



Heriot-Watt University  
Research Gateway

# A General Fine-Grained Reduction Theory for Effect Handlers

**Citation for published version:**

Sieczkowski, F, Pyzik, M & Biernacki, D 2023, 'A General Fine-Grained Reduction Theory for Effect Handlers', *Proceedings of the ACM on Programming Languages*, vol. 7, no. ICFP, 206, pp. 511-540. <https://doi.org/10.1145/3607848>

**Digital Object Identifier (DOI):**

[10.1145/3607848](https://doi.org/10.1145/3607848)

**Link:**

[Link to publication record in Heriot-Watt Research Portal](#)

**Document Version:**

Publisher's PDF, also known as Version of record

**Published In:**

Proceedings of the ACM on Programming Languages

**Publisher Rights Statement:**

©2023 the owner/author(s).

**General rights**

Copyright for the publications made accessible via Heriot-Watt Research Portal 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**

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [open.access@hw.ac.uk](mailto:open.access@hw.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.



# A General Fine-Grained Reduction Theory for Effect Handlers

FILIP SIECZKOWSKI, Heriot-Watt University, United Kingdom

MATEUSZ PYZIK, University of Wrocław, Poland

DARIUSZ BIERNACKI, University of Wrocław, Poland

Effect handlers are a modern and increasingly popular approach to structuring computational effects in functional programming languages. However, while their traditional operational semantics is well-suited to implementation tasks, it is less ideal as a reduction theory. We therefore introduce a fine-grained reduction theory for deep effect handlers, inspired by our existing reduction theory for `shift0`, along with a standard reduction strategy. We relate this strategy to the traditional, non-local operational semantics via a simulation argument, and show that the reduction theory preserves observational equivalence with respect to the classical semantics of handlers, thus allowing its use as a rewriting theory for handler-equipped programming languages – this rewriting system mostly coincides with previously studied type-based optimisations. In the process, we establish theoretical properties of our reduction theory, including confluence and standardisation theorems, adapting and extending existing techniques. Finally, we demonstrate the utility of our semantics by providing the first normalisation-by-evaluation algorithm for effect handlers, and prove its soundness and completeness. Additionally, we establish non-expressibility of the lift operator, found in some effect-handler calculi, by the other constructs.

CCS Concepts: • **Theory of computation** → **Operational semantics; Control primitives.**

Additional Key Words and Phrases: algebraic effect, delimited continuation, reduction, normalization

## ACM Reference Format:

Filip Sieczkowski, Mateusz Pyzik, and Dariusz Biernacki. 2023. A General Fine-Grained Reduction Theory for Effect Handlers. *Proc. ACM Program. Lang.* 7, ICFP, Article 206 (August 2023), 30 pages. <https://doi.org/10.1145/3607848>

## 1 INTRODUCTION

Effect handlers provide a modern, modular approach to integrating effectful computation in functional programming languages. Based on the pioneering work of Plotkin and Power [2004] on algebraic effects, and Plotkin and Pretnar’s idea of reifying algebras as *handlers* [Plotkin and Pretnar 2013], the technique has seen considerable rise in popularity, with several experimental languages developed [Bauer and Pretnar 2015; Biernacki et al. 2019; Brachthäuser et al. 2020; Convent et al. 2020; Leijen 2017a; Zhang and Myers 2019] and integration within industrial functional languages under way.<sup>1</sup> However, despite its highly semantic and theoretical background, certain aspects of the approach, particularly as they relate to operational semantics, remain somewhat understudied.

<sup>1</sup>In December 2022, OCaml 5.0 introduced effect handlers as an experimental feature, c.f. <https://v2.ocaml.org/releases/5.0/manual/effects.html>.

Authors’ addresses: Filip Sieczkowski, [f.sieczkowski@hw.ac.uk](mailto:f.sieczkowski@hw.ac.uk), Heriot-Watt University, School of Mathematical & Computer Sciences, Lochside Walk, Currie, Edinburgh, United Kingdom, EH14 4AS; Mateusz Pyzik, [mateusz.pyzik@cs.uni.wroc.pl](mailto:mateusz.pyzik@cs.uni.wroc.pl), Institute of Computer Science, University of Wrocław, Joliot-Curie 15, Wrocław, Poland, 50-383; Dariusz Biernacki, [dabi@cs.uni.wroc.pl](mailto:dabi@cs.uni.wroc.pl), Institute of Computer Science, University of Wrocław, Joliot-Curie 15, Wrocław, Poland, 50-383.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2023 Copyright held by the owner/author(s).

2475-1421/2023/8-ART206

<https://doi.org/10.1145/3607848>

In this paper, we attempt to reduce this gap by developing a reduction theory for effect handlers, together with some applications.

*Motivating example.* Consider the following program fragment, written in a pseudocode that could be interpreted in most programming languages that provide effect handlers:

```

handle
  iter (fn x -> let y = ask () in tell (x + y)) [1, 2, 3]
with
  | ask () -> resume 5
  | return () -> ()

```

In the above, we assume a standard `iter` function that applies its first argument to the elements of the list in turn (for their side effects), and two *operations*: `ask` and `tell`. These have no interpretation on their own – their behaviour is determined by *handlers* that can be found, dynamically, in their evaluation context. The handler for the `ask` operation is provided: its intended semantics is to resume computation in the place where the operation was encountered with a constant value of 5. The other operation, on the other hand, is left uninterpreted by this fragment of the code – it may well be unknown at this point if, for instance, the snippet is the body of a function.

To a knowledgeable observer, the intent of this toy example is clear: since the provided handler for `ask` is benign and the environment can never insert a different handler between the operations and the provided one, the program fragment should be equivalent to the following, simpler one:

```
tell 6; tell 7; tell 8
```

Indeed, in an appropriate system (with a static type-and-effect system) we could use known reasoning techniques, such as the logical relation of Biernacki et al. [2018], to establish that the two are observationally equivalent.<sup>2</sup> However, this would be an (admittedly trivial) instance of a sophisticated reasoning technique that is not necessarily easy to automate. Maybe a compiler could transform the first snippet into the second one as an optimisation based simply on specialisation of program fragments?

To consider this question, we must first recall how operational semantics of effect handlers is usually defined. The gist of the crucial reduction rule is that when we encounter an operation, the rule looks through the evaluation context until a matching handler is found, and reifies the intervening part as a *resumption*. This resumption, together with the argument of the operation are then passed to the body of the handler. In the case of our program fragment, the evaluation would look as follows: first, the standard definition of `iter` is unfolded, giving rise to the following snippet:

```

handle
  let y = ask () in tell (1 + y);
  iter (fn x -> let y = ask () in tell (x + y)) [2, 3]
with
  | ask () -> resume 5
  | return () -> ()

```

Now, the operation is in the evaluation position, and its matching handler can be found. Since the argument of the operation is trivial, the only interesting aspect is the resumption that gets captured and passed to the handler, giving us the following result:

```

(fn z ->
  handle
    let y = z in tell (1 + y);

```

<sup>2</sup>In fact, it would be easy to prove a more general result that would hold for any input list.

```

    iter (fn x -> let y = ask () in tell (x + y)) [2, 3]
  with
  | ask () -> resume 5
  | return () 5

```

Note that the `handle` expression is retained within the resumption: this is a characteristic feature of *deep* effect handlers, with which we will concern ourselves in this development.

We can now reduce the application of the resumption to its argument, and proceed with evaluation to get the following snippet:

```

handle
  tell 6; iter (fn x -> let y = ask () in tell (x + y)) [2, 3]
with
| ask () -> resume 5
| return () -> ()

```

At this point, however, the operation `tell` is in the evaluation position, and we have no matching handler. No other reduction (save for unfolding `iter` on the tail of the list) can apply, and thus we cannot reach the simpler result.

This example, while clearly artificial, exposes inherent limitations that the standard operational semantics of deep effect handlers imposes on program optimisation through specialisation. The presence of an uninterpreted operation (or, indeed, an uninterpreted variable in the function position of an open program fragment) “freezes” any handlers in its evaluation context, and prevents us from reducing them, thus limiting the possibilities of compiling out known effects into more efficient, pure code. Can we avoid these problems and obtain the simpler program, without any trace of the `ask` operation or its associated handler?

*Fine-grained reduction.* It turns out that if we choose an appropriate reduction theory, such optimisation becomes possible. In fact, one set of rewrite rules that can be used as to optimise programs has been proposed by Saleh [2019] and taken up by Karachalias et al. [2021]. While their focus is on optimising the general form of calls to recursive, effectful functions wrapped in an effect handler — in the context of our motivating example, calls of the shape `handle iter ... xs with ...` where `xs` may be a general expression — this task requires them to consider more local rules for program simplification as well.<sup>3</sup> Independently, Biernacki et al. [2021] proposed a similar set of rules as a *bona fide* reduction theory in the context of delimited-control operators. The core idea behind both sets of rules is to define the reduction of the delimiter (respectively, `$` or `handle` construct for the two types of delimited control) through its interactions with other computation formers, notably the `let` expressions. In this approach, rather than form the resumption only when an operation is encountered, a handler would “distribute” through the `let`-expression, appropriately adjusting its `return` clause (which determines the result if the computation within the handler terminates with a value). In our example, the reduction (after unfolding the `iter` function) could proceed as follows:

```

handle
  let y = ask () in tell (1 + y)
with
| ask () -> resume 5
| return () ->
  handle

```

<sup>3</sup>The general case they tackle is a proper compiler optimisation that cannot be naturally viewed as part of the reduction theory, and we do not aim to tackle it in the present paper.

```

    iter (fn x -> let y = ask () in tell (x + y)) [2, 3]
  with
  | ask () -> resume 5
  | return () -> ()

```

Note how the handler gets duplicated as the iteration over the tail of the list gets incorporated into the return clause. This may seem worrisome from the point of view of evaluation (and we address these concerns in Section 2); however, it also unlocks the possibility of further reduction, which was prohibited by the traditional approach. Following on, we get:

```

  handle ask ()
  with
  | ask () -> resume 5
  | return y ->
    handle tell (1 + y)
    with
    | ask () -> resume 5
    | return () ->
      handle iter (fn x -> let y = ask () in tell (x + y)) [2, 3]
      with
      | ask () -> resume 5
      | return () -> ()

```

Now, since the handler matches the operation, we can continue — but there is no need to construct the resumption, since there can never be any intervening context. Thus, instead of a “resumption” we simply pass the argument and return clause to the body of the handler, obtaining:

```

  handle tell (1 + 5)
  with
  | ask () -> resume 5
  | return () ->
    handle iter (fn x -> let y = ask () in tell (x + y)) [2, 3]
    with
    | ask () -> resume 5
    | return () -> ()

```

This time, the operation within the handler does *not* match the handled operation. Again, there is no need to capture the handler, even though the evaluation of `tell 6` cannot proceed. Instead, we simply sequence the computation of the return clause after the control-stuck operation, obtaining:

```

  tell 6;
  handle iter (fn x -> let y = ask () in tell (x + y)) [2, 3]
  with
  | ask () -> resume 5
  | return () -> ()

```

Now we can continue the reduction process for the tail of the list, which leads to removing all instances of `ask` (and its associated handler) altogether. In this paper, we introduce and study a calculus that makes optimisations like the one described above possible.

*Structure of the paper and contributions.* The remainder of the paper is structured as follows. In Section 2 we develop a core calculus of deep effect handlers, in the style of Piróg et al. [2019]. We equip it with two sets of semantics: a traditional one, based on capturing parts of the evaluation

term	$e, a$	$::=$	$\text{return } v \mid \text{let } x = a \text{ in } e \mid u v \mid \text{do } v \mid \text{handle } e \text{ with } h, r \mid \text{lift } e$
value	$u, v, w, h, r$	$::=$	$x \mid \lambda x . e$
	$(\beta.v)$	$(\lambda x . e) v$	$\mapsto_c e[v/x]$
	$(\text{DH}.v)$	handle return $v$ with $h, r$	$\mapsto_c r v$
	$(\beta.\text{let})$	let $x = \text{return } v$ in $e$	$\mapsto_t e[v/x]$
	$(\text{lift}.v)$	lift return $v$	$\mapsto_t \text{return } v$

Fig. 1. The common syntax and contractions shared by the context-capturing and the fine-grained semantics.

evaluation context	$K ::=$	$[\ ] \mid \text{let } x = K \text{ in } e \mid \text{lift } K \mid \text{handle } K \text{ with } h, r$
$n$ -free context	$J_n ::=$	$(_{n=0}) [\ ] \mid \text{let } x = J_n \text{ in } e \mid (_{n=m+1}) \text{lift } J_m \mid \text{handle } J_{n+1} \text{ with } h, r$
$(\text{DH}.j)$	handle $J_0[\text{do } v]$ with $h, r$	$\mapsto_n \text{let } f = h v \text{ in } f \lambda x . \text{handle } J_0[\text{return } x] \text{ with } h, r$

$$\frac{a \mapsto_c e}{K[a] \mapsto_{\lambda_{n\text{DH}}} K[e]} \quad \frac{a \mapsto_n e}{K[a] \mapsto_{\lambda_{n\text{DH}}} K[e]} \quad \frac{a \mapsto_t e}{K[a] \mapsto_{\lambda_{n\text{DH}}} K[e]}$$

Fig. 2. The traditional, context-capturing operational semantics of  $\lambda_{n\text{DH}}$ .

context as resumptions, and a fine-grained one, which draws on the optimisation rules of Saleh [2019] and the reduction theory of Biernacki et al. [2021]. We also establish a non-expressibility result for the *lift* construct, which was a feature of some of the systems to date, but whose precise expressive power was not formally studied before. In Section 3 we establish bisimilarity between the traditional semantics and the fine-grained rules, treated as evaluation theories. Thus, we provide a connection between the two sets of semantics of the deep handlers. In Sections 4 and 5 we treat the fine-grained contraction rules as a basis of a *reduction* theory, and establish confluence and standardisation results for our calculus (with the weak-head part of the standard reduction corresponding to the evaluation of the previous section). In this, we follow classic approaches from lambda calculus, and extend them to our setup to account for new challenges posed by the effect handlers. These results allow us to establish, in Section 6, fine-grained equational reasoning as a sound tool to reason about observational equivalence with respect to the traditional semantics, formally justifying Saleh’s basic optimisations without recourse to types. Finally, in Section 7 we develop the first normalization-by-evaluation procedure for effect handlers, based on the fine-grained reduction, and prove its soundness and completeness, and we conclude in Section 9. With the exception of the expressibility of lift in Section 2.2 all results presented in this paper are formalised in the Coq interactive theorem prover.

## 2 SYNTAX AND SEMANTICS

In this section, we introduce two calculi of algebraic effects that share the same syntax but are equipped with different operational semantics. The syntax and the common reduction rules are shown in Figure 1. The syntax of the calculi is minimal – it is a fine-grain call-by-value  $\lambda$ -calculus [Levy et al. 2003], extended with algebraic-effect specific constructs – a common setup, used for instance by Hillerström et al. [2020] or Piróg et al. [2019].

