THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Shallow Effect Handlers

Citation for published version:

Hillerström, D & Lindley, S 2018, Shallow Effect Handlers. in *Proceedings of 16th Asian Symposium on Programming Languages and Systems (APLAS) 2018*. Lecture Notes in Computer Science, vol. 11275, Programming and Software Engineering, vol. 11275, Springer, Cham, Wellington, New Zealand, pp. 415-435, 16th Asian Symposium on Programming Languages and Systems, Wellington, New Zealand, 2/12/18. https://doi.org/10.1007/978-3-030-02768-1_22

Digital Object Identifier (DOI):

[10.1007/978-3-030-02768-1_22](https://doi.org/10.1007/978-3-030-02768-1_22)

Link:

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

Document Version:

Peer reviewed version

Published In:

Proceedings of 16th Asian Symposium on Programming Languages and Systems (APLAS) 2018

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.

Shallow Effect Handlers

Daniel Hillerström and Sam Lindley(✉)

The University of Edinburgh, United Kingdom
{Daniel.Hillerstrom,Sam.Lindley}@ed.ac.uk

Abstract. Plotkin and Pretnar’s effect handlers offer a versatile abstraction for modular programming with user-defined effects. Traditional *deep handlers* are defined by folds over computation trees. In this paper we study *shallow handlers*, defined instead by case splits over computation trees. We show that deep and shallow handlers can simulate one another up to specific notions of administrative reduction. We present the first formal accounts of an abstract machine for shallow handlers and a Continuation Passing Style (CPS) translation for shallow handlers taking special care to avoid memory leaks. We provide implementations in the Links web programming language and empirically verify that neither implementation introduces unwarranted memory leaks.

Keywords: effect handlers · abstract machines · continuation passing

1 Introduction

Expressive control abstractions are pervasive in mainstream programming languages, be that *async/await* as pioneered by C#, generators and iterators as commonly found in JavaScript and Python, or coroutines in C++20. Such abstractions may be simulated directly with higher-order functions, but at the expense of writing the entire source program in Continuation Passing Style (CPS). To retain *direct-style*, some languages build in several different control abstractions, e.g., JavaScript has both *async/await* and generators/iterators, but hard-wiring multiple abstractions increases the complexity of the compiler and run-time.

An alternative is to provide a single control abstraction, and derive others as libraries. Plotkin and Pretnar’s effect handlers provide a modular abstraction that subsumes all of the above control abstractions. Moreover, they have a strong mathematical foundation [20,21] and have found applications across a diverse spectrum of disciplines such as concurrent programming [4], probabilistic programming [8], meta programming [24], and more [12].

With effect handlers computations are viewed as trees. Effect handlers come in two flavours *deep* and *shallow*. Deep handlers are defined by folds (specifically *catamorphisms* [18]) over computation trees, whereas shallow handlers are defined as case-splits. Catamorphisms are attractive because they are semantically well-behaved and provide appropriate structure for efficient implementations using optimisations such as fusion [23]. However, they are not always convenient for implementing other structural recursion schemes such as mutual recursion.

Most existing accounts of effect handlers use deep handlers. In this paper we develop the theory of shallow effect handlers.

As shallow handlers impose no particular structural recursion scheme, they can be more convenient. For instance, using shallow handlers it is easy to model Unix pipes as two mutually recursive functions (specifically *mutumorphisms* [7]) that alternate production and consumption of data. With shallow handlers we define a classic demand-driven Unix pipeline operator as follows

$$\begin{aligned}
 \text{pipe} &: \langle \rangle \rightarrow \alpha! \{ \text{Yield} : \beta \rightarrow \langle \rangle, \langle \rangle \rightarrow \alpha! \{ \text{Await} : \beta \} \} \rightarrow \alpha! \emptyset \\
 \text{copipe} &: \langle \beta \rightarrow \alpha! \{ \text{Await} : \beta \}, \langle \rangle \rightarrow \alpha! \{ \text{Yield} : \beta \rightarrow \langle \rangle \} \rangle \rightarrow \alpha! \emptyset \\
 \text{pipe } \langle p, c \rangle &= \text{handle}^\dagger c \langle \rangle \text{ with } \quad \text{copipe } \langle c, p \rangle = \text{handle}^\dagger p \langle \rangle \text{ with} \\
 &\quad \text{return } x \mapsto x \quad \text{return } x \mapsto x \\
 &\quad \text{Await } r \mapsto \text{copipe } \langle r, p \rangle \quad \text{Yield } p \ r \mapsto \text{pipe } \langle r, \lambda \langle \rangle. c \ p \rangle
 \end{aligned}$$

A pipe takes two thunked computations, a producer p and a consumer c . A computation type $A!E$ is a value type A and an effect E , which enumerates the operations that the computation may perform.

The **pipe** function specifies how to *handle* the operations of its arguments and in doing so performs no operations of its own, thus its effect is pure $\emptyset$. Each of the thunks returns a value of type α . The producer can perform the **Yield** operation, which yields a value of type β and the consumer can perform the **Await** operation, which correspondingly awaits a value of type β . The shallow handler runs the consumer. If the consumer returns a value, then the return clause is executed and simply returns that value as is. If the consumer performs the **Await** operation, then the handler is supplied with a special resumption argument r , which is the continuation of the consumer computation reified as a first-class function. The **copipe** is now invoked with r and the producer as arguments.

The **copipe** function is similar. The arguments are swapped and the consumer now expects a value. The shallow handler runs the producer. If it performs the **Yield** operation, then **pipe** is invoked with the resumption of the producer along with a thunk that applies the resumption of the consumer to the yielded value.

As a simple example consider the composition of a producer that yields a stream of ones, and a consumer that awaits a single value.

$$\begin{aligned}
 &\text{pipe } \langle \text{rec ones } \langle \rangle. \text{do Yield } 1; \text{ones } \langle \rangle, \lambda \langle \rangle. \text{do Await} \rangle \\
 \rightsquigarrow^+ &\text{copipe } \langle \lambda x. x, \text{rec ones } \langle \rangle. \text{do Yield } 1; \text{ones } \langle \rangle \rangle \\
 \rightsquigarrow^+ &\text{pipe } \langle \lambda \langle \rangle. \text{rec ones } \langle \rangle. \text{do Yield } 1; \text{ones } \langle \rangle, \lambda \langle \rangle. 1 \rangle \rightsquigarrow^+ 1
 \end{aligned}$$

(The computation **do** ℓ p performs operation ℓ with parameter p .)

The difference between shallow handlers and deep handlers is that in the latter the original handler is implicitly wrapped around the body of the resumption, meaning that the next effectful operation invocation is necessarily handled by the same handler. Shallow handlers allow the freedom to choose how to handle the next effectful operation; deep handlers do not. Pipes provide the quintessential example for contrasting shallow and deep handlers. To implement pipes with deep handlers, we cannot simply use term level recursion, instead we effectively have to *defunctionalise* [22] the shallow version of pipes using recursive types. Following Kammar et al. [12] we define two mutually recursive types

for producers and consumers, respectively.

$$\begin{aligned}\text{Producer } \alpha \beta &= \langle \rangle \rightarrow (\text{Consumer } \alpha \beta \rightarrow \alpha! \emptyset)! \emptyset \\ \text{Consumer } \alpha \beta &= \beta \rightarrow (\text{Producer } \alpha \beta \rightarrow \alpha! \emptyset)! \emptyset\end{aligned}$$

The underlying idea is *state-passing*: the **Producer** type is an alias for a suspended computation which returns a computation parameterised by a **Consumer** computation. Correspondingly, **Consumer** is an alias for a function that consumes an element of type β and returns a computation parameterised by a **Producer** computation. The ultimate return value has type α . Using these recursive types, we can now give types for deep pipe operators and their implementations.

$$\begin{aligned}\text{pipe}' : (\langle \rangle \rightarrow \alpha! \{\text{Await} : \beta\}) &\rightarrow \text{Producer } \alpha \beta \rightarrow \alpha! \emptyset \\ \text{copipe}' : (\langle \rangle \rightarrow \alpha! \{\text{Yield} : \beta \rightarrow \langle \rangle\}) &\rightarrow \text{Consumer } \alpha \beta \rightarrow \alpha! \emptyset \\ \text{pipe}' c = \text{handle } c \langle \rangle \text{ with} &\quad \text{copipe}' p = \text{handle } p \langle \rangle \text{ with} \\ \quad \text{return } x \mapsto \lambda y. x &\quad \text{return } x \mapsto \lambda y. x \\ \text{Await } r \mapsto \lambda p. p \langle \rangle r &\quad \text{Yield } p r \mapsto \lambda c. c p r \\ \text{runPipe } \langle p; c \rangle = \text{pipe}' c (\lambda \langle \rangle. \text{copipe}' p)\end{aligned}$$

Application of the pipe operator is no longer direct as extra plumbing is required to connect the now decoupled handlers. The observable behaviour of `runPipe` is the same as the shallow `pipe`. Indeed, the above example yields the same result.

$$\text{runPipe } \langle \text{rec ones } \langle \rangle. \text{do Yield } 1; \text{ones } \langle \rangle, \lambda \langle \rangle. \text{do Await} \rangle \rightsquigarrow^+ 1$$

In this paper we make five main contributions, each shedding their own light on the computational differences between deep and shallow handlers:

- A proof that shallow handlers with general recursion can simulate deep handlers up to congruence and that, at the cost of performance, deep handlers can simulate shallow handlers up to administrative reductions (§3).
- The first formal account of an abstract machine for shallow handlers (§4).
- The first formal account of a CPS translation for shallow handlers (§5).
- An implementation of both the abstract machine and the CPS translation as backends for the Links web programming language [2].
- Empirical evaluations of our implementations (§6).

