects corresponding to parameterised monads [1].

We are aware of three languages that are specifically designed with effect handlers in mind.

- Eff [3] is a strict language with Hindley-Milner type inference similar in spirit to ML, but extended with effect handlers. It includes a novel feature for supporting fresh generation of effects in order to support effects such as ML-style higher-order state (which has an operation for generating new references). The original version of Eff [3] does not include an effect type system. Later variants do include an effect type system [2, 31] which is considerably more complicated than ours. It makes essential use of subtyping, includes a region system, and a form of effect polymorphism which uses a form of row polymorphism.
- Frank [23] is a language with effect handlers but no separate notion of function: a function is but a special case of a handler. Frank has a bidirectional type system. It includes an effect type system and a novel form of effect polymorphism in which the programmer need never read or write effect variables. Frank’s effect system can be viewed as implementing a form of row polymorphism. Unlike Links, but much like Koka [18], Frank allows multiple occurrences of the same label in a row. In contrast rows in Links are based on Remy’s design in which duplicates are not allowed, but negative information is.
- Shonky [24] amounts to a dynamically-typed variant of Frank. Though it is not statically typed, handlers must be annotated with the names of the effects that they handle. The implementation of Shonky is quite similar to ours in that it uses a generalisation of the CEK machine. The main differences are that Shonky does not use an ANF representation, so has more forms of continuation to handle, and in contrast to our nested continuation structure, Shonky uses a completely flat structure.

Although OCaml itself has no support for effect handlers, a development branch, Multicore OCaml [9], does. Multicore OCaml does not include an effect type system, and handlers are restricted

so that continuations can be invoked at most once. This design admits an efficient implementation, as continuations need never be copied, so they can simply be stored on the stack.

7. Conclusions and Future Work

We have implemented algebraic effects and handlers using row polymorphism and demonstrated that rows allow us to compose effectful computations seamlessly. We have formalised our system as the core calculus $\lambda_{\text{eff}}^\rho$ and shown a correspondence between two semantics: a small-step operational semantics and an abstract machine semantics, the latter of which matches the implementation. We conclude by discussing ongoing and future work.

Effects are pervasive in modern web applications, so we would like to extend our implementation to the client backend of Links. The client backend already produces JavaScript in continuation-passing style in order to implement concurrency. We plan to extend this representation to support handlers.

The overhead incurred by the Links interpreter is significant [12, 14]. To improve performance, we are working on building a compiler backend with support for handlers [13]. One performance bottleneck we envisage is the need to support copying of continuations. It is well-known that one-shot continuations can be implemented efficiently [7]. Links already has a linear type system, which in the future we plan to take advantage of in order to track the linearity of handlers. During code generation we can specialise the run-time representation of handlers according to their linearity.

Links employs a message-passing concurrency model, similar to Erlang, but typed. Taking ideas from Multicore OCaml [9], we are investigating whether we can rebuild the Links concurrency implementation directly in terms of handlers.

Acknowledgments

Nicolas Oury originally suggested the Nim game as an example to demonstrate programming with handlers. Thanks to Simon Fowler, Ohad Kammar, Caoimhín Laoide-Kemp, Craig McLaughlin, Mathias R. Pedersen, and Gabriel Scherer for useful feedback, suggestions, and discussions. The first author was supported by EPSRC grant EP/L01503X/1 (University of Edinburgh CDT in Pervasive Parallelism). The second author was supported by EPSRC grant EP/K034413/1 (A Basis for Concurrency and Distribution).

References

- [1] R. Atkey. Parameterised notions of computation. *Journal of Functional Programming*, 19(3-4):335–376, 2009.
- [2] A. Bauer and M. Pretnar. An effect system for algebraic effects and handlers. *Logical Methods in Computer Science*, 10(4), 2014.
- [3] A. Bauer and M. Pretnar. Programming with algebraic effects and handlers. *Journal of Logical and Algebraic Methods in Programming*, 84(1):108–123, 2015.