More precisely, the syntax is stratified into two categories: values and terms (expressions). Values comprise variables and  $\lambda$ -abstractions. Note that we do not need to handle recursive functions explicitly, as the calculi are untyped, and thus fixed-point combinators are easily expressible. Terms represent possibly effectful computations. A term may return a value, chain computations sequentially with a let-expression, apply a function to an argument (both need to be values), invoke an effect, handle an effectful computation (with  $h$  and  $r$  representing an effect handler and a return clause, respectively), or *lift* a computation. Contrary to the syntax used in the example in Section 1, the algebraic operations or effects are unnamed; we use the lift operator to allow a single expression to use multiple distinct effects. This is achieved through the lifted computation “skipping” the nearest enclosing handler, which allows us to encode examples such as the one presented in the introduction.

*Context-capturing calculus.* The first calculus we consider is called  $\lambda_{n\text{DH}}$ , and it is an untyped version of a calculus of deep handlers considered by Piróg et al. [2019]. The operational semantics of  $\lambda_{n\text{DH}}$  is defined by the reduction rules in Figures 1 and 2 – the contraction relation  $\mapsto$  is split into relations  $\mapsto_c$ ,  $\mapsto_t$  and  $\mapsto_n$  to simplify formulation of some of the technical lemmas in Section 3. The rules  $(\beta.v)$ ,  $(\text{DH}.v)$ ,  $(\beta.\text{let})$  and  $(\text{lift}.v)$  are standard, and define the interaction of functions, handlers, let-expressions and the lift operator with values ( $e[v/x]$  stands for the usual capture-avoiding substitution). In particular, handling a value (rule  $(\text{DH}.v)$ ) consists in passing it to the return value (no effects need to be handled), whereas lifting a value (rule  $(\text{lift}.v)$ ) is an identity operation (no effects need to be lifted).

The most instrumental reduction rule of the calculus is  $(\text{DH}.j)$  and it says that handling an effect is done by capturing an appropriately ‘balanced’ evaluation context in one step. In order to formalize the notion of a balanced context, we index evaluation contexts  $J$  by its freeness level, i.e., the number of handlers surrounding the context hole that the context would need to neutralize the lift operators that surround the context hole to skip handlers. We can see that in the grammar of  $J_n$ , lift increases the index, whereas handle decreases it. Then, balanced contexts are exactly 0-free contexts  $J_0$ , and the handler surrounding such a context is the one that matches an effect invocation within the context. The rule  $(\text{DH}.j)$  further specifies that the handler  $h$  is given the value passed to the effect invocation and a captured resumption (delimited continuation), represented by a 0-free context surrounded by the same handle expression, which is characteristic of deep handlers (as opposed to shallow handlers that capture a handler-free resumption [Hillerström et al. 2020]).

The evaluation relation  $\mapsto_{\lambda_{n\text{DH}}}$  is then defined as a contextual closure of  $\mapsto$  with respect to the entire set of evaluation contexts  $K$  that subsume the  $n$ -free contexts  $J_n$  (as usual,  $K[e]$  stands for  $K$  plugged with  $e$ ). It can be shown that  $\mapsto_{\lambda_{n\text{DH}}}$  is deterministic, and that for a closed term  $e$ , either  $e$  reduces infinitely, or it evaluates to return  $v$ , or it reaches a control-stuck term  $J_n[\text{do } v]$ . Moreover,  $\mapsto_{\lambda_{n\text{DH}}}$  is compatible with respect to evaluation contexts, i.e.,  $K[e] \mapsto_{\lambda_{n\text{DH}}} K[e']$  whenever  $e \mapsto_{\lambda_{n\text{DH}}} e'$ .

*Fine-grained calculus.* The second calculus,  $\lambda_{fg\text{DH}}$ , builds on the twin lineage of Saleh’s optimization rules and Biernacki et al.’s reduction theory for delimited control to avoid explicit capture of contexts, and facilitate optimization. It extends the contraction rules of Figure 1, with the three additional rules shown in Figure 3.

Reduction  $(\text{DH}.do)$  describes a direct interaction of an effect invocation with its surrounding handler, when there is no intermediate computation separating them – the effect gets handled, so handle  $\text{do } v$  with  $h, r$  can be rewritten as  $h v r$  – which, due to the highly stylized presentation of our calculus, we take as syntactic sugar for a let-expression  $\text{let } f = h v \text{ in } f r$ . Note that  $r$  is passed to  $h$  as the second argument, which means that at this stage it represents the resumption. Reduction  $(\text{DH}.lift)$  describes a direct interaction of a lift expression with its surrounding handler – the

$$\begin{array}{l}
\text{evaluation context } E_0 ::= [] \mid E_0[\text{let } x = [] \text{ in } e] \mid E_0[\text{lift } []] \\
E_{n+1} ::= E_n[\text{handle } [] \text{ with } h, r] \\
\\
(\text{DH.}do) \quad \text{handle do } v \text{ with } h, r \quad \mapsto_f \quad \text{let } f = h v \text{ in } f r \\
(\text{DH.lift}) \quad \text{handle lift } e \text{ with } h, r \quad \mapsto_f \quad \text{let } x = e \text{ in } r x \\
(\text{DH.let}) \quad \text{handle let } x = a \text{ in } e \text{ with } h, r \quad \mapsto_f \quad \text{handle } a \text{ with } h, \lambda x. \text{ handle } e \text{ with } h, r \\
\\
\frac{a \mapsto_c e}{E_n[a] \mapsto_{\lambda_{fg}DH} E_n[e]} \quad \frac{a \mapsto_f e}{E_n[a] \mapsto_{\lambda_{fg}DH} E_n[e]} \quad \frac{a \mapsto_t e}{E_0[a] \mapsto_{\lambda_{fg}DH} E_0[e]}
\end{array}$$

Fig. 3. The fine-grained operational semantics of  $\lambda_{fg}DH$ .

handler  $h$  is skipped and the computation  $e$  is passed to the return clause  $r$  of the handler. The most complex contraction of the fine-grained calculus, and the one that enforces its *deep* treatment of handlers, is **(DH.let)**. It states that a handler can be distributed to the subcomputations of a let-expression  $\text{let } x = a \text{ in } e$ , by appropriately updating the return clause  $r$  of the handler with  $e$  guarded by a copy of the handler the reduction started with. Effectively, each use of **(DH.let)** or **(DH.lift)** rule closes the distance between an effect call and its handler until they are facing each other directly, and can be eliminated by the **(DH.do)** rule.

The reduction rules we propose resemble other reduction theories for calculi of delimited control that are based on piece-by-piece consumption of evaluation contexts [Felleisen 1988; Pretnar 2015]. In these works, however, the control operator or the operation aggregates the elementary contexts on the way to its surrounding delimiter or handler. Our calculus, on the other hand, consumes the context in the opposite direction, i.e. from the handler to the operation, and we take advantage of the return clause of the handler to store the elementary contexts guarded by (possibly superfluous) handlers.<sup>4</sup> Both **(DH.do)** and **(DH.let)** correspond directly to rewrite rules given by Saleh [2019] (respectively, rules WITH-HANDLED-OP and WITH-DO in Figure 3.2); **(DH.lift)** plays the same role as his WITH-PURE, but in a purely syntax-directed, untyped setting, which makes it more applicable as a basis for a reduction theory.

The grammar of evaluation contexts  $E_n$  again takes advantage of an index  $n$ , but this time its role is simpler than in the nonlocal calculus –  $n$  is the number of handlers surrounding the context hole. As a matter of fact, we only need to keep track of the fact whether the computation takes place under a handler or not, and, therefore, a boolean flag would suffice. Note that, unlike in the nonlocal calculus, contexts are defined inside-out, which for technical reasons, e.g. in Section 3, is more convenient.

The evaluation relation  $\mapsto_{\lambda_{fg}DH}$  is then defined as a compatible closure of the relations  $\mapsto_c$ ,  $\mapsto_f$ , and  $\mapsto_t$  with respect to evaluation contexts. Notice, however, that the ‘top-level’ reductions  $\mapsto_t$  can only occur in a handle-free context  $E_0$ , since the semantics of let- and lift- expressions depends on the presence of a surrounding handler. Therefore, we have defined a so-called hybrid evaluation strategy [Biernacka et al. 2017; García-Pérez and Nogueira 2014], where the semantics of some constructs depends on the kind of the surrounding context (our contexts are kinded by the indices  $n$  in this sense). Hybrid strategies typically are used in the semantic specification of full normalization

<sup>4</sup>It is worth noting that such architecture of the reduction system immediately rules out a possibility of accounting for shallow handlers whose semantics assumes that the handler is never copied to the resumption, whereas the rule **(DH.let)** eagerly does exactly that.



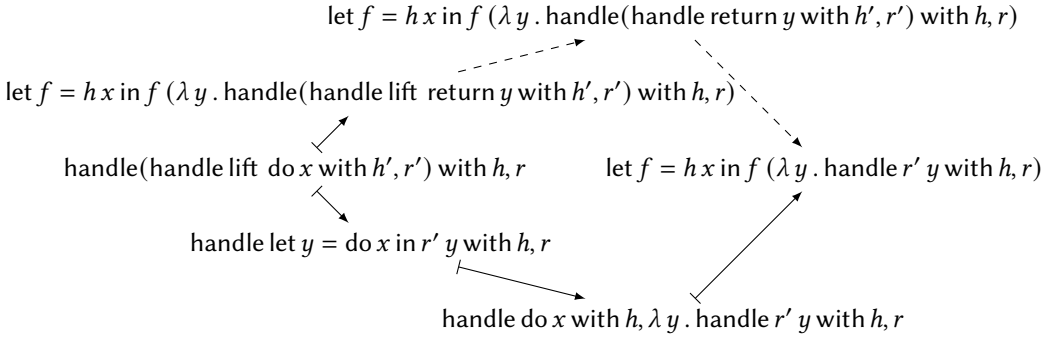


Fig. 4. Comparison of evaluation in the context-capturing and fine-grained semantics. Fine-grained semantics needs three evaluation steps to arrive at a handler call (lower path), while a single nonlocal step is sufficient (upper path). There is a mismatch between the captured resumptions that can be mitigated by two general reduction steps.

in  $\lambda$ -calculi. As for the traditional semantics, it can be shown that  $\mapsto_{\lambda_{fgDH}}$  is deterministic and that the normal forms it computes for closed expressions are returned values and control-stuck terms  $E_0[do v]$ .

The semantics duplicating handlers that we have defined may not be the first choice for an implementation of an evaluator for a calculus of algebraic effects. It is, however, an interesting alternative to the standard semantics that is more powerful when it comes to program rewriting and optimization, as illustrated with the motivating example in Section 1, and as a foundation for a fully-normalizing procedure (developed in Section 7). In order to talk about reduction viewed as code optimization, we introduce a relation  $\rightarrow_{\lambda_{fgDH}}$  that is a compatible closure of the contraction rules  $\mapsto_t$ ,  $\mapsto_c$ , and  $\mapsto_f$  with general term contexts, i.e., it allows for reduction at any position in a term (in Section 4 we show that  $\rightarrow_{\lambda_{fgDH}}$  is confluent). The following example demonstrates both evaluation and general reduction relations in action (such simplification could not be obtained using the traditional reduction rules):

$$\begin{array}{l}
 \text{handle let } y = p x \text{ in do } y \text{ with } h, r \\
 \mapsto_{\lambda_{fgDH}} \text{handle } p x \text{ with } h, \lambda y . \text{handle do } y \text{ with } h, r \\
 \rightarrow_{\lambda_{fgDH}} \text{handle } p x \text{ with } h, \lambda y . \text{let } f = h y \text{ in } f r
 \end{array}$$

(In Sections 5 and 6, we study the formal relation between  $\rightarrow_{\lambda_{fgDH}}$  and  $\mapsto_{\lambda_{fgDH}}$ .)

In the following subsection we present an additional view on the evaluation relation in the fine-grained calculus, offered by a syntactic abstract machine, whereas in Section 3 we prove that the standard and the fine-grained semantics are equivalent as far as evaluation is concerned. The abstract machine and the simulation theorems shed light on the inner workings of the new semantics and explain the architecture of the new calculus in greater detail. For now, it is instructive to compare how the context-capturing and the fine-grained semantics capture the context surrounding an effect call, which is illustrated in Figure 4.

## 2.1 A Syntactic Abstract Machine

In this section we present a *syntactic* abstract machine for the fine-grained calculus that is constructed in the spirit of the CK machine [Felleisen et al. 2009]. The goal here is to show how in the new semantics the expressions are decomposed and what the actual structure of evaluation

stack of handlers	$H ::=$	$\square \mid H[\text{handle } \square \text{ with } h, r]$
top continuation	$T ::=$	$\square \mid T[\text{let } x = \square \text{ in } e] \mid T[\text{lift } \square]$
intermediate value	$i ::=$	$\text{return } v \mid \text{do } v \mid \text{let } x = a \text{ in } e \mid \text{lift } e$
$\langle \text{handle } e \text{ with } h, r \mid H \mid T \rangle_e \triangleright \langle e \mid H[\text{handle } \square \text{ with } h, r] \mid T \rangle_e$ $\langle (\lambda x. e) v \mid H \mid T \rangle_e \triangleright \langle e[v/x] \mid H \mid T \rangle_e$ $\langle i \mid H \mid T \rangle_e \triangleright \langle H \mid i \mid T \rangle_c$ $\langle H[\text{handle } \square \text{ with } h, r] \mid i \mid T \rangle_c \triangleright \langle i \mid H[\text{handle } \square \text{ with } h, r] \mid T \rangle_h$ $\langle \square \mid i \mid T \rangle_c \triangleright \langle i \mid T \rangle_i$ $\langle \text{let } x = a \text{ in } e \mid H[\text{handle } \square \text{ with } h, r] \mid T \rangle_h \triangleright \langle a \mid H[\text{handle } \square \text{ with } h, \lambda x. \text{handle } e \text{ with } h, r] \mid T \rangle_h$ $\langle \text{return } v \mid H[\text{handle } \square \text{ with } h, r] \mid T \rangle_h \triangleright \langle r v \mid H \mid T \rangle_e$ $\langle \text{do } v \mid H[\text{handle } \square \text{ with } h, r] \mid T \rangle_h \triangleright \langle H \mid \text{let } x = h v \text{ in } x r \mid T \rangle_c$ $\langle \text{lift } e \mid H[\text{handle } \square \text{ with } h, r] \mid T \rangle_h \triangleright \langle H \mid \text{let } x = e \text{ in } r x \mid T \rangle_c$ $\langle \text{let } x = a \text{ in } e \mid T \rangle_i \triangleright \langle a \mid \square \mid T[\text{let } x = \square \text{ in } e] \rangle_e$ $\langle \text{lift } e \mid T \rangle_i \triangleright \langle e \mid T[\text{lift } \square] \rangle_i$ $\langle \text{return } v \mid T \rangle_i \triangleright \langle T \mid v \rangle_p$ $\langle T[\text{let } x = \square \text{ in } e] \mid v \rangle_p \triangleright \langle e[v/x] \mid \square \mid T \rangle_e$ $\langle T[\text{lift } \square] \mid v \rangle_p \triangleright \langle T \mid v \rangle_p$		

Fig. 5. An abstract machine for fine-grained operational semantics. The machine starts in the configuration  $\langle e \mid \square \mid \square \rangle_e$ , for a closed expression  $e$ . The final configurations are  $\langle \square \mid v \rangle_p$ , representing a computed value  $v$ , and  $\langle \text{do } v \mid T \rangle_i$ , representing a control-stuck term  $T[\text{do } v]$ .

contexts is. A more realistic abstract machine, equipped with environments and other semantic components can be found in Section 7.