§2 introduces our core calculus of deep and shallow effect handlers. §7 discusses related work and §8 concludes.

2 Handler Calculus

In this section, we present $\lambda^\dagger$, a Church-style row-polymorphic call-by-value calculus for effect handlers. To support comparison within a single language we include both deep and shallow handlers. The calculus is an extension of Hillerström and Lindley’s calculus of extensible deep handlers $\lambda_{\text{eff}}^\rho$ [9] with shallow handlers and recursive functions. Following Hillerström and Lindley, $\lambda^\dagger$ provides a row polymorphic effect type system and is based on fine-grain call-by-value [16], which names each intermediate computation as in A-normal form [6], but unlike A-normal form is closed under β -reduction.

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

Fig. 1: Types, Kinds, and Environments

2.1 Types

The syntax of types, kinds, and environments is given in Fig. 1.

Value Types. Function type $A \rightarrow C$ maps values of type A to computations of type C . Polymorphic type $\forall \alpha^K. C$ is parameterised by a type variable α of kind K . Record type $\langle R \rangle$ represents records with fields constrained by row R . Dually, variant type $[R]$ represents tagged sums constrained by row R .

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

Handler Types. The handler type $C \Rightarrow^\delta D$ represent handlers that transform computations of type C into computations of type D (where δ empty denotes a deep handler and $\delta = \dagger$ a shallow handler).

Row Types. Effect, record, and variant types are given by row types. A *row type* (or just *row*) describes a collection of distinct labels, each annotated by a presence type. A presence type indicates whether a label is *present* with 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 . The row variable in an open row type can be instantiated with additional labels. We identify rows up to reordering of labels. For instance, we consider rows $\ell_1 : P_1; \dots; \ell_n : P_n; \cdot$ and $\ell_n : P_n; \dots; \ell_1 : P_1; \cdot$ equivalent. Absent labels in closed rows are redundant. The unit type is the empty closed record, that is, $\langle \cdot \rangle$. Dually, the empty type is the empty, closed variant $[\cdot]$. Often we omit the $\cdot$ for closed rows.

Kinds. We have six kinds: *Type*, *Comp*, *Effect*, *Handler*, *Row_L*, *Presence*, which respectively classify value types, computation types, effect types, row types, presence types, and handler types. Row kinds are annotated with a set of labels $\mathcal{L}$. The kind of a complete row is *Row_∅*. More generally, *Row_L* denotes a partial row that may not mention labels in $\mathcal{L}$. We write $\ell : A$ as sugar for $\ell : \text{Pre}(A)$.

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

Fig. 2: Term Syntax

Type Variables. We let α, ρ and θ range over type variables. By convention we write 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 (Γ) map term variables to their types and kind environments (Δ) map type variables to their kinds.

2.2 Terms

The terms are given in Fig. 2. We let x, y, z, r, p range over term variables. By convention, we use r to denote resumption names. The syntax partitions terms into values, computations and handlers. Value terms comprise variables (x), lambda abstraction ($\lambda x^A.M$), type abstraction ($\Lambda \alpha^K.M$), the introduction forms for records and variants, and recursive functions ($\mathbf{rec} \, g^{A \rightarrow C} x.M$). 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 .

All elimination forms are computation terms. Abstraction and type abstraction are eliminated using application ($V W$) and type application ($V T$) respectively. The record eliminator ($\mathbf{let} \, \langle \ell = x; y \rangle = V \mathbf{in} \, N$) 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 ($\mathbf{case} \, V \{ \ell \, x \mapsto M; y \mapsto N \}$), 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 ($\mathbf{absurd}^C V$). A trivial computation ($\mathbf{return} \, V$) returns value V . The expression ($\mathbf{let} \, x \leftarrow M \mathbf{in} \, N$) evaluates M and binds the result to x in N .

Operation invocation ($\mathbf{do} \, \ell \, V$)^E performs operation ℓ with value argument V . Handling ($\mathbf{handle}^\delta M \mathbf{with} \, H$) runs a computation M using deep (δ empty) or shallow ($\delta = \dagger$) handler H . A handler definition H consists of a return clause $\{ \mathbf{return} \, x \mapsto M \}$ and a possibly empty set of operation clauses $\{ \ell \, p \, r \mapsto N_\ell \}_{\ell \in \mathcal{L}}$. The return clause defines how to handle the final return value of the handled computation, which is bound to x in M . The operation clause for ℓ binds the operation parameter to p and the resumption r in N_ℓ .

We define three projections on handlers: H^{ret} yields the singleton set containing the return clause of H and H^ℓ yields the set of either zero or one operation

clauses in H that handle the operation ℓ and H^{ops} yields the set of all operation clauses in H . We write $\text{dom}(H)$ for the set of operations handled by H . Various term forms are annotated with type or kind information; we sometimes omit such annotations. We write $\text{Id}(M)$ for **handle** M **with** $\{\text{return } x \mapsto \text{return } x\}$.

Syntactic sugar. We make use of standard syntactic sugar for pattern matching, n -ary record extension, n -ary case elimination, and n -ary tuples.

2.3 Kinding and Typing

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 $\Delta; \Gamma \vdash H : C \Rightarrow^\delta D$ states that handler H has type $C \Rightarrow^\delta D$ 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 $\text{FTV}(\Gamma)$ to be the set of free type variables in Γ . We omit the full kinding and typing rules due to lack of space; they can be found in the extended version of the paper [10, Appendix A]. The interesting rules are those for performing and handling operations.

$$\begin{array}{c}
\text{T-Do} \\
\frac{\Delta; \Gamma \vdash V : A \quad E = \{\ell : A \rightarrow B; R\}}{\Delta; \Gamma \vdash (\text{do } \ell \ V)^E : B!E}
\end{array}
\qquad
\begin{array}{c}
\text{T-HANDLE} \\
\frac{\Delta; \Gamma \vdash M : C \quad \Delta; \Gamma \vdash H : C \Rightarrow^\delta D}{\Delta; \Gamma \vdash \text{handle}^\delta M \text{ with } H : D}
\end{array}$$

$$\begin{array}{c}
\text{T-HANDLER} \\
\begin{array}{l}
C = A!\{(\ell_i : A_i \rightarrow B_i)_i; R\} \\
D = B!\{(\ell_i : P_i)_i; R\} \\
H = \{\text{return } x \mapsto M\} \uplus \{\ell_i \ p \ r \mapsto N_i\}_i \\
\Delta; \Gamma, x : A \vdash M : D \\
[\Delta; \Gamma, p : A_i, r : B_i \rightarrow D \vdash N_i : D]_i
\end{array} \\
\hline
\Delta; \Gamma \vdash H : C \Rightarrow D
\end{array}
\qquad
\begin{array}{c}
\text{T-HANDLER}^\dagger \\
\begin{array}{l}
C = A!\{(\ell_i : A_i \rightarrow B_i)_i; R\} \\
D = B!\{(\ell_i : P_i)_i; R\} \\
H = \{\text{return } x \mapsto M\} \uplus \{\ell_i \ p \ r \mapsto N_i\}_i \\
\Delta; \Gamma, x : A \vdash M : D \\
[\Delta; \Gamma, p : A_i, r : B_i \rightarrow C \vdash N_i : D]_i
\end{array} \\
\hline
\Gamma \vdash H : C \Rightarrow^\dagger D
\end{array}$$

The T-HANDLER and T-HANDLER[†] rules are where most of the work happens. 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 to say whether 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 difference in typing deep and shallow handlers is that the resumption of the former has return type D , whereas the resumption of the latter has return type C .

2.4 Operational Semantics

Figure 3 gives a small-step operational semantics for $\lambda^\dagger$. The reduction relation $\rightsquigarrow$ is defined on computation terms. The interesting rules are the handler rules.