- [4] M. Blume, U. A. Acar, and W. Chae. Exception handlers as extensible cases. In G. Ramalingam, editor, *Programming Languages and Systems, 6th Asian Symposium, APLAS 2008, Bangalore, India, December 9-11, 2008. Proceedings*, volume 5356 of *Lecture Notes in Computer Science*, pages 273–289. Springer, 2008.
- [5] C. L. Bouton. Nim, a game with a complete mathematical theory. *The Annals of Mathematics*, 3(1/4):35–39, 1901. URL <https://www.jstor.org/stable/pdf/1967631.pdf>.
- [6] E. Brady. Programming and reasoning with algebraic effects and dependent types. In Morrisett and Uustalu [27], pages 133–144.
- [7] C. Bruggeman, O. Waddell, and R. K. Dybvig. Representing control in the presence of one-shot continuations. In C. N. Fischer, editor, *Proceedings of the ACM SIGPLAN’96 Conference on Programming Language Design and Implementation (PLDI)*, Philadelphia, Pennsylvania, May 21–24, 1996, pages 99–107. ACM, 1996.
- [8] E. Cooper, S. Lindley, P. Wadler, and J. Yallop. Links: Web programming without tiers. In F. S. de Boer, M. M. Bonsangue, S. Graf, and W. P. de Roever, editors, *Formal Methods for Components and Objects, 5th International Symposium, FMCO 2006, Amsterdam, The Netherlands, November 7–10, 2006, Revised Lectures*, volume 4709 of *Lecture Notes in Computer Science*, pages 266–296. Springer, 2006.
- [9] S. Dolan, L. White, K. Sivaramakrishnan, J. Yallop, and A. Madhavapeddy. Effective concurrency through algebraic effects. OCaml Workshop, September 2015. URL http://kcsrk.info/papers/effects_ocaml15.pdf.
- [10] M. Felleisen and D. P. Friedman. Control operators, the SECD-machine, and the λ -calculus. In *The Proceedings of the Conference on Formal Description of Programming Concepts III, Ebberup, Denmark*, pages 193–217. Elsevier, 1987.
- [11] C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In *Proceedings of the ACM SIGPLAN 1993 Conference on Programming Language Design and Implementation*, PLDI ’93, pages 237–247, New York, NY, USA, 1993. ACM.
- [12] D. Hillerström. Handlers for algebraic effects in Links. Master’s thesis, The University of Edinburgh, Scotland, 2015. URL http://project-archive.inf.ed.ac.uk/msc/20150206/msc_proj.pdf.
- [13] D. Hillerström, S. Lindley, and K. Sivaramakrishnan. Compiling Links effect handlers to the OCaml backend. ML Workshop, 2016. URL <http://homepages.inf.ed.ac.uk/s1467124/papers/mlabstract2016.pdf>.
- [14] S. Holmes. Compiling Links server-side code. Bachelor thesis, The University of Edinburgh, 2009. URL <http://links-lang.org/papers/undergrads/steven.pdf>.
- [15] O. Kammar, S. Lindley, and N. Oury. Handlers in action. In Morrisett and Uustalu [27], pages 145–158.
- [16] O. Kiselyov and H. Ishii. Freer monads, more extensible effects. In B. Lippmeier, editor, *Proceedings of the 8th ACM SIGPLAN Symposium on Haskell, Haskell 2015, Vancouver, BC, Canada, September 3–4, 2015*, pages 94–105. ACM, 2015.
- [17] O. Kiselyov, A. Sabry, and C. Swords. Extensible effects: an alternative to monad transformers. In C. Shan, editor, *Proceedings of the 2013 ACM SIGPLAN Symposium on Haskell, Boston, MA, USA, September 23–24, 2013*, pages 59–70. ACM, 2013.
- [18] D. Leijen. Koka: Programming with row polymorphic effect types. In P. Levy and N. Krishnaswami, editors, *Proceedings 5th Workshop on Mathematically Structured Functional Programming, MSFP, Grenoble, France, April 2014*, volume 153 of *EPTCS*, pages 100–126, 2014.
- [19] X. Leroy and F. Pessaux. Type-based analysis of uncaught exceptions. *ACM Transactions on Programming Languages and Systems*, 22(2): 340–377, 2000.
- [20] P. B. Levy, J. Power, and H. Thielecke. Modelling environments in call-by-value programming languages. *Information and Computation*, 185(2):182–210, 2003.