The syntactic abstract machine in Figure 5 has been obtained through Danvy and Nielsen’s refocusing method [Biernacka et al. 2017; Danvy and Nielsen 2004] from the fine-grained reduction semantics presented above. The structure of the machine coincides with the one of Materzok and Biernacki [2011] for a calculus of delimited control, where the computation is not guaranteed to be delimited by a reset or a handler, as in our case. The evaluation context is divided into two parts. The inner part  $H$  represents a (possibly empty) stack of handlers surrounding a given expression – that it is a stack of handlers rather than a stack of expressions, each guarded by a handler, is one of the most notable differences with the traditional semantics. The outer part  $T$  represents the top-level continuation, i.e., the rest of the context beyond the outermost handler.

The machine is in refocused form which means that all it does is decomposing a term and contracting the next redex it finds, followed by repeating this procedure from the resulting configuration (no term reconstruction is necessary). The machine operates in five modes:  $\langle e \mid H \mid T \rangle_e$  dispatches on the structure of the expression  $e$  being evaluated,  $\langle H \mid i \mid T \rangle_c$  dispatches on the structure of the handler stack  $H$  when given an intermediate value  $i$ ,  $\langle i \mid H \mid T \rangle_h$  dispatches on the structure of the intermediate value  $i$  when under a handler, i.e., when  $H$  is non-empty,  $\langle i \mid T \rangle_i$  dispatches on the structure of the intermediate value  $i$  when not under a handler, and  $\langle T \mid v \rangle_p$

dispatches on the structure of the top-level continuation  $T$  to complete the computation when given a value  $v$ . Each configuration represents a complete state of the machine, e.g.,  $\langle e \mid H \mid T \rangle_e$  should be understood as a term  $T[H[e]]$ . The initial configuration is  $\langle e \mid \square \mid \square \rangle_e$  for a closed expression  $e$ , whereas the final configurations are  $\langle \square \mid v \rangle_p$  and  $\langle \text{do } v \mid T \rangle_i$ .

The hybrid nature of the fine-grained semantics manifests itself in the abstract machine expressly: the intermediate values are interpreted differently in the configurations  $\langle i \mid H \mid T \rangle_h$  and  $\langle i \mid T \rangle_i$ .

## 2.2 The Expressive Power of Lift

The lift operator [Leijen 2017a] is a construct present in several calculi of effect handlers in the literature [Biernacki et al. 2018, 2019; Piróg et al. 2019], but its expressive power does not appear to have been formally investigated. It has been noted, however, that in a calculus with type-and-effect system for shift0, a more traditional delimited-control operator [Danvy and Filinski 1990; Kiselyov and Shan 2007; Materzok 2013], the lift operator can be straightforwardly macro expressed [Piróg et al. 2019]. In this section we show that lift is *not* eliminable, in the sense of Felleisen [1991], in the untyped calculi of deep handlers considered in this work. We work with  $\lambda_{n\text{IDH}}$ , but as we prove in Section 3, the two operational semantics are equivalent and the same reasoning can be carried out in  $\lambda_{fg\text{DH}}$ . This result is new and it justifies the inclusion of lift in the calculus we consider as a construct that increases its expressive power.<sup>5</sup>

More concretely, following Felleisen's work on the expressive power of programming languages [Felleisen 1991], we prove that there does not exist a map  $\phi$  from  $\lambda_{n\text{IDH}}$  to  $\lambda_{n\text{IDH}}^-$ , a restricted version of  $\lambda_{n\text{IDH}}$  where lift has been removed, such that:

- $\phi$  preserves program-ness, i.e., it maps closed expressions to closed expressions
- $\phi$  is a homomorphism on  $\lambda_{n\text{IDH}}^-$
- for any program  $e$  in  $\lambda_{n\text{IDH}}$ , we have that  $e$  terminates if and only, if  $\phi(e)$  terminates.

To this end, we first observe that  $\lambda_{n\text{IDH}}^-$  is a conservative restriction of  $\lambda_{n\text{IDH}}$ , as required by Felleisen's framework, i.e.:

- the term constructors of  $\lambda_{n\text{IDH}}^-$  are a subset of the term constructors of  $\lambda_{n\text{IDH}}$
- the phrases (expressions) of  $\lambda_{n\text{IDH}}^-$  are a subset of the phrases of  $\lambda_{n\text{IDH}}$
- the programs (closed expressions) of  $\lambda_{n\text{IDH}}^-$  are a subset of the programs of  $\lambda_{n\text{IDH}}$
- the semantics of  $\lambda_{n\text{IDH}}^-$  is a restriction of the semantics of  $\lambda_{n\text{IDH}}$ , i.e., the semantics of a program in  $\lambda_{n\text{IDH}}^-$  is given by the rules  $(\beta.v)$ ,  $(\text{DH}.v)$ ,  $(\beta.\text{let})$ , and  $(\text{DH}.j)$ , where, in the absence of the lift operator, the evaluation contexts  $J_0$  are exactly the non-capturing contexts, i.e., contexts where the hole is not guarded by a handler.

Next, let us consider the following program  $e$  in  $\lambda_{n\text{IDH}}$

$$\text{handle}(\text{handle}(\text{lift}(\text{do } v)) \text{ with } \lambda x . \lambda k . \Omega, \lambda x . \Omega) \text{ with } \lambda x . \lambda k . x, \lambda x . x$$

where  $v$  is any closed value. Applying the reduction rules of  $\lambda_{n\text{IDH}}$ , we can see that  $e$  evaluates to  $v$ . Indeed, the inner, diverging, handler is skipped by the lift operator, and the outer handler handles the operation by discarding the resumption and returning the value passed to the operation.

Then, if  $\phi$  is a map from  $\lambda_{n\text{IDH}}$  to  $\lambda_{n\text{IDH}}^-$  that preserves program-ness and is homomorphic on  $\lambda_{n\text{IDH}}^-$ , the image of  $e$  through  $\phi$  is the following program  $e'$

$$\text{handle}(\text{handle}(\phi(\text{lift}(\text{do } v))) \text{ with } \lambda x . \lambda k . \Omega, \lambda x . \Omega) \text{ with } \lambda x . \lambda k . x, \lambda x . x$$

<sup>5</sup>We mean the expressive power in a very precise sense, namely as defined by Felleisen. We only claim that lift cannot be *locally* translated away. Other, *global* translations, e.g., passing the level of freeness from operations to handlers are entirely possible.

Now, the (closed) expression  $\phi(\text{lift}(\text{do } v))$  which is in the evaluation position in  $e'$  can either diverge, or it can evaluate to a value  $u$ , or it can evaluate to a control-stuck term of the form  $J_n[\text{do } u]$  for some non-capturing context  $J_n$  and a value  $u$ . In each case, program  $e'$  diverges in  $\lambda_{n/\text{DH}}^-$ . Therefore, there does not exist a translation from  $\lambda_{n/\text{DH}}$  to  $\lambda_{n/\text{DH}}^-$  that makes it possible to locally express (let alone macro-express<sup>6</sup>) lift in terms of algebraic operations and deep handlers.

### 3 SIMULATION OF FINE-GRAINED EVALUATION BY NONLOCAL EVALUATION

In this section we study the relation of the two calculi defined above as evaluation theories. We establish two simulation properties, namely that:

- (1) nonlocal evaluation can be simulated by fine-grained evaluation, and
- (2) fine-grained evaluation can be simulated by nonlocal evaluation.

These are formally stated below.

**THEOREM 3.1 (SIMULATION OF NONLOCAL EVALUATION).** *If  $a \mapsto_{\lambda_{n/\text{DH}}}^* e$  and  $a \rightsquigarrow a'$ , then there exists  $e'$  such that  $e \rightsquigarrow e'$  and  $a' \mapsto_{\lambda_{\text{fgr}}^*} e'$ .*

**THEOREM 3.2 (SIMULATION OF FINE-GRAINED EVALUATION).** *If  $a \mapsto_{\lambda_{\text{fgr}}^*} e$  and  $a' \rightsquigarrow a$ , then there exists  $e'$  such that  $e' \rightsquigarrow \text{NF}_{\rightarrow}(e)$  and  $a' \mapsto_{\lambda_{n/\text{DH}}}^* e'$ .*

While we only define the similarity relation in the following section, we remark that, as corollaries of these theorems we obtain the fact that if a program  $e$  evaluates to a value  $v$  using one semantics, than it evaluates to some value  $v'$  using the other, and moreover  $v$  and  $v'$  have the same structure (up to  $\eta$ -expansion). From this corollary, in turn, it follows that the evaluation of a given closed term  $e$ , either terminates in both semantics yielding the same kind of result (a value or a control-stuck term), or it diverges in both semantics. Thus, the semantic functions arising from the two definitions can be considered equivalent.

#### 3.1 Similarity Relation and Flattening

Both simulations are formulated in terms of the similarity relation  $\rightsquigarrow$ , defined in Figure 6. The definition consists of three parts: the *structural* relation on expressions ( $\sim$ ), the relation on values ( $\sim_v$ ) and the similarity ( $\rightsquigarrow$ ). These are defined in a way that enforces a structure of terms that generally matches on both sides, but accounts for the distinct nature of the two semantics.

The structural similarity,  $\sim$ , is built congruentially over the relation on values, enforcing a matching structure of related computations. In turn, it is used by  $\rightsquigarrow$  to loosen this relationship by allowing the fine-grained evaluation to perform certain reductions. These “simulation” reductions, written  $\rightarrow$ , comprise the contractions that describe the interactions of the handle expressions with other computation formers – let and lift expressions ((**DH.let**) and (**DH.lift**) respectively) – and values (**DH.v**). We discuss the rationale behind this choice in the following. Finally, the relation on values relates values of the matching structure (which in the case of variables enforces their equality), while allowing the bodies of functions to be related by  $\rightsquigarrow$ . It also allows the nonlocal values to be  $\eta$ -longer than their fine-grained counterparts. Both of these allowances stem from the way the nonlocal calculus handles resumptions. First, their reification introduces “new” lambda-abstractions, which may lead to forms that are  $\eta$ -longer than on the fine-grained side. The second aspect is connected to the rationale behind the simulation reductions.

To understand the intuition behind the inclusion of simulation reductions, consider a (**DH.j**)-redex: a handle expression with a do expression plugged into a 0-free context  $J$ . Since  $J$  may contain

<sup>6</sup>For macro-expressibility, the translation function  $\phi$  is additionally required to be defined compositionally on the expression being eliminated, in this case, the lift expression [Felleisen 1991].

$$\begin{array}{c}
(\rightsquigarrow) \subseteq \lambda_{n\text{DH}} \times \lambda_{f\text{gDH}} \\
\\
\frac{}{x \sim_v x} \text{VAR} \qquad \frac{e \rightsquigarrow e'}{\lambda x . e \sim_v \lambda x . e'} \lambda \qquad \frac{v \sim_v \lambda x . v' x \quad x \notin \text{FV}(v')}{v \sim_v v'} \eta \\
\\
\frac{v \sim_v v'}{\text{return } v \sim \text{return } v'} \text{RET} \qquad \frac{v \sim_v v'}{\text{do } v \sim \text{do } v'} \text{DO} \qquad \frac{u \sim_v u' \quad v \sim_v v'}{uv \sim u'v'} \text{APP} \\
\\
\frac{a \sim a' \quad e \sim e'}{\text{let } x = a \text{ in } e \sim \text{let } x = a' \text{ in } e'} \text{LET} \qquad \frac{e \sim e'}{\text{lift } e \sim \text{lift } e'} \text{LIFT} \\
\\
\frac{e \sim e' \quad h \sim_v h' \quad r \sim_v r'}{\text{handle } e \text{ with } h, r \sim \text{handle } e' \text{ with } h', r'} \text{HANDLE} \\
\\
\frac{e \mapsto_{(\text{DH},v)} e' \vee e \mapsto_{(\text{DH},\text{let})} e' \vee e \mapsto_{(\text{DH},\text{lift})} e'}{E_n[e] \rightsquigarrow E_n[e']} \text{STEP} \qquad \frac{e \sim e' \quad e' \rightsquigarrow^* e''}{e \rightsquigarrow e''} \text{SIM}
\end{array}$$

Fig. 6. Similarity relation  $\rightsquigarrow$ .

let- and lift-expressions, we cannot guarantee that a structurally similar term is a (DH.do)-redex. However, a fine-grained computation starting from a structurally related term will always reach such a redex by repeated application of (DH.let) and (DH.lift) rules. Since  $J$  is 0-free, we can in fact guarantee that it will be appropriately flattened and stored as part of the extended continuation of the original handle expression. Thus, we would in fact reach a (DH.do)-redex with related handlers and arguments, and a significantly extended return clause. However, this extended return clause (which takes the role of the resumption in the reduct of (DH.do)) is not, in most cases, structurally related to the resumption captured by (DH.j), since the reordering of  $J$  was already performed. In order to bridge this gap, we allow the simulation reductions to happen silently on the fine-grained side of the simulation relation and under lambda-expressions related by the similarity of values.

The flattening reduction,  $\rightsquigarrow$ , is very simple, and trivially normalising. This is useful in the proof process, as its normal forms are very simple: they are either a return expression in a 0-level evaluation context (i.e., one that contains no handlers), or an application or do expression, this time in an arbitrary evaluation context. Thus, the slightly imprecise simulation in Theorem 3.2, which only gives us similarity with a fully flattened version of the term, is not particularly problematic.

### 3.2 Simulation of Nonlocal Evaluation

To simulate nonlocal evaluation, we simulate each step separately. Each step is simulated by zero or more steps of fine-grained evaluation. To simulate a single step, we employ a lemma on simulation of a contraction and certain facts about the simulation reduction,  $\rightsquigarrow$ , and flattening.

Most of the cases of simulation of a contraction are simple. The first interesting case, (DH.j), it uses a Context Capture Lemma, described informally above, which relates how any 0-free context in a handle-expression can be moved to resumption part of a handle-expression by flattening. The other non-trivial case, ( $\beta$ .v), relies on a separate lemma proven by structural induction on the derivation of  $\sim_v$ -similarity between a  $\lambda$ -abstraction and a value, which needs to account for possible  $\eta$ -expansions.