S-APP	$(\lambda x.M) V \rightsquigarrow M[V/x]$
S-TYAPP	$(\Lambda \alpha.M) A \rightsquigarrow M[A/\alpha]$
S-SPLIT	$\text{let } \langle \ell = x; y \rangle = \langle \ell = V; W \rangle \text{ in } N \rightsquigarrow N[V/x, W/y]$
S-CASE ₁	$\text{case } \ell V \{ \ell x \mapsto M; y \mapsto N \} \rightsquigarrow M[V/x]$
S-CASE ₂	$\text{case } \ell V \{ \ell' x \mapsto M; y \mapsto N \} \rightsquigarrow N[\ell V/y], \quad \text{if } \ell \neq \ell'$
S-REC	$(\text{rec } g x.M) V \rightsquigarrow M[(\text{rec } g x.M)/g, V/x]$
S-LET	$\text{let } x \leftarrow \text{return } V \text{ in } N \rightsquigarrow N[V/x]$
S-RET	$\text{handle}^\delta(\text{return } V) \text{ with } H \rightsquigarrow N[V/x],$ where $H^{\text{ret}} = \{\text{return } x \mapsto N\}$
S-OP	$\text{handle } \mathcal{E}[\text{do } \ell V] \text{ with } H \rightsquigarrow$ $N[V/p, \lambda y. \text{handle } \mathcal{E}[\text{return } y] \text{ with } H/r],$ where $\ell \notin BL(\mathcal{E})$ and $H^\ell = \{\ell p r \mapsto N\}$
S-OP [†]	$\text{handle}^\dagger \mathcal{E}[\text{do } \ell V] \text{ with } H \rightsquigarrow N[V/p, \lambda y. \mathcal{E}[\text{return } y]/r],$ where $\ell \notin BL(\mathcal{E})$ and $H^\ell = \{\ell p r \mapsto N\}$
S-LIFT	$\mathcal{E}[M] \rightsquigarrow \mathcal{E}[N], \quad \text{if } M \rightsquigarrow N$
Evaluation contexts $\mathcal{E} ::= [\] \mid \text{let } x \leftarrow \mathcal{E} \text{ in } N \mid \text{handle}^\delta \mathcal{E} \text{ with } H$	

Fig. 3: Small-Step Operational Semantics

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

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

The S-RET rule invokes the return clause of a handler. The S-OP^δ rules handle an operation by invoking the appropriate operation clause. The constraint $\ell \notin BL(\mathcal{E})$ asserts that no handler in the evaluation context handles the operation: a handler reaches past any other inner handlers that do not handle ℓ . The difference between S-OP and S-OP[†] is that the former rewraps the handler about the body of the resumption. We write R^+ for transitive closure of relation R .

Definition 1. *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 $\ell \notin BL(\mathcal{E})$.*

Theorem 2 (Type Soundness). *If $\vdash M : A!E$ then either $M \not\rightsquigarrow^*$ or there exists $\vdash N : A!E$ such that $M \rightsquigarrow^+ N \not\rightsquigarrow$ and N is normal with respect to E .*

3 Deep as Shallow and Shallow as Deep

In this section we show that shallow handlers and general recursion can simulate deep handlers up to congruence, and that deep handlers can simulate shallow handlers up to administrative reduction. Both translations are folklore, but we believe the precise simulation results are novel.

3.1 Deep as Shallow

The implementation of deep handlers using shallow handlers (and recursive functions) is by a rather direct local translation. Each handler is wrapped in a recursive function and each resumption has its body wrapped in a call to this recursive function. Formally, the translation $\mathcal{S}[-]$ is defined as the homomorphic extension of the following equations to all terms.

$$\begin{aligned}\mathcal{S}[\text{handle } M \text{ with } H] &= (\text{rec } h \text{ f. handle}^\dagger f \langle \rangle \text{ with } \mathcal{S}[H]h) (\lambda \langle \rangle. \mathcal{S}[M]) \\ \mathcal{S}[H]h &= \mathcal{S}[H^{\text{ret}}]h \uplus \mathcal{S}[H^{\text{ops}}]h \\ \mathcal{S}[\{\text{return } x \mapsto N\}]h &= \{\text{return } x \mapsto \mathcal{S}[N]\} \\ \mathcal{S}[\{\ell \text{ p } r \mapsto N_\ell\}_{\ell \in \mathcal{L}}]h &= \{\ell \text{ p } r \mapsto \text{let } r \leftarrow \text{return } \lambda x. h (\lambda \langle \rangle. r x) \text{ in } \mathcal{S}[N_\ell]\}_{\ell \in \mathcal{L}}\end{aligned}$$

Theorem 3. *If $\Delta; \Gamma \vdash M : C$ then $\Delta; \Gamma \vdash \mathcal{S}[M] : C$.*

In order to obtain a simulation result, we allow reduction in the simulated term to be performed under lambda abstractions (and indeed anywhere in a term), which is necessary because of the redefinition of the resumption to wrap the handler around its body. Nevertheless, the simulation proof makes minimal use of this power, merely using it to rename a single variable. We write R_{cong} for the compatible closure of relation R , that is the smallest relation including R and closed under term constructors for $\lambda^\dagger$.

Theorem 4 (Simulation up to Congruence). *If $M \rightsquigarrow N$ then $\mathcal{S}[M] \rightsquigarrow_{\text{cong}}^+ \mathcal{S}[N]$.*

Proof. By induction on $\rightsquigarrow$ using a substitution lemma. The interesting case is S-DEEP-OP, which is where we apply a single β -reduction, renaming a variable, under the lambda abstraction representing the resumption.

3.2 Shallow as Deep

Implementing shallow handlers in terms of deep handlers is slightly more involved than the other way round. It amounts to the encoding of a case split by a fold and involves a translation on handler types as well as handler terms. Formally, the translation $\mathcal{D}[-]$ is defined as the homomorphic extension of the following equations to all types, terms, and type environments.

$$\begin{aligned}\mathcal{D}[C \Rightarrow D] &= \mathcal{D}[C] \Rightarrow \langle \rangle \rightarrow \mathcal{D}[C], \langle \rangle \rightarrow \mathcal{D}[D] \rangle \\ \mathcal{D}[\text{handle}^\dagger M \text{ with } H] &= \text{let } z \leftarrow \text{handle } \mathcal{D}[M] \text{ with } \mathcal{D}[H] \text{ in} \\ &\quad \text{let } \langle f, g \rangle = z \text{ in } g \langle \rangle \\ \mathcal{D}[H] &= \mathcal{D}[H^{\text{ret}}] \uplus \mathcal{D}[H^{\text{ops}}] \\ \mathcal{D}[\{\text{return } x \mapsto N\}] &= \{\text{return } x \mapsto \text{return } \langle \lambda \langle \rangle. \text{return } x, \lambda \langle \rangle. \mathcal{D}[N] \rangle\} \\ \mathcal{D}[\{\ell \text{ p } r \mapsto N\}_{\ell \in \mathcal{L}}] &= \{\ell \text{ p } r \mapsto \\ &\quad \text{let } r = \lambda x. \text{let } z \leftarrow r x \text{ in let } \langle f, g \rangle = z \text{ in } f \langle \rangle \text{ in} \\ &\quad \text{return } \langle \lambda \langle \rangle. \text{let } x \leftarrow \text{do } \ell \text{ p } \text{ in } r x, \lambda \langle \rangle. \mathcal{D}[N] \rangle\}_{\ell \in \mathcal{L}}\end{aligned}$$

Each shallow handler is encoded as a deep handler that returns a pair of thunks. The first forwards all operations, acting as the identity on computations. The second interprets a single operation before reverting to forwarding.

Theorem 5. *If $\Delta; \Gamma \vdash M : C$ then $\mathcal{D}[\Delta]; \mathcal{D}[\Gamma] \vdash \mathcal{D}[M] : \mathcal{D}[C]$.*

As with the implementation of deep handlers as shallow handlers, the implementation is again given by a local translation. However, this time the administrative overhead is more significant. Reduction up to congruence is insufficient and we require a more semantic notion of administrative reduction.

Definition 6 (Administrative Evaluation Contexts). *An evaluation context $\mathcal{E}$ is administrative, $\text{admin}(\mathcal{E})$, iff*

1. *For all values V , we have: $\mathcal{E}[\mathbf{return} V] \rightsquigarrow^* \mathbf{return} V$*
2. *For all evaluation contexts $\mathcal{E}'$, operations $\ell \in BL(\mathcal{E}) \setminus BL(\mathcal{E}')$, values V :*

$$\mathcal{E}[\mathcal{E}'[\mathbf{do} \ell V]] \rightsquigarrow^* \mathbf{let} x \leftarrow \mathbf{do} \ell V \mathbf{in} \mathcal{E}[\mathcal{E}'[\mathbf{return} x]]$$

The intuition is that an administrative evaluation context behaves like the empty evaluation context up to some amount of administrative reduction, which can only proceed once the term in the context becomes sufficiently evaluated. Values annihilate the evaluation context and handled operations are forwarded.