- [21] S. Liang, P. Hudak, and M. P. Jones. Monad transformers and modular interpreters. In R. K. Cytron and P. Lee, editors, *Conference Record of POPL’95: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Francisco, California, USA, January 23–25, 1995*, pages 333–343. ACM Press, 1995.
- [22] S. Lindley and J. Cheney. Row-based effect types for database integration. In B. C. Pierce, editor, *Proceedings of TLDI 2012: The Seventh ACM SIGPLAN Workshop on Types in Languages Design and Implementation, Philadelphia, PA, USA, Saturday, January 28, 2012*, pages 91–102. ACM, 2012.
- [23] S. Lindley, C. McBride, and C. McLaughlin. Do be do be do. URL <http://homepages.inf.ed.ac.uk/slindley/papers/frankly-draft-july2016.pdf>. Draft, March 2016.
- [24] C. McBride. Shonky, 2016. <https://github.com/pigworker/shonky>.
- [25] J. G. Morris. Variations on variants. In B. Lippmeier, editor, *Proceedings of the 8th ACM SIGPLAN Symposium on Haskell, Haskell 2015, Vancouver, BC, Canada, 2015*, pages 71–81. ACM, 2015.
- [26] J. G. Morris and M. P. Jones. Instance chains: type class programming without overlapping instances. In P. Hudak and S. Weirich, editors, *Proceeding of the 15th ACM SIGPLAN International Conference on Functional Programming, ICFP 2010, Baltimore, Maryland, USA, September 27–29, 2010*, pages 375–386. ACM, 2010.
- [27] G. Morrisett and T. Uustalu, editors. *ACM SIGPLAN International Conference on Functional Programming, ICFP’13, Boston, MA, USA - September 25 - 27, 2013*, 2013. ACM.
- [28] D. A. Orchard and T. Petricek. Embedding effect systems in Haskell. In W. Swierstra, editor, *Proceedings of the 2014 ACM SIGPLAN symposium on Haskell, Gothenburg, Sweden, September 4–5, 2014*, pages 13–24. ACM, 2014.
- [29] G. D. Plotkin and J. Power. Adequacy for algebraic effects. In F. Honsell and M. Miculan, editors, *Foundations of Software Science and Computation Structures, 4th International Conference, FOSSACS Genova, Italy, April 2–6, 2001, Proceedings*, volume 2030 of *Lecture Notes in Computer Science*, pages 1–24. Springer, 2001.
- [30] G. D. Plotkin and M. Pretnar. Handling algebraic effects. *Logical Methods in Computer Science*, 9(4), 2013.
- [31] M. Pretnar. Inferring algebraic effects. *Logical Methods in Computer Science*, 10(3), 2014.
- [32] M. Pretnar. An introduction to algebraic effects and handlers. invited tutorial paper. *Electronic Notes in Theoretical Computer Science*, 319: 19–35, 2015.
- [33] D. Rémy. Theoretical aspects of object-oriented programming. chapter Type Inference for Records in Natural Extension of ML, pages 67–95. MIT Press, Cambridge, MA, USA, 1994.
- [34] W. Swierstra. Data types à la carte. *Journal of Functional Programming*, 18(4):423–436, 2008.
- [35] N. Wu and T. Schrijvers. Fusion for free - efficient algebraic effect handlers. In R. Hinze and J. Voigtländer, editors, *Mathematics of Program Construction - 12th International Conference, MPC 2015, Königswinter, Germany, 2015. Proceedings*, volume 9129 of *Lecture Notes in Computer Science*, pages 302–322. Springer, 2015.
- [36] N. Wu, T. Schrijvers, and R. Hinze. Effect handlers in scope. In *Proceedings of the 2014 ACM SIGPLAN Symposium on Haskell, Haskell ’14*, pages 1–12, New York, NY, USA, 2014. ACM.
- [37] B. A. Yorgey, S. Weirich, J. Cretin, S. L. P. Jones, D. Vytiniotis, and J. P. Magalhães. Giving Haskell a promotion. In B. C. Pierce, editor, *Proceedings of TLDI 2012: The Seventh ACM SIGPLAN Workshop on Types in Languages Design and Implementation, Philadelphia, PA, USA, Saturday, January 28, 2012*, pages 53–66. ACM, 2012.