Simulation of a single step uses also an inversion lemma for  $\sim$ , where the left hand side is a term plugged in a nonlocal evaluation context. It work in tandem with a procedure and lemmas for context flattening to reduce problem to simulation of a contraction.

### 3.3 Simulation of Fine-Grained Evaluation

For the purpose of this proof, we divide the fine-grained reductions into two categories: the flattening reductions ( $\rightarrow$ ), already discussed above, and the remaining, which we dub *real*. Since the fine-grained term may be significantly reordered by flattening, we proceed by decomposing the non-local term, to establish its shape. As in the other directions, the most interesting cases arise when the expression is a  $(\mathbf{DH}.j)$  or  $(\beta.v)$  redex in a context, and the shape of the non-local term restricts the real step that the fine-grained side could have taken (any simulation steps can be folded into the simulation relation without taking any steps on the non-local side). Although this restriction is not entirely precise (for instance, a  $(\beta.let)$  redex on the non-local side can give rise to either  $(\beta.let)$  or  $(\beta.v)$  redexes on the fine-grained side, since the let-expression may turn into a handle through flattening), it is sufficient to recover the necessary elements of the simulation. In the case of a  $(\mathbf{DH}.j)$ -redex we need to account for the flattening of the 0-free context in a manner similar to the other direction of the proof, and the cases of  $(\beta.let)$ - and  $(lift.v)$ -redexes are only involved due to the complexities arising from flattening of the evaluation context.

The remaining two cases, however, introduce substantial complexity. This is due to the fact that both  $(\beta.v)$  and  $(\mathbf{DH}.v)$ -redexes in the non-local calculus lead to  $(\beta.v)$  redexes on the fine-grained side – and the similarity relation for values may introduce an unbounded number of  $\eta$ -expansions on the non-local side. Moreover, since the bodies of the functions may be similar only up to flattening, this may force the non-local side to perform a significant amount of computation in order to obtain an appropriately similar substitution. The resulting technical lemma turns out to be rather complex, and we present it below.

**LEMMA 3.3.** *Let  $e_1 \rightsquigarrow_{\text{NF}}^{\perp} e_2 \triangleq \exists e'_1. e_1 \mapsto_{\lambda_{\text{ndH}}}^* e'_1 \wedge e'_1 \rightsquigarrow \text{NF}_{\rightarrow}(e_2)$ , and let  $\sim_K$  denote the natural extension of the similarity relation  $\sim$  to evaluation contexts  $K$  of non-local semantics. Then, the following implications hold:*

- if  $K_1 \sim_K K_2, v_1 \sim_v v_2$  and  $u \sim_v \eta^n(\lambda x. e)$ , then  $K_1[u v_1] \rightsquigarrow_{\text{NF}}^{\perp} K_2[e[v_2/x]]$ ;
- if  $e_1 \rightsquigarrow E[(\eta^n(\lambda x. e)) v]$ , then  $e_1 \rightsquigarrow_{\text{NF}}^{\perp} E[e[v/x]]$ ;
- if  $K_1 \sim_K K_2, K_2[e_2] \rightsquigarrow^* E[(\eta^n(\lambda x. e)) v]$ ,  $e_1 \sim e_2$  and  $\text{wf}(e_1)$ , then  $K_1[e_1] \rightsquigarrow_{\text{NF}}^{\perp} E[e[v/x]]$ .

Above,  $\eta^n(v)$  denotes  $n$ -fold  $\eta$ -expansion of  $v$ , and the expression is considered well-formed when, assuming it is a value in an evaluation context, the context terminates with a 0-free part contained in a handler, i.e.,

$$\text{wf}(e) \triangleq \forall v, K. e = K[\text{return } v] \Rightarrow \exists K', h, r, J_0. K = K'[\text{handle } J \text{ with } h, r].$$

The proof proceeds by mutual induction on the three aspects of the similarity definition, with the dependency structure matching the one of the definition. For the first lemma, the case for variables is impossible, and the one for  $\eta$ -expansion is solved inductively. For lambda-abstraction, we can either establish the conclusion directly, if  $n = 0$ , or use the second lemma in the other case. The second lemma establishes well-formedness of its argument and uses the third, while the third proceeds inductively in most cases. The interesting cases here are application, when we can use the first lemma inductively, and the case for handle if the context  $K'$  resulting from the well-formedness is empty (and thus  $e_1 = \text{handle } J_0 \text{ with } h, r$ ). Here, depending on the shape of  $J$  we can either use the first lemma inductively (if  $J$  is empty), or establish the conclusion directly, since in the other cases no  $\eta$ -expansions are possible.

$$\begin{array}{c}
\frac{}{x \Rightarrow_v x} \text{S-VAR} \qquad \frac{e \Rightarrow_e e'}{\lambda x . e \Rightarrow_v \lambda x . e'} \text{S-LAM} \qquad \frac{v \Rightarrow_v v'}{\text{return } v \Rightarrow_e \text{return } v'} \text{S-RET} \\
\frac{}{(x, e_1, e_2) \Rightarrow_{\text{let}} \text{let } x = e_1 \text{ in } e_2} \text{NL-LET} \qquad \frac{}{\text{let } x = \text{return } v_1 \text{ in } e_2 \Rightarrow_{\text{let}} e_2[v_1/x]} \text{NL-RED} \\
\frac{e_1 \Rightarrow_e e'_1 \quad e_2 \Rightarrow_e e'_2 \quad (x, e'_1, e'_2) \Rightarrow_{\text{let}} e}{\text{let } x = e_1 \text{ in } e_2 \Rightarrow_e e} \text{S-LET}
\end{array}$$

Fig. 7. Selected rules from the definition of the super-parallel reduction  $\Rightarrow$  for the pure fragment of the calculus.

#### 4 CONFLUENCE OF MULTI-STEP FINE-GRAINED REDUCTION

In this section we present a proof sketch of one of the most fundamental properties of the reduction relation  $\rightarrow_{\lambda_{\text{fgDH}}}$  one expects, namely the confluence of the reduction relation or the Church-Rosser theorem [Barendregt 1985], expressed as the following diamond property for  $\rightarrow_{\lambda_{\text{fgDH}}}^*$ :

**THEOREM 4.1 (CONFLUENCE).** *If  $e \rightarrow_{\lambda_{\text{fgDH}}}^* e_1$  and  $e \rightarrow_{\lambda_{\text{fgDH}}}^* e_2$ , then there exists a term  $e'$ , such that  $e_1 \rightarrow_{\lambda_{\text{fgDH}}}^* e'$  and  $e_2 \rightarrow_{\lambda_{\text{fgDH}}}^* e'$ .*

We prove this theorem using a technique of *super-parallel reductions* and *complete super developments* [Aczel 1978; McKinna 2022; McKinna and Pollack 1999; Takahashi 1995; van Raamsdonk 1993], which is a simplification of the classic proof method by *parallel reductions* due to Tait and Martin-Löf [Barendregt 1985]. The idea is to define a super-parallel reduction relation  $\Rightarrow_e$  on expressions (along with a super-parallel reduction relation  $\Rightarrow_v$  on values, with analogous properties), such that

- (1) if  $e \rightarrow_{\lambda_{\text{fgDH}}} e'$  then  $e \Rightarrow_e e'$
- (2) if  $e \Rightarrow_e e'$  then  $e \rightarrow_{\lambda_{\text{fgDH}}}^* e'$
- (3) if  $e_1 \Rightarrow_e e'_1$  and  $e_2 \Rightarrow_e e'_2$  then  $e_1[e_2/x] \Rightarrow_e e'_1[e'_2/x]$  for any  $x$  that may occur freely in  $e_2, e'_2$ .

By (1) and (2), we have that  $\rightarrow_{\lambda_{\text{fgDH}}}^*$  and  $\Rightarrow_e^*$  coincide, and therefore, to establish the confluence of  $\rightarrow_{\lambda_{\text{fgDH}}}$ , it suffices to show that  $\Rightarrow_e$  satisfies the diamond property, which requires (3). All three properties are proved by a routine induction on either  $\rightarrow_{\lambda_{\text{fgDH}}}$  or  $\Rightarrow_e$ .

In Figure 7, we present a couple of representative rules defining  $\Rightarrow_e$  for the pure fragment of the calculus. We can observe that the definition is almost identical with the one, we would write if we decided to use a parallel reduction relation instead. The only difference is in the combination of the rules S-LET and NL-RED, where a potentially newly created redex  $\text{let } x = \text{return } v_1 \text{ in } e'_2$  is contracted on the fly. In contrast, the parallel reduction relation contracts only redexes that are present in the original term. So, the parallel reduction rule that would correspond to the combination of S-LET and NL-RED would read:

$$\frac{v_1 \Rightarrow_v \text{return } v'_1 \quad e_2 \Rightarrow_e e'_2}{\text{let } x = \text{return } v_1 \text{ in } e_2 \Rightarrow_e e'_2[v_1/x]}$$

This extra eagerness of the super-parallel reductions is not critical for the proof of the diamond property of  $\Rightarrow_e$ , but in conjunction with a more eager version of the complete developments that we use in the proof [McKinna 2022], it slightly simplifies the Coq formalization – we save on some inversions.

$$\begin{array}{c}
\frac{}{(e, h, r) \Rightarrow_{\text{handle}} \text{handle } e \text{ with } h, r} \text{NH-HANDLE} \qquad \frac{(r, v) \Rightarrow_{\text{app}} e}{(\text{return } v, h, r) \Rightarrow_{\text{handle}} e} \text{NH-RET} \\
\\
\frac{x \notin \text{FV}(h, r)}{(\text{let } x = a \text{ in } e, h, r) \Rightarrow_{\text{handle}} \text{handle } a \text{ with } h, \lambda x . \text{handle } e \text{ with } h, r} \text{NH-LET} \\
\\
\frac{(r, x) \Rightarrow_{\text{app}} e_1 \quad (x, e, e_1) \Rightarrow_{\text{let}} e_2 \quad x \notin \text{FV}(r)}{(\text{lift } e, h, r) \Rightarrow_{\text{handle}} e_2} \text{NH-LIFT} \\
\\
\frac{x \notin \text{FV}(r)}{(\text{do } v, h, r) \Rightarrow_{\text{handle}} \text{let } x = h v \text{ in } x r} \text{NH-LET} \\
\\
\frac{e_1 \Rightarrow e_2 \quad h_1 \Rightarrow_v h_2 \quad r_1 \Rightarrow_v r_2 \quad (e_2, h_2, r_2) \Rightarrow_{\text{handle}} e}{\text{handle } e_1 \text{ with } h_1, r_1 \Rightarrow_e e} \text{S-HANDLE}
\end{array}$$

Fig. 8. Excerpt from definition of superparallel reduction  $\Rightarrow$ .

The remaining redexes are treated along the same lines. In particular, in Figure 8 we present the rules for the handle construct. The auxiliary relation  $\Rightarrow_{\text{handle}}$  non-deterministically either reconstructs a handle expression from the given subexpressions or it contracts the redex introduced by the relation  $\Rightarrow_e$  on the subexpression of the original term.

As for the proof of the diamond property of  $\Rightarrow_e$ , we take advantage of the notion of the complete super development of a term, which is a *deterministic* relation  $\Rightarrow_e^c$  ( $\Rightarrow_e^c$ ) that simultaneously contracts all the redexes existing in a term as well as those introduced by the rules such as NL-RED. More precisely,  $\Rightarrow_e^c$  is a subrelation of  $\Rightarrow_e$  obtained from  $\Rightarrow_e$  by restricting the rules NL-LET, NH-HANDLE, etc., so that they are not allowed to introduce a top-level redex. For example, in NL-LET,  $e_1$  is not allowed to have the form  $\text{return } v$ .

Since the relation  $\Rightarrow_e^c$  contracts at least as many redexes as  $\Rightarrow_e$ , we are in a position to prove by induction on  $\Rightarrow_e^c$  the *triangle lemma* of Takahashi [1995]:

LEMMA 4.2. *If  $e \Rightarrow_e e_1$  and  $e \Rightarrow_e^c e_2$ , then  $e_1 \Rightarrow_e e_2$ .*

For the proof of this lemma, an essential role in the definition of  $\Rightarrow_e$  is played by the auxiliary relations such as  $\Rightarrow_{\text{handle}}$ . They allow the relation  $\Rightarrow_e$  to catch up with  $\Rightarrow_e^c$  in the presence of critical pairs such as, e.g., in the expression

$$\text{handle let } x = \text{return } v \text{ in } e \text{ with } h, r$$

where reducing the outer redex leads to an intermediate term

$$\text{handle return } v \text{ with } h, \lambda x . \text{handle } e \text{ with } h, r$$

that can be then reduced in two steps to

$$\text{handle } e[v/x] \text{ with } h, r$$

thanks to the rule NH-RET. This rule takes advantage of the auxiliary relation  $\Rightarrow_{\text{app}}$ , not shown here, which can either reconstruct an application of two values or, as in our example,  $\beta$ -contract it.

Now, using the facts that  $\Rightarrow_e^c$  is functional, i.e., deterministic and total, if  $e \Rightarrow_e e_1$  and  $e \Rightarrow_e e_2$ , we can construct two triangles given by Lemma 4.2, and compose them into the desired diamond,



such that  $e_1 \Rightarrow_e e'$  and  $e_2 \Rightarrow_e e'$ , where  $e'$  is the unique term such that  $e \Rightarrow_e^c e'$  ( $e'$  does not depend on  $e_1$  or  $e_2$ ).

## 5 STANDARDISATION THEOREM

The standardisation theorem, in its classical form [Mitschke 1979], states that any reduction sequence that reaches a head-normal form can be separated into a head-reduction sequence that reaches a normal form, followed by a number of *internal* reductions. As noted by McKinna and Pollack [1999], this can be adapted to the more practically useful notion of weak-head reductions, and since the important consequences follow already from Mitschke's semi-standardisation lemma, this is the problem we focus on in this section.

To prove standardisation, we adapt the methodology developed by Takahashi [1995] and refined by McKinna and Pollack [McKinna 2022; McKinna and Pollack 1999]. The crux of the idea is to use *internal parallel* reductions to describe the internal (i.e., non-weak-head) reductions of the standard sequence. In its refined form, the *standard* reduction sequences can then be defined as a composition of weak-head reduction sequences with internal-parallel reductions, and semi-standardisation simply states that any reduction sequence can be mimicked by a standard one.

While the methodology is very refined in the case of pure lambda calculus, the fact that our notion of evaluation (with which we want the weak-head reduction to coincide) is a hybrid strategy introduces new challenges. This is because the *top* contractions ( $(\beta.let)$  and  $(lift.v)$ ), while part of the weak-head reduction in appropriate contexts, become *frozen* if performed within a handle expression — and both *let* and *lift* constructs cease to act as evaluation contexts within handle. Thus, we lose a clear, structural distinction between evaluation positions, where any weak-head reductions remain weak-head, and the remaining positions, where no weak-head reduction is possible.