Definition 7 (Approximation up to Administrative Reduction). *Define $\gtrsim$ as the compatible closure of the following inference rules.*

$$\frac{}{M \gtrsim M} \qquad \frac{M \rightsquigarrow M' \quad M' \gtrsim N}{M \gtrsim N} \qquad \frac{\text{admin}(\mathcal{E}) \quad M \gtrsim N}{\mathcal{E}[M] \gtrsim N}$$

We say that M approximates N up to administrative reduction if $M \gtrsim N$.

Approximation up to administrative reduction captures the property that administrative reduction may occur anywhere within a term. The following lemma states that the forwarding component of the translation is administrative.

Lemma 8. *For all shallow handlers H , the following context is administrative:*

$$\mathbf{let} z \leftarrow \mathbf{handle} [\] \mathbf{with} \mathcal{D}[H] \mathbf{in} \mathbf{let} \langle f; _ \rangle = z \mathbf{in} f \langle \rangle$$

Theorem 9 (Simulation up to Administrative Reduction). *If $M' \gtrsim \mathcal{D}[M]$ and $M \rightsquigarrow N$ then there exists N' such that $N' \gtrsim \mathcal{D}[N]$ and $M' \rightsquigarrow^+ N'$.*

Proof. By induction on $\rightsquigarrow$ using a substitution lemma and Lemma 8. The interesting case is $\text{S-Op}^\dagger$, which uses Lemma 8 to approximate the body of the resumption up to administrative reduction.

4 Abstract Machine

In this section we develop an abstract machine that supports deep and shallow handlers *simultaneously*. We build upon prior work [9] in which we developed an abstract machine for deep handlers by generalising the continuation structure of a CEK machine (Control, Environment, Kontinuation) [5]. In our prior work we sketched an adaptation for shallow handlers. It turns out that this adaptation has a subtle flaw. We fix the flaw here with a full development of shallow handlers along with a proof of correctness.

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

Fig. 4: Abstract Machine Syntax

The informal account. A machine continuation is a list of handler frames. A handler frame is a pair of a *handler closure* (handler definition) and a *pure continuation* (a sequence of let bindings). Handling an operation amounts to searching through the continuation for a matching handler. The resumption is constructed during the search by reifying each handler frame. The resumption is assembled in one of two ways depending on whether the matching handler is deep or shallow. For a deep handler, the current handler closure is included, and a deep resumption is a reified continuation. An invocation of a deep resumption amounts to concatenating it with the current machine continuation. For a shallow handler, the current handler closure must be discarded leaving behind a dangling pure continuation, and a shallow resumption is a pair of this pure continuation and the remaining reified continuation. (By contrast, the prior flawed adaptation prematurely precomposed the pure continuation with the outer handler in the current resumption.) An invocation of a shallow resumption again amounts to concatenating it with the current machine continuation, but taking care to concatenate the dangling pure continuation with that of the next frame.

The formal account. The abstract machine syntax is given in Fig. 4. A configuration $\mathcal{C} = \langle M \mid \gamma \mid \kappa \circ \kappa' \rangle$ of our abstract machine is a quadruple of a computation term (M), an environment (γ) mapping free variables to values, and two continuations (κ) and (κ'). The latter continuation is always the identity, except when forwarding an operation, in which case it is used to keep track of the extent to which the operation has been forwarded. We write $\langle M \mid \gamma \mid \kappa \rangle$ as syntactic sugar for $\langle M \mid \gamma \mid \kappa \circ [] \rangle$ where $[]$ is the identity continuation.

Values consist of function closures, type function closures, records, variants, and captured continuations. A continuation κ is a stack of frames $[\theta_1, \dots, \theta_n]$. We annotate captured continuations with input types in order to make the results of §4.1 easier to state. Each frame $\theta = (\sigma, \chi)$ represents pure continuation σ , corresponding to a sequence of let bindings, inside handler closure χ . A pure continuation is a stack of pure frames. A pure 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 $s ++ s'$ for the concatenation of stack s on top of s' . We use pattern matching to deconstruct stacks.

Transition function	
M-INIT	$M \longrightarrow \langle M \mid \emptyset \mid [(\emptyset, \{\text{return } x \mapsto x\})] \rangle$
M-APPCLOSURE	$\langle V \mid W \mid \gamma \mid \kappa \rangle \longrightarrow \langle M \mid \gamma[x \mapsto \llbracket W \rrbracket \gamma] \mid \kappa \rangle,$
M-APPREC	$\langle V \mid W \mid \gamma \mid \kappa \rangle \longrightarrow \langle M \mid \gamma[g \mapsto (\gamma', \text{rec } g^{A \rightarrow C} x.M), x \mapsto \llbracket W \rrbracket \gamma] \mid \kappa \rangle,$
M-APPCONT	$\langle V \mid W \mid \gamma \mid \kappa \rangle \longrightarrow \langle \text{return } W \mid \gamma \mid \kappa' \dashv \kappa \rangle,$
M-APPCONT [†]	$\langle V \mid W \mid \gamma \mid (\sigma, \chi) :: \kappa \rangle \longrightarrow \langle \text{return } W \mid \gamma \mid \kappa' \dashv ((\sigma' \dashv \sigma, \chi) :: \kappa) \rangle,$
M-APPTYPE	$\langle V \mid A \mid \gamma \mid \kappa \rangle \longrightarrow \langle M[A/\alpha] \mid \gamma' \mid \kappa \rangle,$
M-SPLIT	$\langle \text{let } \ell = x; y \rangle = V \text{ in } N \mid \gamma \mid \kappa \rangle \longrightarrow \langle N \mid \gamma[x \mapsto v, y \mapsto w] \mid \kappa \rangle,$
M-CASE	$\langle \text{case } V \mid \{ \ell x \mapsto M; y \mapsto N \} \mid \gamma \mid \kappa \rangle \longrightarrow \begin{cases} \langle M \mid \gamma[x \mapsto v] \mid \kappa \rangle, \\ \langle N \mid \gamma[y \mapsto \ell' v] \mid \kappa \rangle, \end{cases}$
M-LET	$\langle \text{let } x \leftarrow M \text{ in } N \mid \gamma \mid (\sigma, \chi) :: \kappa \rangle \longrightarrow \langle M \mid \gamma \mid ((\gamma, x, N) :: \sigma, \chi) :: \kappa \rangle$
M-HANDLE	$\langle \text{handle}^\delta M \text{ with } H \mid \gamma \mid \kappa \rangle \longrightarrow \langle M \mid \gamma \mid ([\gamma], (\gamma, H)^\delta) :: \kappa \rangle$
M-RETCONT	$\langle \text{return } V \mid \gamma \mid ((\gamma', x, N) :: \sigma, \chi) :: \kappa \rangle \longrightarrow \langle N \mid \gamma'[x \mapsto \llbracket V \rrbracket \gamma] \mid (\sigma, \chi) :: \kappa \rangle$
M-RETHANDLER	$\langle \text{return } V \mid \gamma \mid ([\gamma], (\gamma', H)) :: \kappa \rangle \longrightarrow \langle M \mid \gamma'[x \mapsto \llbracket V \rrbracket \gamma] \mid \kappa \rangle,$
M-RETTOP	$\langle \text{return } V \mid \gamma \mid [] \rangle \longrightarrow \llbracket V \rrbracket \gamma$
M-DO	$\langle (\text{do } \ell V)^E \mid \gamma \mid ((\sigma, (\gamma', H)) :: \kappa) \circ \kappa' \rangle \longrightarrow \langle M \mid \gamma[x \mapsto \llbracket V \rrbracket \gamma, r \mapsto (\kappa' \dashv \llbracket \sigma, (\gamma', H) \rrbracket^B)] \mid \kappa \rangle,$
M-DO [†]	$\langle (\text{do } \ell V)^E \mid \gamma \mid ((\sigma, (\gamma', H)^\dagger) :: \kappa) \circ \kappa' \rangle \longrightarrow \langle M \mid \gamma[x \mapsto \llbracket V \rrbracket \gamma, r \mapsto (\kappa', \sigma)^{B^\dagger}] \mid \kappa \rangle,$
M-FORWARD	$\langle (\text{do } \ell V)^E \mid \gamma \mid ((\sigma, (\gamma', H)^\delta) :: \kappa) \circ \kappa' \rangle \longrightarrow \langle (\text{do } \ell V)^E \mid \gamma \mid \kappa \circ (\kappa' \dashv \llbracket \sigma, (\gamma', H)^\delta \rrbracket) \rangle, \text{ if } H^\ell = \emptyset$
Value interpretation	
$\llbracket x \rrbracket \gamma = \gamma(x)$	$\llbracket \lambda x^A. M \rrbracket \gamma = (\gamma, \lambda x^A. M)$
$\llbracket \langle \rangle \rrbracket \gamma = \langle \rangle$	$\llbracket \ell = V; W \rrbracket \gamma = \langle \ell = \llbracket V \rrbracket \gamma; \llbracket W \rrbracket \gamma \rangle$
	$\llbracket A\alpha^K. M \rrbracket \gamma = (\gamma, A\alpha^K. M)$
	$\llbracket \text{rec } g^{A \rightarrow C} x. M \rrbracket \gamma = (\gamma, \text{rec } g^{A \rightarrow C} x. M)$

Fig. 5: Abstract Machine Semantics

The abstract machine semantics defining the transition function $\longrightarrow$ is given in Fig. 5. It 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. The rules (M-APPCLOSURE), (M-APPREC), (M-APPCONT), (M-APPCONT[†]), (M-APPTYPE), (M-SPLIT), and (M-CASE) enact the elimination of values. 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; (M-RETHANDLER) invokes the return clause of a handler if the pure continuation is empty; and (M-RETTOP) returns a final value if the continuation is empty. The rule (M-DO) applies the current handler to an operation if the label matches one of the operation clauses. The captured continuation is assigned the forwarding continuation with the current frame appended to the end of it. The rule (M-DO[†]) is much like (M-DO), except it constructs a shallow resumption, discarding the current handler but keeping the current pure continuation. The rule (M-FORWARD) appends the current continuation frame onto the end of the forwarding continuation.

4.1 Correctness

The (M-INIT) rule provides a canonical way to map a computation term onto a configuration. Fig. 6 defines an inverse mapping $\langle - \rangle$ from configurations to computation 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 $\langle - \rangle$ 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 a feature of the operational semantics. The rules (M-APPCONT^δ), (M-LET), (M-HANDLE), and (M-FORWARD) are administrative in that $\langle - \rangle$ is invariant under them. This leaves β -rules (M-APPCLOSURE), (M-APPREC), (M-APPTYPE), (M-SPLIT), (M-CASE), (M-RETCONT), (M-RETHANDLER), (M-DO[†]), and (M-DO[†]), each of which corresponds directly to 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 $\longrightarrow_a^* \longrightarrow_\beta$.

Each reduction in the operational semantics is simulated by a sequence of administrative steps followed by a single β -step in the abstract machine. The *Id* handler (§2.2) implements the top-level identity continuation.

Theorem 10 (Simulation). *If $M \rightsquigarrow N$, then for any C such that $\langle C \rangle = \text{Id}(M)$ there exists C' such that $C \Longrightarrow C'$ and $\langle C' \rangle = \text{Id}(N)$.*

Proof. By induction on the derivation of $M \rightsquigarrow N$.

Corollary 11. *If $\vdash M : A!E$ and $M \rightsquigarrow^+ N \not\rightsquigarrow$, then $M \longrightarrow^+ C$ with $\langle C \rangle = N$.*

Configurations

$$\langle\langle M \mid \gamma \mid \kappa \circ \kappa' \rangle\rangle = \langle\kappa' \dashv \kappa\rangle(\langle M \rangle \gamma) = \langle\kappa'\rangle(\langle\kappa\rangle(\langle M \rangle \gamma))$$

Pure continuations

$$\langle\langle \rangle\rangle M = M \quad \langle\langle (\gamma, x, N) :: \sigma \rangle\rangle M = \langle\sigma\rangle(\mathbf{let} \ x \leftarrow M \ \mathbf{in} \ \langle N \rangle(\gamma \setminus \{x\}))$$

Continuations

$$\langle\langle \rangle\rangle M = M \quad \langle\langle (\sigma, \chi) :: \kappa \rangle\rangle M = \langle\kappa\rangle(\langle\chi\rangle(\langle\sigma\rangle(M)))$$

Handler closures

$$\langle\langle (\gamma, H) \rangle\rangle^\delta M = \mathbf{handle}^\delta M \ \mathbf{with} \ \langle H \rangle \gamma$$

Computation terms

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

Handler definitions

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

Value terms and values

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

Fig. 6: Mapping from Abstract Machine Configurations to Terms

5 Higher-Order CPS Translation

In this section we formalise a CPS translation for deep and shallow handlers. We adapt the higher-order translation of Hillerström et al. [11]. They formalise a translation for deep handlers and then briefly outline an extension for shallow handlers. Alas, there is a bug in their extension. Their deep handler translation takes advantage of the rewinding of the body of a resumption with the current handler to combine the current return clause with the current pure continuation. Their shallow handler translation attempts to do the same, but the combination

Syntax

$$\begin{array}{ll}
\text{Values} & V, W ::= x \mid \underline{\lambda} x k.M \mid \mathbf{rec} \, g \, x k.M \mid \ell \mid \underline{\langle V, W \rangle} \mid \mathbf{res}^\delta \\
\text{Computations} & M, N ::= V \mid U \underline{\textcircled{@}} V \underline{\textcircled{@}} W \mid \mathbf{let} \, \underline{\langle x, y \rangle} = V \mathbf{in} \, N \\
& \mid \mathbf{case} \, V \{ \ell \mapsto M; x \mapsto N \} \mid \mathbf{app} \, V \, W
\end{array}$$

Syntactic sugar

$$\begin{array}{ll}
\mathbf{let} \, x = V \mathbf{in} \, N \equiv N[V/x] & \langle \rangle \equiv \ell_\langle \rangle \quad [] \equiv \ell_\square \\
\ell \, V \equiv \underline{\langle \ell, V \rangle} & \langle \ell = V; W \rangle \equiv \ell \, \underline{\langle V, W \rangle} \quad V \underline{\textcircled{@}} W \equiv \ell \underline{\textcircled{@}} \underline{\langle V, W \rangle} \\
\mathbf{case} \, V \{ \ell \, x \mapsto M; y \mapsto N \} \equiv & \mathbf{let} \, \langle \ell = x; y \rangle = V \mathbf{in} \, N \equiv \\
\mathbf{let} \, y = V \mathbf{in} \, \mathbf{let} \, \underline{\langle z, x \rangle} = y \mathbf{in} & \mathbf{let} \, \underline{\langle z, z' \rangle} = V \mathbf{in} \, \mathbf{let} \, \underline{\langle x, y \rangle} = z' \mathbf{in} \\
\mathbf{case} \, z \{ \ell \mapsto M; z \mapsto N \} & \mathbf{case} \, z \{ \ell \mapsto N; z \mapsto \ell_\perp \}
\end{array}$$

Reductions

$$\begin{array}{ll}
\text{U-APP} & (\underline{\lambda} x k.M) \underline{\textcircled{@}} V \underline{\textcircled{@}} W \rightsquigarrow M[V/x, W/k] \\
\text{U-REC} & (\mathbf{rec} \, g \, x k.M) \underline{\textcircled{@}} V \underline{\textcircled{@}} W \rightsquigarrow M[\mathbf{rec} \, g \, x k.M/g, V/x, W/k] \\
\text{U-SPLIT} & \mathbf{let} \, \underline{\langle x, y \rangle} = \underline{\langle V, W \rangle} \mathbf{in} \, N \rightsquigarrow N[V/x, W/y] \\
\text{U-CASE}_1 & \mathbf{case} \, \ell \{ \ell' \mapsto M; x \mapsto N \} \rightsquigarrow M \\
\text{U-CASE}_2 & \mathbf{case} \, \ell \{ \ell' \mapsto M; x \mapsto N \} \rightsquigarrow N[\ell/x], \quad \text{if } \ell \neq \ell' \\
\text{U-KAPPNIL} & \mathbf{app} \, (\underline{\langle \rangle}, \underline{\langle v, e \rangle}) \underline{\textcircled{@}} k \rightsquigarrow v \underline{\textcircled{@}} V \underline{\textcircled{@}} k \\
\text{U-KAPPCONS} & \mathbf{app} \, (\underline{\langle f \underline{\textcircled{@}} s, h \rangle}) \underline{\textcircled{@}} k \rightsquigarrow f \underline{\textcircled{@}} V \underline{\textcircled{@}} (\underline{\langle s, h \rangle}) \underline{\textcircled{@}} k \\
\text{U-RES} & \mathbf{res} \, (q_n \underline{\textcircled{@}} \dots \underline{\textcircled{@}} q_1 \underline{\textcircled{@}} []) \rightsquigarrow \lambda x k. \mathbf{app} \, (q_1 \underline{\textcircled{@}} \dots \underline{\textcircled{@}} q_n \underline{\textcircled{@}} k) \, x \\
\text{U-RES}^\dagger & \mathbf{res}^\dagger \, (\underline{\langle f_1 \underline{\textcircled{@}} \dots \underline{\textcircled{@}} f_m, h \rangle}) \underline{\textcircled{@}} q_n \underline{\textcircled{@}} \dots \underline{\textcircled{@}} q_1 \underline{\textcircled{@}} [] \rightsquigarrow \\
& \lambda x k. \mathbf{let} \, (\underline{\langle s', h' \rangle}) \underline{\textcircled{@}} k' = k \mathbf{in} \\
& \mathbf{app} \, (q_1 \underline{\textcircled{@}} \dots \underline{\textcircled{@}} q_n \underline{\textcircled{@}} \underline{\langle f_1 \underline{\textcircled{@}} \dots \underline{\textcircled{@}} f_m \underline{\textcircled{@}} s', h' \rangle}) \underline{\textcircled{@}} k' \, x
\end{array}$$

Fig. 7: Untyped Target Calculus

is now unsound as the return clause must be discarded by the resumption. We fix the bug by explicitly separating out the return continuation. Moreover, our translation is carefully designed to avoid memory leaks. The key insight is that to support the typical tail-recursive pattern of shallow handlers without generating useless identity continuations it is essential that we detect and eliminate them. We do so by representing pure continuations as lists of pure frames whereby the identity continuation is just an empty list, much like the abstract machine of §4.