These considerations led us to the definitions of weak-head, frozen, internal-parallel and standard reductions presented in Figure 9. To obtain a well-behaved definition of frozen reductions we repeat McKinna's observation that standard sequences can be *defined* as composition of weak-head and internal reductions and adapt it as a definition of weak-head sequences: these are now formed as a composition of *inner* reductions ( $\mapsto_1^*$ ), which are always allowed within a handle expression, and the frozen reductions. The latter are defined to be reflexive and encompass weak-head reduction sequences under *let*, *lift*, and those including the two *top* contractions. It is easily shown that this notion of weak-head reduction corresponds to the notion of fine-grained evaluation ( $\mapsto_{\lambda/\beta/DH}^*$ ) defined in Section 2.

This leaves us with the definition of standard and internal-parallel reductions. These follow the recent presentation by McKinna [2022], adapted to the fine-grained call-by-value setting. Thus, we get the notions of internal-parallel reductions for both expressions ( $\mapsto_e^!$ ) and values ( $\mapsto_v^!$ ), with the definitions following the structure of terms and allowing:

- *internal* reductions for any value sub-terms;
- *standard* reductions for any expression sub-terms outside of the evaluation position;
- *internal* reductions for expression sub-terms in evaluation position of *let* and *lift*;
- *frozen* reductions followed by *internal* reductions for the evaluation position of handle.

The last case accounts for the hybrid nature of the weak-head reduction, as discussed above. Finally, the standard sequences are defined as composition of weak-head and internal reductions.

The crucial step in the methodology we follow is establishing transitivity of the relations defined as above. For this, it suffices to know that standard reductions *absorb* contractions and weak-head reductions from the right (cf. Lemma 5.5 for finer detail). This is the heart of the argument that reductions can be appropriately reordered, and is always somewhat involved, as internal reductions can become unblocked through the commutation of the weak-head reduction in front of them.

$$\begin{array}{c}
\frac{a \rightarrow^w \text{return } v \quad e[v/x] \rightarrow^w e'}{\text{let } x = a \text{ in } e \rightarrow^f e'} \text{ F-LET-RET} \qquad \frac{a \rightarrow^w a'}{\text{let } x = a \text{ in } e \rightarrow^f \text{let } x = a' \text{ in } e} \text{ F-LET} \\
\\
\frac{e \rightarrow^w \text{return } v}{\text{lift } e \rightarrow^f v} \text{ F-LIFT-RET} \qquad \frac{e \rightarrow^w e'}{\text{lift } e \rightarrow^f \text{lift } e'} \text{ F-LIFT} \qquad \frac{}{\text{return } v \rightarrow^f \text{return } v} \text{ F-RET} \\
\\
\frac{}{uv \rightarrow^f uv} \text{ F-APP} \qquad \frac{}{\text{do } v \rightarrow^f \text{do } v} \text{ F-DO} \qquad \frac{}{\text{handle } e \text{ with } h, r \rightarrow^f \text{handle } e \text{ with } h, r} \text{ F-HANDLE} \\
\\
\frac{e_1 \mapsto_1^* e_2 \quad e_2 \rightarrow^f e_3}{e_1 \rightarrow^w e_3} \text{ W} \\
\\
\frac{}{x \rightarrow_v^i x} \text{ I-VAR} \qquad \frac{e_1 \rightarrow^s e_2}{\lambda x. e_1 \rightarrow_v^i \lambda x. e_2} \text{ I-LAM} \qquad \frac{v_1 \rightarrow_v^i v_2}{\text{return } v_1 \rightarrow_e^i \text{return } v_2} \text{ I-RET} \\
\\
\frac{a_1 \rightarrow_e^i a_2 \quad e_1 \rightarrow^s e_2}{\text{let } x = a_1 \text{ in } e_1 \rightarrow_e^i \text{let } x = a_2 \text{ in } e_2} \text{ I-LET} \qquad \frac{e_1 \rightarrow_e^i e_2}{\text{lift } e_1 \rightarrow_e^i \text{lift } e_2} \text{ I-LIFT} \qquad \frac{v_1 \rightarrow_v^i v_2}{\text{do } v_1 \rightarrow_e^i \text{do } v_2} \text{ I-DO} \\
\\
\frac{u_1 \rightarrow_v^i u_2 \quad v_1 \rightarrow_v^i v_2}{u_1 v_1 \rightarrow_e^i u_2 v_2} \text{ I-APP} \qquad \frac{e_1 \rightarrow^f e_2 \quad e_2 \rightarrow_e^i e_3 \quad h_1 \rightarrow_v^i h_2 \quad r_1 \rightarrow_v^i r_2}{\text{handle } e_1 \text{ with } h_1, r_1 \rightarrow_e^i \text{handle } e_3 \text{ with } h_2, r_2} \text{ I-HANDLE} \\
\\
\frac{e_1 \rightarrow^w e_2 \quad e_2 \rightarrow_e^i e_3}{e_1 \rightarrow^s e_3} \text{ S}
\end{array}$$

Fig. 9. Standard reduction ( $\rightarrow^s$ ), internal reduction ( $\rightarrow_e^i$ )/( $\rightarrow_v^i$ ), frozen weak head reduction ( $\rightarrow^f$ ), and weak head reduction ( $\rightarrow^w$ ).

Most of these cases are handled in a standard way: what remains, are the frozen reductions in the evaluation position of handle, which can also get unblocked by the following reduction. This leads us to the following four lemmas, corresponding to the four contractions of handle.

LEMMA 5.1 (HANDLING OF VALUES). *The following implications hold:*

- If  $e \rightarrow^f \text{return } v$ , then  $\text{handle } e \text{ with } h, r \rightarrow^w r v$ .
- If  $e \rightarrow^w \text{return } v$ , then  $\text{handle } e \text{ with } h, r \rightarrow^w r v$ .

LEMMA 5.2 (HANDLING OF EFFECTS). *The following implications hold:*

- If  $e \rightarrow^f$  do  $v$ , then handle  $e$  with  $h, r \rightarrow^w$  let  $x = hv$  in  $xr$ .
- If  $e \rightarrow^w$  do  $v$ , then handle  $e$  with  $h, r \rightarrow^w$  let  $x = hv$  in  $xr$ .

LEMMA 5.3 (HANDLING OF LIFTS). *The following implications hold:*

- If  $e \rightarrow^f$  lift  $e'$ , then handle  $e$  with  $h, r \rightarrow^w$  let  $x = e'$  in  $rx$ .
- If  $e \rightarrow^w$  lift  $e'$ , then handle  $e$  with  $h, r \rightarrow^w$  let  $x = e'$  in  $rx$ .

LEMMA 5.4 (HANDLING OF LET-EXPRESSIONS). *The following implications hold:*

- If  $e \rightarrow^f$  let  $x = a$  in  $e'$ , then exists  $a'$  such that handle  $e$  with  $h, r \rightarrow^w$  handle  $a'$  with  $h, \lambda x$ . handle  $e'$  with  $h, r$  and  $a' \rightarrow^f a$ .
- If  $e \rightarrow^w$  let  $x = a$  in  $e'$ , then exists  $a'$  such that handle  $e$  with  $h, r \rightarrow^w$  handle  $a'$  with  $h, \lambda x$ . handle  $e'$  with  $h, r$  and  $a' \rightarrow^f a$ .

In the first three cases, these simply state that a frozen reduction under handle followed by an appropriate contraction can be simulated with a weak-head reduction sequence. In the final one, however, we separate the original reduction sequence into two parts: a reduction sequence that exposes a let-expression, and any frozen reductions within its evaluation position. While the former sequence can be followed with a weak-head reduction, the latter become internal to the resulting handler when absorbed as part of the standard reduction. Note that these lemmas correspond to some of the critical pairs in the reduction theory, which we needed to consider in the previous section: in a sense, they provide a resolution of these critical pairs in concrete cases that arise from evaluation.

Having handled the unfreezing of reductions frozen within the handle expression, we can now return to establishing transitivity of internal and standard reduction relations. We prove, in turn, that internal reductions on expressions absorb contractions, weak-head and frozen reductions from the right. Most of these give rise to standard reduction sequences, but we establish more precise commutation results for *top* contractions and frozen reductions. These allow us to establish transitivity for internal reductions on expressions, where a lack of precision would only give us standard reduction sequence as a result. While not strictly necessary, this is a good litmus test of our definition of internal reductions for expressions. With transitivity established, it is easy to prove semi-standardisation.

LEMMA 5.5 (RECOMPOSITION). *Let  $\rightarrow$  denote left-to-right relation composition. The following inclusions hold:*

- $(\rightarrow_e^i); (\rightarrow_c) \subseteq (\rightarrow^s)$ ,
- $(\rightarrow_e^i); (\rightarrow_f) \subseteq (\rightarrow^s)$ ,
- $(\rightarrow_e^i); (\rightarrow_t) \subseteq (\rightarrow^f); (\rightarrow_e^i)$ ,
- $(\rightarrow_e^i); (\rightarrow_i) \subseteq (\rightarrow^s)$ ,
- $(\rightarrow_e^i); (\rightarrow_i^*) \subseteq (\rightarrow^s)$ ,
- $(\rightarrow_e^i); (\rightarrow^f) \subseteq (\rightarrow^f); (\rightarrow_e^i)$ ,
- $(\rightarrow_e^i); (\rightarrow^w) \subseteq (\rightarrow^s)$ ,
- $(\rightarrow_e^i); (\rightarrow_e^i) \subseteq (\rightarrow_e^i)$ ,
- $(\rightarrow_v^i); (\rightarrow_v^i) \subseteq (\rightarrow_v^i)$ ,
- $(\rightarrow^s); (\rightarrow^s) \subseteq (\rightarrow^s)$ .

THEOREM 5.6 (SEMI-STANDARDISATION). *If  $e \rightarrow_{\lambda_{\text{gDH}}}^* e'$ , then  $e \rightarrow^s e'$ .*

PROOF. Since standard reduction is reflexive and transitive, it remains to show that single-step reduction is a subrelation of standard reduction, which follows by simple induction.  $\square$

This is, to the best of our knowledge, the first explicitly stated standardisation result for a theory of effect handlers. Having said that, a careful reading of McLaughlin’s proof of the Context Lemma for ELLA [McLaughlin 2020, Chapter 6], a core calculus of the Frank programming language [Convent et al. 2020] and thus substantially different in terms of structuring and handling effects, reveals the structure of weak-head and internal reductions characteristic to standardisation.

## 6 EQUATIONAL REASONING

The results presented thus far license the fine-grained reduction theory we have developed to be used as a foundation for an equational theory that is sound with respect to the standard nonlocal semantics. To formally state this, let us define  $=_{\lambda_{fgDH}}$  as a symmetric closure of  $\rightarrow_{\lambda_{fgDH}}$ . Let us also define the notion of observational equivalence  $\approx_{\lambda_{nIDH}}$  as follows:  $a \approx_{\lambda_{nIDH}} e$  ( $a$  is observationally equivalent to  $e$ ), if for all contexts  $C$  such that both  $C[a]$  and  $C[e]$  are closed terms,  $(\exists v, C[a] \mapsto_{\lambda_{nIDH}}^* \text{return } v)$  if and only if  $(\exists u, C[e] \mapsto_{\lambda_{nIDH}}^* \text{return } u)$ .<sup>7</sup>

**COROLLARY 6.1 (SOUNDNESS OF FINE-GRAINED EQUATIONAL REASONING).** *If  $a =_{\lambda_{fgDH}} e$  then  $a \approx_{\lambda_{nIDH}} e$ .*

**PROOF.** Let us assume that  $a =_{\lambda_{fgDH}} e$  and  $C[a] \mapsto_{\lambda_{nIDH}}^* v$ . From the former assumption, we know that  $C[a] =_{\lambda_{fgDH}} C[e]$ . From the latter assumption, we get that  $C[a] \mapsto_{\lambda_{fgDH}}^* v'$  by the simulation of nonlocal evaluation (Theorem 3.1). Therefore, we have that  $C[e] =_{\lambda_{fgDH}} v'$ , which by confluence of  $\rightarrow_{\lambda_{fgDH}}$  (Theorem 4.1) gives us that  $C[e] \rightarrow_{\lambda_{fgDH}}^* w$  for some value  $w$ . Now, applying the standardisation theorem (Theorem 5.6), we obtain that  $C[e] \mapsto_{\lambda_{fgDH}}^* w'$  for some value  $w'$ . Finally, by the simulation of fine-grained evaluation (Theorem 3.2), we arrive at the conclusion  $C[e] \mapsto_{\lambda_{nIDH}}^* u$  for some value  $u$ .  $\square$

From the above corollary it follows that the directed reduction theory defined as  $\rightarrow_{\lambda_{fgDH}}$  and understood as a program optimization, as in the motivating example in Section 1, is sound.

## 7 NORMAL FORMS AND NORMALISATION

One of the useful consequences of the introduction of a reduction theory for effect handlers is the ability to study its normal forms and normalisation processes. There are two traditional applications of such results: firstly, in partial evaluation, or specialisation of programs – such as the example presented in the introduction – and as a means to determine convertibility in dependent type theories. While only the first of these is directly applicable to our current setup (due to its untyped character), we believe that the results presented below may be useful in attempts to marry effects and effect handlers with dependent types.

We begin by characterising normal forms of our calculus, presented in Figure 10. The normal computations ( $\mu$ ) consist of either normal values ( $\xi$ ) or *neutral computations* ( $\nu$ ), with normal values encompassing variables and lambda-abstractions with normalised bodies. The remainder of our normal forms are the neutral computations: expressions that are either *control-stuck*, i.e., containing a do expression unenclosed by a handle, or a *environment-stuck*, i.e., containing an application of a variable to a (normalised) value in a (normalised) evaluation context. The distinction in the allowable contexts is achieved, by indexing them with the number of handlers, and preventing the appearance of control-stuck terms if there are handlers present. Note that this description is somewhat removed from a traditional separation into normal and neutral forms of terms, where the neutral class encompasses both elimination forms and variables. In our case, variables  $x$  and their

<sup>7</sup>It can be shown, straightforwardly adjusting the argument from [Biernacki et al. 2020], that this notion of observational equivalence coincides with the standard one, in which termination, rather than termination to a value, is observed.

$$\begin{aligned}
\mu &::= \xi \mid v_0 && \text{(normal expressions)} \\
\xi &::= x \mid \lambda x . \mu && \text{(normal values)} \\
v_n &::= \text{do } \xi_{(n=0)} \mid \text{let } x = v_n \text{ in } \mu_{(n=0)} \mid \text{lift } v_{n(n=0)} \mid x \xi \mid \text{handle } v_{n+1} \text{ with } \xi, \xi' && \text{(neutral expressions)}
\end{aligned}$$

Fig. 10. Normal forms for fine-grained reduction

eliminations  $x \xi$  belong to two different classes. This is due to their distinct reduction behaviour: the term  $\text{let } x = y \text{ in } e$  cannot evaluate further (although reductions within  $v$  or  $e$  may still be allowed), while  $\text{let } x = y \text{ in } e$  is a valid redex. We cannot extend the notion of substitutions and allow the first computation to reduce, since a substitution of certain values for  $y$  may cause effects – and in such a case the behaviour of a hypothetical term  $e[y v/x]$  could easily differ from the original! Therefore, the let-bindings, lift and handle expressions that organise the computation need to be preserved if they enclose a neutral computation, but not if they enclose a value: a substitution in a value cannot turn a stuck term into a redex.

We formalise the notion that our definitions characterise normal forms by the following lemmas (note that  $\mu$  and  $v$  trivially inject into expressions, and  $\xi$  into values):

LEMMA 7.1. *The sets  $\mu$  and  $\xi$  are precisely normal forms of expressions and values, respectively, with respect to fine-grained reduction  $\rightarrow_{\lambda_{\text{fgrDH}}}$ , i.e.,*

- (1) *Normal forms  $(\mu, v_n, \xi)$  do not reduce;*
- (2) *Any expression  $e$  (respectively, value  $v$ ) either reduces to some expression  $e'$  (value  $v'$ ) or is a normal form  $\mu = e$  ( $\xi = v$ ).*

## 7.1 Normalisation by Evaluation

Since the equational theory of our calculus is given through explicitly directed, confluent reductions, it is natural to compute normal forms of expressions (where they exist, as the calculus is, in general, non-terminating) through a combination of (open-term) evaluation and readback functions. We begin by presenting the evaluator in two styles: as a partial function, implementable in a general-purpose programming language, such as OCaml, and as its graph, implementable as an inductively-defined relation; the latter corresponds to our formalisation in Coq. Both evaluators work utilise environment semantics rather than substitutions; therefore, we begin by introducing the necessary semantic artifacts.

*Semantic constructs.* The necessary components of an environment-based evaluator comprise semantic values, *functions* (lambda-like values), environments, closures (computations), stacks and answers; these are presented in Figure 11. Note that syntactic expressions appear only as parts of the closures, and the remainder of the instrumentation is achieved through simple inductive definitions.<sup>8</sup> For the sake of uniformity, the evaluation will operate on pairs of stacks and *configurations*, and produce an answer as a result. The first three of these configurations (annotated with E, A and F) will inspect the first element of the configuration, while the remaining four will inspect the stack.

Note that additional semantic functions and computations are present, beyond the expected closures for functions and expressions. These arise derivationally as semantic counterparts of the code that reduction rules create, as will become apparent once we see the evaluator. To

<sup>8</sup>The precise implementation of the environments depends on the representation of variables, which we abstract from in this presentation.

$SVal \ni v ::= x \mid f$	(semantic values)
$SFun \ni f ::= \langle x, e, \rho \rangle \mid v- \mid -v \mid H(f, v, v')$	(semantic functions)
$\rho \in SEnv \triangleq SVal^{Var}$	(environments)
$SComp \ni c ::= \langle e, \rho \rangle \mid v v'$	(closures/semantic computations)
$SCont \ni E_n ::= \square_{(n=0)} \mid E_n[\text{let } f]_{(n=0)} \mid E_n[\text{lift}]_{(n=0)} \mid E_m[\text{handle } v_h v_r]_{(n=m+1)}$	(stacks)
$SAns \ni a ::= v \mid E_0[\text{do } v] \mid E_n[x v]$	(answers)
$Conf \ni C ::= \langle c \rangle_E \mid \langle v_1, v_2 \rangle_A \mid \langle f, v \rangle_F \mid \langle v \rangle_{CV} \mid \langle c, f \rangle_{CL} \mid \langle c \rangle_{CF} \mid \langle v \rangle_{CD}$	(configurations)

Fig. 11. Syntax instrumented for environment-based semantics.

$\llbracket - \rrbracket_v : SVal \rightarrow Val$	$\llbracket - \rrbracket_c : Conf \rightarrow Exp$
$\llbracket x \rrbracket_v = x$	$\llbracket \langle c \rangle_E \rrbracket_c = \llbracket c \rrbracket_c$
$\llbracket f \rrbracket_v = \lambda x. \llbracket f \rrbracket_f^x$	$\llbracket \langle v_1, v_2 \rangle_A \rrbracket_c = \llbracket v_1 \rrbracket_v \llbracket v_2 \rrbracket_v$
$\llbracket - \rrbracket_f^- : SFun \times Var \rightarrow Exp$	$\llbracket \langle f, v \rangle_F \rrbracket_c = \llbracket f \rrbracket_f^x [\llbracket v \rrbracket_v / x]$
$\llbracket \langle x, e, \rho \rangle \rrbracket_f^y = ((a \mapsto \llbracket \rho(a) \rrbracket_v) [x \mapsto y]) (e)$	$\llbracket \langle v \rangle_{CV} \rrbracket_c = \text{return } \llbracket v \rrbracket_v$
$\llbracket v- \rrbracket_f^x = \llbracket v \rrbracket_v x$	$\llbracket \langle c, f \rangle_{CL} \rrbracket_c = \text{let } x = \llbracket c \rrbracket_c \text{ in } \llbracket f \rrbracket_f^x$
$\llbracket -v \rrbracket_f^x = x \llbracket v \rrbracket_v$	$\llbracket \langle c \rangle_{CF} \rrbracket_c = \text{lift } \llbracket c \rrbracket_c$
$\llbracket H(f, v, v') \rrbracket_f^x = \text{handle } \llbracket f \rrbracket_f^x \text{ with } \llbracket v \rrbracket_v, \llbracket v' \rrbracket_v$	$\llbracket \langle v \rangle_{CD} \rrbracket_c = \text{do } \llbracket v \rrbracket_v$
$\llbracket - \rrbracket_c : SComp \rightarrow Exp$	$V : Val \times SEnv \rightarrow SVal$
$\llbracket \langle e, \rho \rangle \rrbracket_c = (a \mapsto \llbracket \rho(a) \rrbracket_v) (e)$	$V(x, \rho) = \rho(x)$
$\llbracket v v' \rrbracket_c = \llbracket v \rrbracket_v \llbracket v' \rrbracket_v$	$V(\lambda x. e, \rho) = \langle x, e, \rho \rangle$

Fig. 12. Erasure of semantic components of the evaluator and the interpretation of values.

provide a better intuition for the interpretation of these constructs (and since they will be useful in the following) we present erasure functions for semantic values, functions, computations and configurations that map these into values or expressions, as appropriate. These can be naturally extended to the stacks (which share the structure with evaluation contexts) and answers.

*Evaluation and Read-Back.* With the necessary semantic components defined, we now turn to the definition of the evaluation and read-back functions. Since both of these are partial, we present them via their graphs, as inductively defined predicates that are functional. In this, we follow some of the spirit of Abel’s presentation of untyped normalisation-by-evaluation via partial applicative structures [Abel 2013], although due to the call-by-value nature of the reduction theory and the presence of control operators the spaces defined above do not form partial applicative structures as such.

The evaluation, presented in Figure 13, is an abstract-machine-style big-step semantics, reminiscent of Harper *et al.*’s account of call/cc in ML [Harper *et al.* 1993].<sup>9</sup> The evaluation proceeds

<sup>9</sup>Defunctionalizing a higher-order evaluator in continuation-passing style [Reynolds 1998a], expressed in a suitable meta-language, e.g., a domain theory meta-language [Reynolds 1998b] or an ML family representative [Ager *et al.* 2003], we

$$\begin{array}{c}
\frac{\langle V(v, \rho) \rangle_{CV}, E \searrow a}{\langle \langle \text{return } v, \rho \rangle \rangle_E, E \searrow a} \quad \frac{\langle V(v_1, \rho), V(v_2, \rho) \rangle_A, E \searrow a}{\langle \langle v_1 v_2, \rho \rangle \rangle_E, E \searrow a} \quad \frac{\langle \langle a, \rho \rangle, \langle x, e, \rho \rangle \rangle_{CL}, E \searrow a}{\langle \langle \text{let } x = a \text{ in } e, \rho \rangle \rangle_E, E \searrow a} \\
\frac{\langle \langle e, \rho \rangle \rangle_{CF}, E \searrow a}{\langle \langle \text{lift } e, \rho \rangle \rangle_E, E \searrow a} \quad \frac{\langle V(v, \rho) \rangle_{CD}, E \searrow a}{\langle \langle \text{do } v, \rho \rangle \rangle_E, E \searrow a} \quad \frac{\langle \langle e, \rho \rangle \rangle_E, E[\text{handle } V(h, \rho) V(r, \rho)] \searrow a}{\langle \langle \text{handle } e \text{ with } h, r, \rho \rangle \rangle_E, E \searrow a} \\
\frac{\langle v_1, v_2 \rangle_A, E \searrow a}{\langle v_1 v_2 \rangle_E, E \searrow a} \quad \frac{}{\langle x, v \rangle_A, E \searrow E[xv]} \quad \frac{\langle f, v \rangle_F, E \searrow a}{\langle f, v \rangle_A, E \searrow a} \\
\frac{\langle \langle e, \rho[x \mapsto v] \rangle \rangle_E, E \searrow a}{\langle \langle x, e, \rho \rangle, v \rangle_F, E \searrow a} \quad \frac{\langle f, v \rangle_F, E[\text{handle } v_1 v_2] \searrow a}{\langle H(f, v_1, v_2), v \rangle_F, E \searrow a} \quad \frac{\langle v_1, v_2 \rangle_A, E \searrow a}{\langle v_1 -, v_2 \rangle_F, E \searrow a} \\
\frac{\langle v_1, v_2 \rangle_A, E \searrow a}{\langle -v_2, v_1 \rangle_F, E \searrow a} \quad \frac{}{\langle v \rangle_{CD}, E_0 \searrow E_0[\text{do } v]} \quad \frac{\langle v_h v, -v_r \rangle_{CL}, E \searrow a}{\langle v \rangle_{CD}, E[\text{handle } v_h v_r] \searrow a} \\
\frac{\langle c \rangle_E, E_0[\text{let } f] \searrow a}{\langle c, f \rangle_{CL}, E_0 \searrow a} \quad \frac{\langle c \rangle_E, E[\text{handle } v_h H(f, v_h, v_r)] \searrow a}{\langle c, f \rangle_{CL}, E[\text{handle } v_h v_r] \searrow a} \\
\frac{\langle c \rangle_E, E_0[\text{lift}] \searrow a}{\langle c \rangle_{CF}, E_0 \searrow a} \quad \frac{\langle c, v_r - \rangle_{CL}, E \searrow a}{\langle c \rangle_{CF}, E[\text{handle } v_h v_r] \searrow a} \\
\frac{}{\langle v \rangle_{CV}, \square \searrow v} \quad \frac{\langle f, v \rangle_F, E \searrow a}{\langle v \rangle_{CV}, E[\text{let } f] \searrow a} \quad \frac{\langle v \rangle_{CV}, E \searrow a}{\langle v \rangle_{CV}, E[\text{lift}] \searrow a} \quad \frac{\langle v_r, v \rangle_A, E \searrow a}{\langle v \rangle_{CV}, E[\text{handle } v_h v_r] \searrow a}
\end{array}$$

Fig. 13. The evaluation judgment. The depth indices on stacks are omitted where unconstrained or clear from context.

by inspecting the configuration and – in the case of the four “continue” configurations – the stack. The eval configuration (denoted by E) inspects the closure, decomposes the expression and chooses the appropriate continuation or application configuration. The apply configuration (A) either recognises that the computation is environment-stuck, producing an appropriate answer, or continues with the application of the function. This transition is the analogue of the  $(\beta.v)$  reduction of the reduction theory. The function-application configuration (F) proceeds depending on the shape of the function: if it is a closure, the argument is added to the environment, and the evaluation proceeds in the usual fashion. The two “application” functions apply one of the arguments to the other, according to which one was missing. Finally, the H function reconstructs the larger context within which its first argument can now be applied to the given value.

This leaves us with the continuation configurations. Three of these, for do (denoted with CD), let (CL) and lift (CF) only perform non-trivial computation if the context contains a handler. If it does not, a do expression is control-stuck, and thus returns the answer, while the other two

---

obtain a first-order tail-recursive evaluator that lends itself to two equivalent interpretations: either as an abstract machine whose transition function is defined according to the tail-calls of the first-order evaluator, or, as in this work, as a stylized single-premised big-step semantics in which derivation rules, read bottom-up, correspond to the tail-calls of the first-order evaluator.

simply extend the stack and return to the eval configuration. If the stack does contain a handler, however, the computation performs the analogues of  $(\text{DH.do})$ ,  $(\text{DH.let})$  and  $(\text{DH.lift})$  reductions, respectively. Note that these three transitions are responsible for all the non-standard functions and computations of the machine: the reduction of the let expression creates the H function as the modified resumption of the handler, the reduction of the lift sequences the computation through what is in effect an  $\eta$ -expansion of the resumption, while the complex result of the reduction of do is responsible for the final two constructs. Finally, the remaining configuration (CV) represents continuation of the evaluation when a value has been encountered: this can either lead to the final result of the computation when the stack is empty, or the remaining three contractions:  $(\text{lift.v})$ ,  $(\beta.\text{let})$  or  $(\text{DH.v})$ . Note that the relation is obviously functional in its two arguments, thus representing a graph of a partial function.

Before turning to the read-back function and the process of full normalisation, we remark that the evaluation judgment defined above corresponds with the weak-head reduction defined in Section 5.

**LEMMA 7.2.** *If any configuration  $C$  with a stack  $E$  evaluates to an answer  $a$ , i.e.,  $C, E \searrow a$ , then its erasure reduces to the erasure of the answer:  $[E][[C]] \rightarrow^w [a]$ .*

**LEMMA 7.3.** *For any two configurations  $C_1, C_2$  such that  $[C_1] \rightarrow^w [C_2]$  and two stacks  $E_1 \sim_E E_2$  the two evaluations are equivalent as partial functions, i.e.,*

- for any  $a_1$ , if  $C_1, E_1 \searrow a_1$ , then  $C_2, E_2 \searrow a_2$  for some  $a_2$  such that  $a_1 \sim_a a_2$ ;
- for any  $a_2$ , if  $C_2, E_2 \searrow a_2$ , then  $C_1, E_1 \searrow a_1$  for some  $a_1$  such that  $a_1 \sim_a a_2$ .

*The similarity relations  $\sim_E$  and  $\sim_a$  are given as appropriate congruences over equality of erasures of semantic values and functions.*

Since any expression that cannot progress in terms of weak-head reduction is an answer, the latter lemma leads to a simple corollary that any reduction to a weak-head normal form can be computed (up to erasure) by the evaluator.