Following Hillerström et al. [11], we present a higher-order uncurried CPS translation into an untyped lambda calculus. In the style of Danvy and Nielsen [3], we adopt a two-level lambda calculus notation to distinguish between *static* lambda abstraction and application in the meta language and *dynamic* lambda abstraction and application in the target language: underline denotes a static syntax constructor; underline denotes a dynamic syntax constructor. To facilitate this notation we write application as an infix “at” symbol ($\underline{\textcircled{@}}$). We assume the meta language is pure and hence respects the usual β and η equivalences.

5.1 Target Calculus

The target calculus is given in Fig. 7. As in $\lambda^\dagger$ there is a syntactic distinction between values (V) and computations (M). Values (V) comprise: lambda abstractions ($\underline{\lambda}x\ k.M$) and recursive functions ($\mathbf{rec}\ g\ x\ k.M$), each of which take an additional continuation parameter; first-class labels (ℓ); pairs ($\underline{\langle}V, W\underline{\rangle}$); and two special convenience constructors for building deep ($\mathbf{res}\ V$) and shallow ($\mathbf{res}^\dagger V$) resumptions, which we will explain shortly. Computations (M) comprise: values (V); applications ($U\ @\ V\ @\ W$); pair elimination ($\mathbf{let}\ \underline{\langle}x, y\underline{\rangle} = V\ \mathbf{In}\ N$); label elimination ($\mathbf{case}\ V\ \{\ell \mapsto M; x \mapsto N\}$); and a special convenience constructor for continuation application ($\mathbf{app}\ V\ W$).

Lambda abstraction, pairs, application, and pair elimination are underlined to distinguish them from equivalent constructs in the meta language. We define syntactic sugar for variant values, record values, list values, let binding, variant eliminators, and record eliminators. We assume standard n -ary generalisations and use pattern matching syntax for deconstructing variants, records, and lists.

The reductions for functions, pairs, and first-class labels are standard. To explain the reduction rules for continuations, we first explain the encoding of continuations. Much like the abstract machine, a continuation (k) is given by a list of continuation frames. A continuation frame ($\underline{\langle}s, h\underline{\rangle}$) consists of a pair of a pure continuation (s) and a handler (h). A pure continuation is a list of pure continuation frames (f). A handler is a pair of a return continuation (v) and an effect continuation (e) which dispatches on the operations provided by a handler. There are two continuation reduction rules, both of which inspect the first frame of the continuation. If the pure continuation of this frame is empty then the return clause is invoked (U-KAPPNIL). If the pure continuation of this frame is non-empty then the first pure continuation frame is invoked (U-KAPPCONS). A crucial difference between our representation of continuations and that of Hillerström et al. [11] is that they use a flat list of frames whereas we use a nested structure in which each pure continuation is a list of pure frames.

To explain the reduction rules for continuations, we first explain the encoding of resumptions. Reified resumptions are constructed frame-by-frame as reversed continuations — they grow a frame at a time as operations are forwarded through the handler stack. Hillerström et al. [11] adopt such an intensional representation in order to obtain a relatively tight simulation result. We take further advantage of this representation to discard the handler when constructing a shallow handler's resumption. The resumption reduction rules turn reified resumptions into actual resumptions. The deep rule (U-RES) simply appends the reified resumption onto the continuation. The shallow rule (U-RES †) appends the tail of the reified resumption onto the continuation after discarding the topmost handler from the resumption and appending the topmost pure continuation from the resumption onto the topmost pure continuation of the continuation.

The continuation application and resumption constructs along with their reduction rules are macro-expressible in terms of the standard constructs. We choose to build them in order to keep the presentation relatively concise.

Values

$$\begin{aligned}
\llbracket x \rrbracket &= x & \llbracket \lambda \alpha. M \rrbracket &= \lambda z k. \llbracket M \rrbracket \bar{\textcircled{a}} k & \llbracket \langle \ell = V; W \rangle \rrbracket &= \langle \ell = \llbracket V \rrbracket; \llbracket W \rrbracket \rangle \\
\llbracket \lambda x. M \rrbracket &= \lambda x k. \llbracket M \rrbracket \bar{\textcircled{a}} k & \llbracket \langle \rangle \rrbracket &= \langle \rangle & \llbracket \ell V \rrbracket &= \ell \llbracket V \rrbracket \\
\llbracket \text{rec } g x. M \rrbracket &= \text{rec } g x k. \llbracket M \rrbracket \bar{\textcircled{a}} k
\end{aligned}$$

Computations

$$\begin{aligned}
\llbracket V W \rrbracket &= \bar{\lambda} \kappa. \llbracket V \rrbracket \bar{\textcircled{a}} \llbracket W \rrbracket \bar{\textcircled{a}} \downarrow \kappa \\
\llbracket V T \rrbracket &= \bar{\lambda} \kappa. \llbracket V \rrbracket \bar{\textcircled{a}} \langle \rangle \bar{\textcircled{a}} \downarrow \kappa \\
\llbracket \text{let } \langle \ell = x; y \rangle = V \text{ in } N \rrbracket &= \bar{\lambda} \kappa. \text{let } \langle \ell = x; y \rangle = \llbracket V \rrbracket \text{ in } \llbracket N \rrbracket \bar{\textcircled{a}} \kappa \\
\llbracket \text{case } V \{ \ell x \mapsto M; y \mapsto N \} \rrbracket &= \bar{\lambda} \kappa. \text{case } \llbracket V \rrbracket \{ \ell x \mapsto \llbracket M \rrbracket \bar{\textcircled{a}} \kappa; y \mapsto \llbracket N \rrbracket \bar{\textcircled{a}} \kappa \} \\
\llbracket \text{absurd } V \rrbracket &= \bar{\lambda} \kappa. \text{absurd } \llbracket V \rrbracket \\
\llbracket \text{return } V \rrbracket &= \bar{\lambda} \kappa. \text{app } \downarrow \kappa \llbracket V \rrbracket \\
\llbracket \text{let } x \leftarrow M \text{ in } N \rrbracket &= \bar{\lambda} \langle s, \chi \rangle \bar{\textcircled{a}} \kappa. \llbracket M \rrbracket \bar{\textcircled{a}} (\langle \lambda x k. \llbracket N \rrbracket \bar{\textcircled{a}} k \rangle \bar{\textcircled{a}} s, \chi \bar{\textcircled{a}} \kappa) \\
\llbracket \text{do } \ell V \rrbracket &= \bar{\lambda} \langle s, \langle v, e \rangle \rangle \bar{\textcircled{a}} \kappa. e \bar{\textcircled{a}} (\ell \llbracket V \rrbracket, \langle s, \langle v, e \rangle \rangle \bar{\textcircled{a}} \kappa) \bar{\textcircled{a}} \downarrow \kappa \\
\llbracket \text{handle}^\delta M \text{ with } H \rrbracket &= \bar{\lambda} \kappa. \llbracket M \rrbracket \bar{\textcircled{a}} \langle \rangle, \llbracket H \rrbracket^\delta \bar{\textcircled{a}} \kappa \\
\llbracket H \rrbracket^\delta &= \langle \llbracket H^{\text{ret}} \rrbracket, \llbracket H^{\text{ops}} \rrbracket^\delta \rangle \\
\llbracket \{ \text{return } x \mapsto N \} \rrbracket &= \lambda x k. \llbracket N \rrbracket \bar{\textcircled{a}} k \\
\llbracket \{ \langle \ell p r \mapsto N_\ell \rangle_{\ell \in \mathcal{L}} \} \rrbracket^\delta &= \lambda x k. \text{let } \langle z, \langle p, rk \rangle \rangle = x \text{ in} \\
&\quad \text{case } z \{ (\ell \mapsto \text{let } r = \text{res}^\delta rk \text{ in } \llbracket N_\ell \rrbracket \bar{\textcircled{a}} k)_{\ell \in \mathcal{L}} \} \\
&\quad y \mapsto M_{\text{forward}}((y, p, rk), k) \} \\
M_{\text{forward}}((y, p, rk), k) &= \text{let } \langle s', \langle v', e' \rangle \rangle \bar{\textcircled{a}} k' = k \text{ in} \\
&\quad e' \bar{\textcircled{a}} (y \langle p, \langle s', \langle v', e' \rangle \rangle \bar{\textcircled{a}} rk \rangle) \bar{\textcircled{a}} k'
\end{aligned}$$

Top-level program

$$\top \llbracket M \rrbracket = \llbracket M \rrbracket \bar{\textcircled{a}} (\langle \rangle, \langle \lambda x k. x, \lambda z k. \text{absurd } z \rangle \bar{\textcircled{a}} \kappa)$$

Fig. 8: Higher-Order Uncurried CPS Translation of $\lambda^\dagger$

5.2 Static Terms

Redexes marked as **static** are reduced as part of the translation (at compile time), whereas those marked as **dynamic** are reduced at runtime.