**COROLLARY 7.4.** *For any expressions  $e_1, e_2$  such that  $e_1 \rightarrow^w e_2$  and  $e_2 \not\rightarrow^w$ , there exists an answer  $a$  such that  $\langle\langle e_1, x \mapsto x \rangle\rangle_E, \square \searrow a$  and  $[a] = e_2$ .*

More important, from the normalisation perspective, is the relationship of evaluation to general reduction. Instead of tackling this problem directly, we utilise the standardisation theorem, which helps us order the general reduction sequence. In order to state the lemmas, we first define the notions of reduction for the components of the evaluator by taking  $v \rightarrow_v v' \triangleq [v] \rightarrow^i [v']$  and  $f \rightarrow_f f' \triangleq [f]^x \rightarrow^s [f']^x$ , and defining the notions for stacks and answers (as  $\rightarrow_E$  and  $\rightarrow_a$  respectively) as a congruential closure over the constructors of the notions for values and functions. This lets us state the properties that relate reduction and evaluation.

**LEMMA 7.5.** *For any configurations  $C_1, C_2$  and stacks  $E_1, E_2$  such that  $[E_1][[C_1]] \rightarrow^s [E_2][[C_2]]$ :*

- (1) *if  $C_1, E_1 \searrow a_1$  for some answer  $a_1$ , there exists  $a_2$  such that  $C_2, E_2 \searrow a_2$  and  $a_1 \rightarrow_a a_2$ ;*
- (2) *if  $C_2, E_2 \searrow a_2$  for some answer  $a_2$ , there exists  $a_1$  such that  $C_1, E_1 \searrow a_1$  and  $a_1 \rightarrow_a a_2$ .*

With the evaluation aspect defined, we now turn to the read-back function, whose graph is defined inductively in Figure 14. It consists of four mutually inductive predicates: reading back answers as normal forms, semantic values as normal values, semantic functions (with a variable name as a second parameter) as normal forms, and  $n$ -deep stacks, together with  $n$ -deep neutral computations as 0-deep neutral computations. Most of the process proceeds structurally, with all the recursive evaluation in the internal structure delegated to reading back of functions – which, as expected, proceeds by applying the function to the fresh variable (chosen to match the appropriate



$$\begin{array}{c}
\frac{RB_v v \searrow \xi}{RB_a v \searrow \xi} \quad \frac{RB_v v \searrow \xi \quad RB_E^n E[x \xi] \searrow v}{RB_a E_n[x v] \searrow v} \quad \frac{RB_v v \searrow \xi \quad RB_E^0 E[\text{do } \xi] \searrow v}{RB_a E_0[\text{do } v] \searrow v} \\
\\
\frac{}{RB_v x \searrow x} \quad \frac{RB_f f^x \searrow \mu}{RB_v f \searrow \lambda x . \mu} \quad \frac{\langle f, x \rangle_F, \square \searrow a \quad RB_a a \searrow \mu}{RB_f f^x \searrow \mu} \\
\\
\frac{}{RB_E^0 \square[v] \searrow v} \quad \frac{RB_f f^x \searrow \mu \quad RB_E^0 E[\text{let } x = v \text{ in } \mu] \searrow v'}{RB_E^0 E[\text{let } f][v] \searrow v'} \quad \frac{RB_E^0 E[\text{lift } v] \searrow v'}{RB_E^0 E[\text{lift}][v] \searrow v'} \\
\\
\frac{RB_v v \searrow \xi \quad RB_v v' \searrow \xi' \quad RB_E^n E[\text{handle } v \text{ with } \xi, \xi'] \searrow v'}{RB_E^{n+1} E[\text{handle } v v'][v] \searrow v'}
\end{array}$$

Fig. 14. The read-back judgment.

binder in the resulting normal form). As expected, the following properties are rather simple to obtain, given the lemmas about evaluation above.

**LEMMA 7.6.** *The read-back process only performs internal reductions, i.e., for any answer  $a$  and normal form  $\mu$ , if  $RB_a a \searrow \mu$ , then  $[a] \rightarrow^1 \mu$ .*

**LEMMA 7.7.** *The read-back process is closed under reductions of answers, i.e., for any answers  $a_1, a_2$  such that  $a_1 \rightarrow_a a_2$*

- if  $RB_a a_1 \searrow \mu$  for some  $\mu$ , then  $RB_a a_2 \searrow \mu$ ;
- if  $RB_a a_2 \searrow \mu$  for some  $\mu$ , then  $RB_a a_1 \searrow \mu$ .

The proofs proceed by induction on the structure of the derivation of read-back, mutually with analogous lemmas for values, functions and contexts, and utilise appropriate lemmas about evaluation in the case for functions. Together with the standardisation theorem and the lemmas about behaviour of evaluation with respect to standard reduction, we obtain soundness and completeness of the normalisation process.

**THEOREM 7.8 (SOUNDNESS AND COMPLETENESS OF NBE).** *The (partial) normalisation function, defined as*

$$\text{norm}(e) \searrow \mu \iff \exists a. \langle \langle e, x \mapsto x \rangle \rangle_E, \square \searrow a \wedge RB_a a \searrow \mu$$

is sound and complete, i.e.,

- for any  $e, \mu$ , if  $\text{norm}(e) \searrow \mu$  then  $e \rightarrow_{\lambda_{f\&DH}}^* \mu$ ;
- for any  $e_1, e_2$  if  $e_1 =_{\lambda_{f\&DH}} e_2$  and  $\text{norm}(e_1) \searrow \mu$  for some  $\mu$ , then  $\text{norm}(e_2) \searrow \mu$ .

## 8 RELATED WORK

### 8.1 Reduction Semantics for Effect Handlers

Operational semantics of effect handlers is typically presented as a small-step reduction relation that either takes advantage of explicitly represented evaluation contexts and captures resumptions at once, or it builds the resumption piece-by-piece with a helper construct that holds a prefix of the context to be captured in search of the matching handler. In each case, the resumption is built inside-out.

The former, more global approach can be found, e.g., in Leijen’s work on Koka [Leijen 2017b], in Hillerström et al.’s work on effect handlers in the context of the Links programming language [Hillerström et al. 2020], and in a series of papers coauthored by the first and third authors revolving around the Helium programming language [Biernacki et al. 2018, 2019; Piróg et al. 2019]. Such semantics is an adaptation of a classical context-sensitive reduction semantics for delimited control introduced by Felleisen et al. [Biernacka et al. 2005a] and directly implementable on a CEK-like abstract machine [Biernacka et al. 2005b; Dybvig et al. 2007; Felleisen and Friedman 1987].

The latter, more local approach can be found, e.g., in Bauer and Pretnar’s calculi modelling the Eff programming language [Bauer and Pretnar 2015; Pretnar 2015] or in Schuster et al.’s foundations of the Effekt programming language, based on the notion of capability [Schuster et al. 2022], where the continuation is captured by unwinding the stack of a CEK-like abstract machine. Again, such semantics is an adaptation of a classical operational, based on local reduction rules, semantics for calculi with control operators [Felleisen et al. 1986], first applied to delimited control by Felleisen [1988].

In principle, it is relatively easy to see the correspondence between the global and local semantics of a calculus and to define one given the other one. In both cases performing the effect operation initiates the resumption capture. The fine-grained operational semantics of Section 2 is based on an entirely different principle. Here, the handler plays the leading role and the semantics of the program hinges on the reduction rules that describe the interaction of the handler with other constructs. As a consequence, the resumption is built outside-in. Altogether, this new design makes it non-trivial to establish the correctness of the semantics but, at the same time, opens new possibilities for applications. Moreover, it does not require the runtime system to search or unwind the stack or to keep track of the freeness level for contexts, otherwise critical in the presence of lift [Biernacki et al. 2018, 2019; Piróg et al. 2019]. This form of operational semantics is already present in the authors’ work on the control operator `shift0` [Biernacki et al. 2021].

## 8.2 Rewriting Rules for Handlers as Code Optimisation

In [Pretnar 2015] Bauer and Pretnar present a set of basic equivalences for labeled handlers that are sound with respect to observational equivalence in their calculus. In his PhD dissertation [Saleh 2019], Saleh goes further and defines a set of optimisation rules for handlers that is divided into the simplification rules and handler-reduction rules. The former serve as administrative rules to expose handlers, whereas the latter aim at eliminating or distributing handlers. Three of them coincide with our rules (DH.v), (DH.do), and (DH.let), one is dedicated to a treatment of recursive functions, which we do not need to handle explicitly, and the last one deals with pure computations. The latter rule takes advantage of a type-and-effect system – we instead work in the untyped setting. These optimisation rules were then used by Karachalias et al. [2021] in their study of efficient compilation of effect handlers, again, heavily relying on a type-and-effect system (with subtyping).

None of the above works treats the optimisation rules as a reduction theory proper. In particular, no properties such as confluence or standardisation have been established before for the reduction rules that are the topic of the present work. Moreover, Saleh and Karachalias et al. only prove soundness of their reduction rules with respect to contextual equivalence, whereas we establish both their soundness and completeness with respect to the standard operational semantics.

## 8.3 Confluence and Standardisation Proofs

Our proofs of confluence and standardisation theorems follow the most recent Agda formalization of these results for the pure  $\lambda$ -calculus by McKinna, presented in his unpublished lecture notes [McKinna 2022]. The confluence proof is based on the technique of super-parallel reductions and complete super developments [Aczel 1978; McKinna 2022; McKinna and Pollack 1999;

[Takahashi 1995; van Raamsdonk 1993], which is a simplification of the classic proof method by parallel reductions due to Tait and Martin-Löf [Barendregt 1985]. The proof of the standardisation theorem [Mitschke 1979], in turn, is based on the methodology introduced by Takahashi [1995] and refined by McKinna and Pollack [McKinna 2022; McKinna and Pollack 1999]. McLaughlin’s proof of the Context Lemma for ELLA [McLaughlin 2020, Chapter 6], a core calculus of the Frank programming language [Convent et al. 2020], takes advantage of some of the techniques characteristic to such standardisation proofs.

#### 8.4 Normalization by Evaluation for Algebraic Effects

The only result that appears to directly take up the subject of a fully normalizing procedure for a calculus with algebraic effects is [Ahman and Staton 2013]. The normalization-by-evaluation algorithm presented in that work is rather distant from ours: it relies on a (simple) type system and the grammar of terms does not include effect handlers, just algebraic operations to be interpreted at the meta level, in a residualizing monad.

Normalization by evaluation for the untyped lambda calculus has been studied in [Abel 2013; Aehlig and Joachimski 2004; Filinski and Rohde 2005]. We followed Abel’s work that relies on partial applicative structures and big-step operational semantics [Kahn 1987]. These normalizers are in the functional correspondence [Ager et al. 2003] with fully normalizing abstract machines [Biernacka et al. 2020; García-Pérez and Nogueira 2014]. A read-back function was first presented as a key ingredient of a compiled fully normalizing abstract machine in [Grégoire and Leroy 2002].

### 9 CONCLUSION AND FUTURE WORK

We have introduced an operational semantics for deep effect handlers based on a notion of fine-grained reduction that describes the interaction of control delimiters (handlers) with other computation formers. Based on this notion of contraction, we have developed a reduction theory, with classic properties such as confluence and standardisation — some of the earliest such results for reduction theories of effect handlers and, to the best of our knowledge, the first explicitly stated. Due to the complexities of the reduction rules and the hybrid nature of the evaluation, certain technical innovations were introduced to the classic techniques used in these proofs. We believe that they may generalise to other theories that exhibit similar complexities.

To connect our semantics to the traditional, context-capturing notion of evaluation we have established a bisimilarity between the two; thus, since our reductions are more local and finer-grained, we provide a more powerful equational theory that is nonetheless compatible with the traditional evaluators. As a further application of our semantics to partial evaluation, we have developed the first formalised normalisation-by-evaluation procedure for an effect handler calculus, which can be used in program specialisation, or to establish convertibility of expressions involving handlers. These developments fill a significant gap in theoretical foundations of programming with effect handlers.

*Formalisation.* As noted in the introduction, the bulk of the results presented above has been formalised using the Coq interactive theorem prover. In order to aid clarity, we have used explicit variable names with a hygienic convention in play throughout the paper, where the formalisation uses a functorial, intrinsically well-scoped representation; this, however, is clearly an adequate choice. Except for this point, the presentation follows the formalisation closely, although there are cosmetic differences in the presentation of the evaluator in Section 7. We believe that our formalisation effort presents additional evidence for scalability of techniques such as super-parallel reductions (in Section 4) or parallel internal reductions (in Section 5) to complex reduction theories: while we had to extend those in certain ways to account for the particular critical pairs or the fact

that our weak-head reduction is a hybrid strategy, the fact that they *could* be naturally extended speaks to their robustness.

*Future work.* While significant theoretical properties of our calculus have been established, there are still major directions for future work. Firstly, the impact of the fine-grained semantics on the connection between effect handlers and pure lambda calculus, through a transformation to continuation-passing style, should be explored. Such a development would help ensure that the reduction rules are not just internally coherent, but tie in appropriately with a compilation process to a pure calculus. A separate area for investigation is the interplay between the fine-grained rules and shallow handlers: it is clear that in the current presentation the rules are not easily applicable to shallow handlers.<sup>10</sup> If adaptation of the fine-grained approach to shallow handlers proves impossible, characterising and explaining such a failure could provide a valuable insight into the nature of the two varieties of control operators. Finally, our normalisation algorithm could be used – in a statically-typed, terminating context – as a basis for convertibility checking, which could help guide integration of effect handlers (with type-and-effect-systems) and dependent types.

## DATA-AVAILABILITY STATEMENT

The formalisation associated with the paper, which covers all the main results except for Section 2.2, as well as an OCaml implementation of the normaliser described in Section 7 are available online [Sieczkowski et al. 2023]. For reproduction purposes, we note that the formalisation was checked with Coq v.8.16.1, while the implementation was developed using OCaml v.4.14.0.

## ACKNOWLEDGMENTS

We thank the anonymous reviewers for their valuable comments on the presentation of our results, and in particular on related work. Thanks are also due to James McKinna for valuable discussions, particularly on the standardisation theorem. This work was partially funded by the National Science Centre of Poland, under grant number 2019/33/B/ST6/00289.

## REFERENCES

- Andreas Abel. 2013. Normalization by Evaluation: Dependent Types and Impredicativity. Habilitation thesis.
- Peter Aczel. 1978. *A General Church-Rosser Theorem*. Technical Report. University of Manchester.
- Klaus Aehlig and Felix Joachimski. 2004. Operational aspects of untyped Normalisation by Evaluation. *Math. Struct. Comput. Sci.* 14, 4 (2004), 587–611. <https://doi.org/10.1017/S096012950400427X>
- Mads Sig Ager, Dariusz Biernacki, Olivier Danvy, and Jan Midtgaard. 2003. A functional correspondence between evaluators and abstract machines. In *Proceedings of the 5th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming, 27-29 August 2003, Uppsala, Sweden*. ACM, 8–19. <https://doi.org/10.1145/888251.888254>
- Danel Ahman and Sam Staton. 2013. Normalization by Evaluation and Algebraic Effects. In *Proceedings of the Twenty-ninth Conference on the Mathematical Foundations of Programming Semantics, MFPS 2013, New Orleans, LA, USA, June 23-25, 2013 (Electronic Notes in Theoretical Computer Science, Vol. 298)*, Dexter Kozen and Michael W. Mislove (Eds.). Elsevier, 51–69. <https://doi.org/10.1016/j.entcs.2013.09.007>
- Hendrik Pieter Barendregt. 1985. *The lambda calculus - its syntax and semantics*. Studies in logic and the foundations of mathematics, Vol. 103. North-Holland.
- Andrej Bauer and Matija Pretnar. 2015. Programming with algebraic effects and handlers. *J. Log. Algebr. Meth. Program.* 84, 1 (2015), 108–123. <https://doi.org/10.1016/j.jlamp.2014.02.001>
- Malgorzata Biernacka, Dariusz Biernacki, Witold Charatonik, and Tomasz Drab. 2020. An Abstract Machine for Strong Call by Value. In *Programming Languages and Systems - 18th Asian Symposium, APLAS 2020, Fukuoka, Japan, November 30 - December 2, 2020, Proceedings (Lecture Notes in Computer Science, Vol. 12470)*, Bruno C. d. S. Oliveira (Ed.). Springer, 147–166. [https://doi.org/10.1007/978-3-030-64437-6\\_8](https://doi.org/10.1007/978-3-030-64437-6_8)

<sup>10</sup>The same applies to delimited control operators, where shift0/reset admit fine-grained rules, while the status is unclear for control0/prompt.

- Malgorzata Biernacka, Dariusz Biernacki, and Olivier Danvy. 2005a. An Operational Foundation for Delimited Continuations in the CPS Hierarchy. *Log. Methods Comput. Sci.* 1, 2 (2005). [https://doi.org/10.2168/LMCS-1\(2:5\)2005](https://doi.org/10.2168/LMCS-1(2:5)2005)
- Malgorzata Biernacka, Dariusz Biernacki, and Olivier Danvy. 2005b. An Operational Foundation for Delimited Continuations in the CPS Hierarchy. *CoRR abs/cs/0508048* (2005). arXiv:cs/0508048 <http://arxiv.org/abs/cs/0508048>
- Malgorzata Biernacka, Witold Charatonik, and Klara Zielinska. 2017. Generalized Refocusing: From Hybrid Strategies to Abstract Machines. In *2nd International Conference on Formal Structures for Computation and Deduction, FSCD 2017, September 3-9, 2017, Oxford, UK (LIPIcs, Vol. 84)*, Dale Miller (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 10:1–10:17. <https://doi.org/10.4230/LIPIcs.FSCD.2017.10>
- Dariusz Biernacki, Sergueï Lenglet, and Piotr Polesiuk. 2020. A Complete Normal-Form Bisimilarity for Algebraic Effects and Handlers. In *5th International Conference on Formal Structures for Computation and Deduction, FSCD 2020, June 29-July 6, 2020, Paris, France (Virtual Conference) (LIPIcs, Vol. 167)*, Zena M. Ariola (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 7:1–7:22. <https://doi.org/10.4230/LIPIcs.FSCD.2020.7>
- Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2018. Handle with care: Relational interpretation of algebraic effects and handlers. *PACMPL* 2, POPL (2018), 8:1–8:30. <https://doi.org/10.1145/3158096>
- Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2019. Abstracting algebraic effects. *Proc. ACM Program. Lang.* 3, POPL (2019), 6:1–6:28. <https://doi.org/10.1145/3290319>
- Dariusz Biernacki, Mateusz Pyzik, and Filip Sieczkowski. 2021. Reflecting Stacked Continuations in a Fine-Grained Direct-Style Reduction Theory. In *PPDP 2021: 23rd International Symposium on Principles and Practice of Declarative Programming, Tallinn, Estonia, September 6-8, 2021*, Niccolò Veltri, Nick Benton, and Silvia Ghilezan (Eds.). ACM, 4:1–4:13. <https://doi.org/10.1145/3479394.3479399>
- Jonathan Immanuel Brachthäuser, Philipp Schuster, and Klaus Ostermann. 2020. Effects as Capabilities: Effect Handlers and Lightweight Effect Polymorphism. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 126 (nov 2020), 30 pages. <https://doi.org/10.1145/3428194>
- Lukas Convent, Sam Lindley, Conor McBride, and Craig McLaughlin. 2020. Doo bee doo bee doo. *J. Funct. Program.* 30 (2020), e9. <https://doi.org/10.1017/S0956796820000039>
- Olivier Danvy and Andrzej Filinski. 1990. Abstracting Control. In *Proceedings of the 1990 ACM Conference on LISP and Functional Programming, LFP 1990, Nice, France, 27-29 June 1990*, Gilles Kahn (Ed.). ACM, 151–160. <https://doi.org/10.1145/91556.91622>
- Olivier Danvy and Lasse R. Nielsen. 2004. Refocusing in Reduction Semantics. *BRICS Report Series* 11, 26 (Nov. 2004). <https://doi.org/10.7146/brics.v11i26.21851>
- R. Kent Dybvig, Simon L. Peyton Jones, and Amr Sabry. 2007. A monadic framework for delimited continuations. *J. Funct. Program.* 17, 6 (2007), 687–730. <https://doi.org/10.1017/S0956796807006259>
- Matthias Felleisen. 1988. The Theory and Practice of First-Class Prompts. In *Conference Record of the Fifteenth Annual ACM Symposium on Principles of Programming Languages, San Diego, California, USA, January 10-13, 1988*, Jeanne Ferrante and Peter Mager (Eds.). ACM Press, 180–190. <https://doi.org/10.1145/73560.73576>
- Matthias Felleisen. 1991. On the Expressive Power of Programming Languages. *Sci. Comput. Program.* 17, 1-3 (1991), 35–75. [https://doi.org/10.1016/0167-6423\(91\)90036-W](https://doi.org/10.1016/0167-6423(91)90036-W)
- Matthias Felleisen, Robert Bruce Findler, and Matthew Flatt. 2009. *Semantics Engineering with PLT Redex*. MIT Press. <http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=11885>
- Matthias Felleisen and Daniel P. Friedman. 1987. Control operators, the SECD-machine, and the  $\lambda$ -calculus. In *Formal Description of Programming Concepts - III: Proceedings of the IFIP TC 2/WG 2.2 Working Conference on Formal Description of Programming Concepts - III, Eberup, Denmark, 25-28 August 1986*, Martin Wirsing (Ed.). North-Holland, 193–222.
- Matthias Felleisen, Daniel P. Friedman, Eugene E. Kohlbecker, and Bruce F. Duba. 1986. Reasoning with Continuations. In *Proceedings of the Symposium on Logic in Computer Science (LICS '86), Cambridge, Massachusetts, USA, June 16-18, 1986*. IEEE Computer Society, 131–141.
- Andrzej Filinski and Henning Korsholm Rohde. 2005. Denotational aspects of untyped normalization by evaluation. *RAIRO Theor. Informatics Appl.* 39, 3 (2005), 423–453. <https://doi.org/10.1051/ita:2005026>
- Álvaro García-Pérez and Pablo Nogueira. 2014. On the syntactic and functional correspondence between hybrid (or layered) normalisers and abstract machines. *Sci. Comput. Program.* 95 (2014), 176–199. <https://doi.org/10.1016/j.scico.2014.05.011>
- Benjamin Grégoire and Xavier Leroy. 2002. A compiled implementation of strong reduction. In *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming (ICFP '02), Pittsburgh, Pennsylvania, USA, October 4-6, 2002*, Mitchell Wand and Simon L. Peyton Jones (Eds.). ACM, 235–246. <https://doi.org/10.1145/581478.581501>
- Robert Harper, Bruce F. Duba, and David B. MacQueen. 1993. Typing First-Class Continuations in ML. *J. Funct. Program.* 3, 4 (1993), 465–484. <https://doi.org/10.1017/S095679680000085X>
- Daniel Hillerström, Sam Lindley, and Robert Atkey. 2020. Effect handlers via generalised continuations. *J. Funct. Program.* 30 (2020), e5. <https://doi.org/10.1017/S0956796820000040>

- Gilles Kahn. 1987. Natural Semantics. In *STACS 87, 4th Annual Symposium on Theoretical Aspects of Computer Science, Passau, Germany, February 19-21, 1987, Proceedings (Lecture Notes in Computer Science, Vol. 247)*, Franz-Josef Brandenburg, Guy Vidal-Naquet, and Martin Wirsing (Eds.). Springer, 22–39. <https://doi.org/10.1007/BFb0039592>
- Georgios Karachalias, Filip Koprivec, Matija Pretnar, and Tom Schrijvers. 2021. Efficient compilation of algebraic effect handlers. *Proc. ACM Program. Lang.* 5, OOPSLA (2021), 1–28. <https://doi.org/10.1145/3485479>
- Oleg Kiselyov and Chung-chieh Shan. 2007. A Substructural Type System for Delimited Continuations. In *Typed Lambda Calculi and Applications, 8th International Conference, TLCA 2007, Paris, France, June 26-28, 2007, Proceedings (Lecture Notes in Computer Science, Vol. 4583)*, Simona Ronchi Della Rocca (Ed.). Springer, 223–239. [https://doi.org/10.1007/978-3-540-73228-0\\_17](https://doi.org/10.1007/978-3-540-73228-0_17)
- Daan Leijen. 2017a. Type directed compilation of row-typed algebraic effects. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 486–499. <https://doi.org/10.1145/3009837.3009872>
- Daan Leijen. 2017b. Type directed compilation of row-typed algebraic effects. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 486–499. <https://doi.org/10.1145/3009837>
- Paul Blain Levy, John Power, and Hayo Thielecke. 2003. Modelling environments in call-by-value programming languages. *Inf. Comput.* 185, 2 (2003), 182–210. [https://doi.org/10.1016/S0890-5401\(03\)00088-9](https://doi.org/10.1016/S0890-5401(03)00088-9)
- Marek Materzok. 2013. Axiomatizing Subtyped Delimited Continuations. In *Computer Science Logic 2013 (CSL 2013), CSL 2013, September 2-5, 2013, Torino, Italy (LIPIcs, Vol. 23)*, Simona Ronchi Della Rocca (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 521–539. <https://doi.org/10.4230/LIPIcs.CSL.2013.521>
- Marek Materzok and Dariusz Biernacki. 2011. Subtyping delimited continuations. In *Proceeding of the 16th ACM SIGPLAN international conference on Functional Programming, ICFP 2011, Tokyo, Japan, September 19-21, 2011*, Manuel M. T. Chakravarty, Zhenjiang Hu, and Olivier Danvy (Eds.). ACM, 81–93. <https://doi.org/10.1145/2034773.2034786>
- James McKinna. 2022. The Lambda Calculus, Formalised: the Church-Rosser and Standardisation Theorems, with Applications. <https://www.macs.hw.ac.uk/splv/splv-2022/> Unpublished Lecture Notes from the Scottish Programming Languages and Verification Summer School (SPLV 2022), Edinburgh.
- James McKinna and Robert Pollack. 1999. Some Lambda Calculus and Type Theory Formalized. *J. Autom. Reason.* 23, 3-4 (1999), 373–409. <https://doi.org/10.1023/A:1006294005493>
- Craig McLaughlin. 2020. *Relational reasoning for effects and handlers*. Ph. D. Dissertation. University of Edinburgh, UK. <https://doi.org/10.7488/era/537>
- Gerd Mitschke. 1979. The Standardization Theorem for  $\lambda$ -Calculus. *Math. Log. Q.* 25, 1-2 (1979), 29–31. <https://doi.org/10.1002/malq.19790250104>
- Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2019. Typed Equivalence of Effect Handlers and Delimited Control. In *4th International Conference on Formal Structures for Computation and Deduction, FSCD 2019, June 24-30, 2019, Dortmund, Germany (LIPIcs, Vol. 131)*, Herman Geuvers (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 30:1–30:16. <https://doi.org/10.4230/LIPIcs.FSCD.2019.30>
- Gordon D. Plotkin and A. John Power. 2004. Computational Effects and Operations: An Overview. *Electronic Notes in Theoretical Computer Science* 73 (2004), 149–163. <https://doi.org/10.1016/j.entcs.2004.08.008>
- Gordon D. Plotkin and Matija Pretnar. 2013. Handling Algebraic Effects. *Logical Methods in Computer Science* 9, 4:23 (2013), 1–36. [https://doi.org/10.2168/LMCS-9\(4:23\)2013](https://doi.org/10.2168/LMCS-9(4:23)2013)
- Matija Pretnar. 2015. An Introduction to Algebraic Effects and Handlers. Invited tutorial paper. In *The 31st Conference on the Mathematical Foundations of Programming Semantics, MFPS 2015, Nijmegen, The Netherlands, June 22-25, 2015 (Electronic Notes in Theoretical Computer Science, Vol. 319)*, Dan R. Ghica (Ed.). Elsevier, 19–35. <https://doi.org/10.1016/j.entcs.2015.12.003>
- John C. Reynolds. 1998a. Definitional Interpreters for Higher-Order Programming Languages. *Higher-Order and Symbolic Computation* 11, 4 (1998), 363–397. <https://doi.org/10.1023/A:1010027404223>
- John C. Reynolds. 1998b. *Theories of programming languages*. Cambridge University Press.
- Amr Hany Saleh. 2019. *Efficient Algebraic Effect Handlers*. Ph. D. Dissertation. Arenberg Doctoral School, Faculty of Engineering Science, KU Leuven, Belgium.
- Philipp Schuster, Jonathan Immanuel Brachthäuser, Marius Müller, and Klaus Ostermann. 2022. A typed continuation-passing translation for lexical effect handlers. In *PLDI '22: 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, San Diego, CA, USA, June 13 - 17, 2022*, Ranjit Jhala and Isil Dillig (Eds.). ACM, 566–579. <https://doi.org/10.1145/3519939.3523710>
- Filip Sieczkowski, Mateusz Pyzik, and Dariusz Biernacki. 2023. *A General Fine-Grained Reduction Theory for Effect Handlers: Formalisation*. <https://doi.org/10.5281/zenodo.7993545>
- Masako Takahashi. 1995. Parallel Reductions in lambda-Calculus. *Inf. Comput.* 118, 1 (1995), 120–127. <https://doi.org/10.1006/inco.1995.1057>

- Femke van Raamsdonk. 1993. Confluence and Superdevelopments. In *Rewriting Techniques and Applications, 5th International Conference, RTA-93, Montreal, Canada, June 16-18, 1993, Proceedings (Lecture Notes in Computer Science, Vol. 690)*, Claude Kirchner (Ed.). Springer, 168–182. [https://doi.org/10.1007/978-3-662-21551-7\\_14](https://doi.org/10.1007/978-3-662-21551-7_14)
- Yizhou Zhang and Andrew C. Myers. 2019. Abstraction-Safe Effect Handlers via Tunneling. *Proc. ACM Program. Lang.* 3, POPL, Article 5 (jan 2019), 29 pages. <https://doi.org/10.1145/3290318>

Received 2023-03-01; accepted 2023-06-27