We make use of static lambda abstractions, pairs, and lists. We let κ range over static continuations and χ range over static handlers. We let $\mathcal{V}, \mathcal{W}$ range over meta language values, $\mathcal{M}$ range over meta language expressions, and $\mathcal{P}, \mathcal{Q}$ over meta language patterns. We use list and record pattern matching in the meta language.

$$\begin{aligned}
(\bar{\lambda} \langle \mathcal{P}, \mathcal{Q} \rangle. \mathcal{M}) \bar{\textcircled{a}} \langle \mathcal{V}, \mathcal{W} \rangle &= (\bar{\lambda} \mathcal{P}. \bar{\lambda} \mathcal{Q}. \mathcal{M}) \bar{\textcircled{a}} \mathcal{V} \bar{\textcircled{a}} \mathcal{W} = (\bar{\lambda} (\mathcal{P} \bar{\textcircled{a}} \mathcal{Q}). \mathcal{M}) \bar{\textcircled{a}} (\mathcal{V} \bar{\textcircled{a}} \mathcal{W}) \\
(\bar{\lambda} \langle \mathcal{P}, \mathcal{Q} \rangle. \mathcal{M}) \bar{\textcircled{a}} V &= \text{let } \langle f, s \rangle = V \text{ in } (\bar{\lambda} \mathcal{P}. \bar{\lambda} \mathcal{Q}. \mathcal{M}) \bar{\textcircled{a}} f \bar{\textcircled{a}} s = (\bar{\lambda} (\mathcal{P} \bar{\textcircled{a}} \mathcal{Q}). \mathcal{M}) \bar{\textcircled{a}} V
\end{aligned}$$

A meta language value V can be reified as a target language value $\downarrow V$.

$$\downarrow V = V \quad \downarrow (\mathcal{V} \bar{\textcircled{a}} \mathcal{W}) = \downarrow \mathcal{V} \bar{\textcircled{a}} \downarrow \mathcal{W} \quad \downarrow \langle \mathcal{V}, \mathcal{W} \rangle = \langle \downarrow \mathcal{V}, \downarrow \mathcal{W} \rangle$$

5.3 The Translation

The CPS translation is given in Fig. 8. Its behaviour on constructs for introducing and eliminating values is standard. Where necessary static continuations in

the meta language are reified as dynamic continuations in the target language. The translation of **return** V applies the continuation to $\llbracket V \rrbracket$. The translation of **let** $x \leftarrow M$ **in** N adds a frame to the pure continuation on the topmost frame of the continuation. The translation of **do** ℓ V dispatches the operation to the effect continuation at the head of the continuation. The resumption is initialised with the topmost frame of the continuation. The translations of deep and shallow handling each add a new frame to the continuation. The translation of the operation clauses of a handler dispatches on the operation. If a match is found then the reified resumption is turned into a function and made available in the body of the operation clause. If there is no match, then the operation is forwarded by unwinding the continuation, transferring the topmost frame to the head of the reified resumption before invoking the next effect continuation. The only difference between the translations of a deep handler and a shallow handler is that the reified resumption of the latter is specially marked in order to ensure that the handler is disposed of in the body of a matching operation clause.

Example. The following example illustrates how the higher-order CPS translation avoids generating administrative redexes by performing static reductions.

$$\begin{aligned} \top \llbracket \text{handle } (\text{do Await } \langle \rangle) \text{ with } H \rrbracket &= \llbracket \text{handle } (\text{do Await } \langle \rangle) \text{ with } H \rrbracket @ \mathcal{K}_\top \\ &= \llbracket \text{do Await } \langle \rangle \rrbracket @ \langle \rangle, \llbracket H \rrbracket :: \mathcal{K}_\top \\ &= \llbracket \text{do Await } \langle \rangle \rrbracket @ \langle \rangle, \langle \llbracket H^{\text{ret}} \rrbracket, \llbracket H^{\text{ops}} \rrbracket \rangle :: \mathcal{K}_\top \\ &= \llbracket H^{\text{ops}} \rrbracket @ \text{Await } \langle \rangle, \langle \rangle, \llbracket H \rrbracket :: \langle \rangle @ \downarrow \mathcal{K}_\top \end{aligned}$$

where $\mathcal{K}_\top = (\langle \rangle, \langle \lambda x k. x, \lambda z k. \text{absurd } z \rangle :: \langle \rangle)$. The resulting term passes **Await** directly to the dispatcher that implements the operation clauses of H .

5.4 Correctness

The translation naturally lifts to evaluation contexts.

$$\begin{aligned} \llbracket [] \rrbracket &= \bar{\lambda}_{\kappa. \kappa} \\ \llbracket \text{let } x \leftarrow \mathcal{E} \text{ in } N \rrbracket &= \bar{\lambda}_{\kappa} \langle s, \chi \rangle :: \kappa. \llbracket \mathcal{E} \rrbracket @ \langle \rangle (\langle \lambda x k. \llbracket N \rrbracket @ k \rangle :: s, \chi) :: \kappa \\ \llbracket \text{handle}^\delta \mathcal{E} \text{ with } H \rrbracket &= \bar{\lambda}_{\kappa} \llbracket \mathcal{E} \rrbracket @ \langle \rangle, \llbracket H \rrbracket^\delta :: \kappa \end{aligned}$$

Lemma 12 (Decomposition). $\llbracket \mathcal{E}[M] \rrbracket @ (\mathcal{V} :: \mathcal{W}) = \llbracket M \rrbracket @ (\llbracket \mathcal{E} \rrbracket @ (\mathcal{V} :: \mathcal{W}))$

Though it eliminates static administrative redexes, the translation still yields administrative redexes that cannot be eliminated statically, as they only appear at run-time, which arise from deconstructing a reified stack of continuations. We write $\rightsquigarrow_a$ for the compatible closure of U-SPLIT, U-CASE₁ and U-CASE₂.

The following lemma is central to our simulation theorem. It characterises the sense in which the translation respects the handling of operations.

Lemma 13 (Handling). *If $\ell \notin BL(\mathcal{E})$ and $H^\ell = \{\ell p r \mapsto N_\ell\}$ then:*

1. $\llbracket \text{do } \ell V \rrbracket @ (\llbracket \mathcal{E} \rrbracket @ \langle \rangle, \llbracket H \rrbracket :: \mathcal{W}) \rightsquigarrow^+ \rightsquigarrow_a^* (\llbracket N_\ell \rrbracket @ \mathcal{W})[\llbracket V \rrbracket / p, \lambda y k. \llbracket \text{return } y \rrbracket @ (\llbracket \mathcal{E} \rrbracket @ \langle \rangle, \llbracket H \rrbracket :: k)] / r$

$$\begin{aligned}
2. \quad & \llbracket \text{do } \ell \ V \rrbracket @ (\llbracket \mathcal{E} \rrbracket @ (\llbracket \cdot \rrbracket, \llbracket H \rrbracket^+ \rrbracket \rrbracket \mathcal{W})) \rightsquigarrow^+ \rightsquigarrow_a^* \\
& (\llbracket N_\ell \rrbracket @ \mathcal{W}) (\llbracket V \rrbracket / p, \lambda y \ k. \text{let } (\langle s, \langle v, e \rangle \rrbracket \rrbracket k) = k \text{ in} \\
& \quad \llbracket \text{return } y \rrbracket @ (\llbracket \mathcal{E} \rrbracket @ (\llbracket \cdot \rrbracket, \langle s, \langle v, e \rangle \rrbracket \rrbracket k)) / r]
\end{aligned}$$

We now give a simulation result in the style of Plotkin [19]. The theorem shows that the only extra behaviour exhibited by a translated term is the necessary bureaucracy of dynamically deconstructing the continuation stack.

Theorem 14 (Simulation). *If $M \rightsquigarrow N$ then for all static values $\mathcal{V}$ and $\mathcal{W}$, we have $\llbracket M \rrbracket @ (\mathcal{V} \rrbracket \mathcal{W}) \rightsquigarrow^+ \rightsquigarrow_a^* \llbracket N \rrbracket @ (\mathcal{V} \rrbracket \mathcal{W})$.*

Proof. By induction on the reduction relation ($\rightsquigarrow$) using Lemma 13.

As a corollary, we obtain that the translation simulates full reduction to a value.

Corollary 15. $M \rightsquigarrow^* V$ iff $\top \llbracket M \rrbracket \rightsquigarrow^* \rightsquigarrow_a^* \top \llbracket V \rrbracket$.

6 Empirical Evaluation

We conducted a basic empirical evaluation using an experimental branch of the Links web programming language [2] extended with support for shallow handlers and JavaScript backends based on the CEK machine (§4) and CPS translation (§5). We omit the full details due to lack of space; they can be found in the extended version of the paper [10, Appendix B]. Here we give a brief high-level summary. Our benchmarks are adapted from Kammar et al. [12], comprising: pipes, a count down loop, and n -Queens. Broadly, our results align with those of Kammar et al. Specifically, the shallow implementation of pipes outperforms the deep implementation. The shallow-as-deep translation fails to complete most benchmarks as it runs out of memory. The memory usage pattern exhibited by deep, shallow, and shallow-as-deep implementations are all stable.

Deep handlers perform slightly better than shallow handlers except on the pipes benchmark (CEK and CPS) and the countdown benchmark on the CEK machine. The former is hardly surprising given the inherent indirection in the deep implementation of pipes, which causes unnecessary closure allocations to happen when sending values from one end of the pipe to the other. We conjecture that the relatively poor performance of deep handlers on the CEK version of the countdown benchmark is also due to unnecessary closure allocation in the interpretation of state. Kammar et al. avoid this problem by adopting *parameterised handlers*, which thread a parameter through each handler.

7 Related Work

Shallow Handlers. Most existing accounts of effect handlers use deep handlers. Notable exceptions include Haskell libraries based on free monads [12,13,14], and the Frank programming language [17]. Kiselyov and Ishii [13] optimise their implementation by allowing efficient implementations of catenable lists to be used to support manipulation of continuations. We conjecture that both our abstract machine and our CPS translation could benefit from a similar representation.

Abstract Machines for Handlers. Lindley et al. [17] implement Frank using an abstract machine similar to the one described in this paper. Their abstract machine is not formalised and differs in several ways. In particular, continuations are represented by a single flattened stack, rather than a nested stack like ours, and Frank supports multihandlers, which handle several computations at once. Biernacki et al. [1] present an abstract machine for deep effect handlers similar to that of Hillerström and Lindley [9] but factored slightly differently.

CPS for Handlers. Leijen [15] implements a selective CPS translation for deep handlers, but does not go all the way to plain lambda calculus, relying on a special built in handling construct.

8 Conclusion and Future Work

We have presented the first comprehensive formal analysis of shallow effect handlers. We introduced the handler calculus $\lambda^\dagger$ as a uniform calculus of deep and shallow handlers. We specified formal translations back and forth between deep and shallow handlers within $\lambda^\dagger$, an abstract machine for $\lambda^\dagger$, and a higher-order CPS translation for $\lambda^\dagger$. In each case we proved a precise simulation result, drawing variously on different notions of administrative reduction. We have implemented the abstract machine and CPS translation as backends for Links and evaluated the performance of deep and shallow handlers and their encodings, measuring both execution time and memory consumption. Though deep and shallow handlers can always encode one another, the results suggest that the shallow-as-deep encoding is not viable in practice due to administrative overhead, whereas the deep-as-shallow encoding may be viable. In future we intend to perform a more comprehensive performance evaluation for a wider range of effect handler implementations.

Another outstanding question is to what extent shallow handlers are really needed at all. We have shown that we can encode them generically using deep handlers, but the resulting cruft hinders performance in practice. Extensions to deep handlers not explored in this paper, such as *parameterised handlers* [12,21] or a deep version of the *multihandlers* of Lindley et al. [17], offer the potential for expressing certain shallow handlers without the cruft. Parameterised handlers thread a parameter through each handler, avoiding unnecessary closure allocation. Deep multihandlers directly capture mutumorphisms over computations, allowing a direct implementation of pipes. In future we plan to study the precise relationship between shallow handlers, parameterised handlers, deep multihandlers, and perhaps handlers based on other structural recursion schemes.

Acknowledgements. We would like to thank John Longley for insightful discussions about the inter-encodings of deep and shallow handlers. Daniel Hillerström was supported by EPSRC grant EP/L01503X/1 (EPSRC Centre for Doctoral Training in Pervasive Parallelism). Sam Lindley was supported by EPSRC grant EP/K034413/1 (From Data Types to Session Types—A Basis for Concurrency and Distribution).

References

1. Biernacki, D., Piróg, M., Polesiuk, P., Sieczkowski, F.: Handle with care: relational interpretation of algebraic effects and handlers. *PACMPL* **2**(POPL), 8:1–8:30 (2018)
2. Cooper, E., Lindley, S., Wadler, P., Yallop, J.: Links: Web programming without tiers. In: *FMCO. LNCS*, vol. 4709, pp. 266–296. Springer (2006)
3. Danvy, O., Nielsen, L.R.: A first-order one-pass CPS transformation. *Theor. Comput. Sci.* **308**(1-3), 239–257 (2003)
4. Dolan, S., White, L., Sivaramakrishnan, K., Yallop, J., Madhavapeddy, A.: Effective concurrency through algebraic effects. *OCaml Workshop* (2015)
5. Felleisen, M., Friedman, D.P.: Control operators, the SECD-machine, and the λ -calculus. In: *Formal Description of Programming Concepts III*. pp. 193–217 (1987)
6. Flanagan, C., Sabry, A., Duba, B.F., Felleisen, M.: The essence of compiling with continuations. In: *PLDI*. pp. 237–247. ACM (1993)
7. Fokkinga, M.M.: Tupling and mutomorphisms. *The Squiggolist* **1**(4), 81–82 (1990)
8. Goodman, N.: Uber AI Labs open sources Pyro, a deep probabilistic programming language (Nov 2017), <https://eng.uber.com/pyro/>
9. Hillerström, D., Lindley, S.: Liberating effects with rows and handlers. In: *TyDe@ICFP*. pp. 15–27. ACM (2016)
10. Hillerström, D., Lindley, S.: Shallow effect handlers (extended version) (2018), <http://homepages.inf.ed.ac.uk/slindley/papers/shallow-extended.pdf>
11. Hillerström, D., Lindley, S., Atkey, R., Sivaramakrishnan, K.C.: Continuation passing style for effect handlers. In: *FSCD. LIPIcs*, vol. 84, pp. 18:1–18:19 (2017)
12. Kammrath, O., Lindley, S., Oury, N.: Handlers in action. In: *ICFP*. pp. 145–158. ACM (2013)
13. Kiselyov, O., Ishii, H.: Freer monads, more extensible effects. In: *Haskell*. pp. 94–105. ACM (2015)
14. Kiselyov, O., Sabry, A., Swords, C.: Extensible effects: an alternative to monad transformers. In: *Haskell*. pp. 59–70. ACM (2013)
15. Leijen, D.: Type directed compilation of row-typed algebraic effects. In: *POPL*. pp. 486–499. ACM (2017)
16. Levy, P.B., Power, J., Thielecke, H.: Modelling environments in call-by-value programming languages. *Inf. Comput.* **185**(2), 182–210 (2003)
17. Lindley, S., McBride, C., McLaughlin, C.: Do be do be do. In: *POPL*. pp. 500–514. ACM (2017)
18. Meijer, E., Fokkinga, M.M., Paterson, R.: Functional programming with bananas, lenses, envelopes and barbed wire. In: *FPCA. Lecture Notes in Computer Science*, vol. 523, pp. 124–144. Springer (1991)
19. Plotkin, G.D.: Call-by-name, call-by-value and the lambda-calculus. *Theor. Comput. Sci.* **1**(2), 125–159 (1975)
20. Plotkin, G.D., Power, J.: Adequacy for algebraic effects. In: *FoSSaCS. LNCS*, vol. 2030, pp. 1–24. Springer (2001)
21. Plotkin, G.D., Pretnar, M.: Handling algebraic effects. *Logical Methods in Computer Science* **9**(4) (2013)
22. Reynolds, J.C.: Definitional interpreters for higher-order programming languages. *Higher-Order and Symbolic Computation* **11**(4), 363–397 (1998)
23. Wu, N., Schrijvers, T.: Fusion for free - efficient algebraic effect handlers. In: *MPC. Lecture Notes in Computer Science*, vol. 9129, pp. 302–322. Springer (2015)
24. Yallop, J.: Staged generic programming. *PACMPL* **1**(ICFP), 29:1–29:29 (2017)